# Information Resonance

Ulrike Malmendier, Ryan Oprea and Laura Veldkamp*

July 8, 2025

**Abstract**

Information resonance—the extra weight people give to messages from those that feel like them —shows up in trivial choices and macroeconomic patterns. A new lab experiment explores how resonance works, by suppressing trust, relevance or memory as mechanisms. It establishes that participants are significantly more likely to follow advice from a stranger who has a common belief, preference, or (to a lesser degree) demographic characteristic. This extra weight given to resonant advice permeates a variety of choices and characteristics, even when the characteristics are obviously irrelevant, random assignments undermine trust and the information is salient at the moment of decision. To examine the consequences of resonance, we show how to re-purpose a standard information diffusion model, replacing geographic proximity with characteristic similarity. The model explains why role models matter, how social media can erode authority, and why promotion rules that look neutral can still produce homogeneous leadership teams. Finally, field evidence connects micro biases to potential macro outcomes: young workers disproportionately enter occupations already staffed by ethnically similar incumbents, and disproportionately leave those occupations when ethnic elders experience negative labor market shocks. By weaving together micro behavior, theory tools and macro consequences, the paper uncovers new levers for educators, firms, and policymakers to meaningfully influence their constituencies.

Individuals and communities with different socioeconomic and demographic backgrounds make systematically different choices across key domains, including education, finance, healthcare, and housing. These differences manifest in major life decisions—whether to attend college, start a family, buy a home, invest in stocks, launch a business, or move to a higher-wage city. Research has consistently shown that individuals are more likely to trust and act upon information when it comes from people they identify with – someone who shares characteristics with the listener.[1] But is that information perceived to be more accurate or relevant? Is this trust or memory? Does resonance operate only along dimensions like geography, ethnicity, and gender, or does it extend further? What consequences does it have? This paper systematically examines the concept of

*Preliminary. Please do not cite or circulate. Experimental estimates subject to change. We thank Nikki Azerang, Rishi Chandran, Pietro Risso and Xiaobo Yu for excellent research assistance. Keywords: information, race, gender, homophily, social networks.

[1]For instance, Stolper and Walter (2019) found that clients were more likely to follow financial advice from advisors with whom they shared demographic characteristics. Educational choices diverge across communities, house price beliefs have a strong local component, and vaccine uptake depends on social context (Fogli and Guerrieri, 2019; Kindermann et al., 2021; Kuchler et al., 2020b).

information resonance — how it shapes choices, belief diffusion, and aggregate outcomes — while probing the boundaries of its influence.

Information resonance is not about information access, limited attention, memory or trust. Recipients might even be able to reproduce and recite the expert information in question, but it does not resonate with them because they do not identify with the person conveying the information. As a result, they put less weight on it in their decision-making. This is what we mean by *information resonance*.

This paper takes a multi-pronged approach to studying resonance. First, we use an experiment to isolate resonance and distinguish it from information relevance, memory and trust. The experiment teaches us that many, but not all, personal characteristics can trigger information resonance. The strongest triggers are not the usual demographic suspects. Then, we provide a simple framework to trace out the logical consequences of "information resonance." We connect resonance to the efficacy of role models, the ability of social media to undermine expertise, and a new form of workplace discrimination. Finally, we present suggestive evidence that resonance matters for aggregate labor market decisions, while highlighting the difficulty of cleanly identifying resonance from non-experimental economic data.

Section 2 presents our lab experiment, which explores what forms resonance takes and what social characteristics trigger it. One might wonder whether trust, information relevance or some memory might explain the results. With these questions in mind, we design an incentivized experiment where subjects need to choose between books they'd like to receive, movies they'd like to watch, charities they'd like to donate to and even small investments made in one of two funds, based on the advice of a stranger. The experiment is designed to shut down some channels that might trigger homophily to identify the extra information weighting we call resonance. We provide one piece of evidence about that stranger. The evidence could be their race or ethnicity. But it could also be an array of other characteristics like whether they are a morning or night person, prefer cats or dogs, or even drink Coke or Pepsi. Our experiment makes trust impossible by having a recommender give a recommendation, without knowing who will view it. The receiver of the recommendation knows only one fact about this person. To avoid the possibility that information from people like the participant is viewed as more relevant, the choices are constructed to be irrelevant to any of the personal attributes. To rule out similarity as a device that implants information in memory, we make the recommendation visible at the time when the decision needs to be made. Yet, resonance persists.

While the existing studies make clear that race and gender matter in decision making, are these the only dimensions along which resonance is strong? We find that resonance is strongly significant, across most characteristics and all of our choice domains. People prefer the option recommended by someone with characteristics similar to them, even when those characteristics have no plausible

connection to the choice at hand. The most resonant attributes were not demographics, but beliefs and preferences. For example, dog lovers are more likely to accept the recommendation of another dog lover than to accept the recommendation of someone with their race or gender.

We try to push resonance to its breaking point: How minute and irrelevant can the characteristic be, and still trigger resonance? Where resonance fails is with experiences that are neither chosen nor formative. For example, having divorced parents is not something one chooses and reflects little about the kind of person they are. It is an example of an experience that does not trigger resonance.

Knowing that resonance takes the form of additional weight on information from sources the characteristic similarity, we construct a model to trace out the logical consequences of this behavior. Section 4 proposes a simple mathematical representation of resonance. It uses a model of information diffusion, when the proximity of the characteristic vectors of the speaker and listener determines how much information "resonates," i. e., how much weight individuals put on an action-outcome pair they observe. Using standard tools of spatial information diffusion in a Bayesian learning model, social proximity serves as the distance metric and tilts the updating process.

The first result shows how role models are powerful because their experiences resonate with many others in their community. The effect is like being a central node in a social network, but without needing any direct social connection between the role model and the observer. Evidence of the power of role models comes from many domains, including finance, education and health.[2]

Next, we use this model to study how this resonance effect interacts with social media. We model the effect of social media as relaxing geographic constraints on information access. Before the advent of social media, someone could only observe their neighbors' experiences. Now, they can connect with anyone. This creates more information for all. However, it is particularly powerful for people in minorities or with extreme social characteristics. Before, they would have found only few geographically close people like them to learn from. They would have been forced to learn from more centrist members of their community, and such learning would have been slow. With the advent of social media, they can weight similar extreme people more, and put correspondingly less weight on the actions of their physical community or on the advice of outside experts.

[2]In the macro-finance realm, D'Acunto et al. (2021) show that women and Black men update their unemployment and inflation expectations in line with FOMC forecasts when the presence of a woman or Black man on the FOMC is made salient, but White women and Black men *both* respond to the presence of a White woman or a Black man. In the context of education, a plethora of research documents the effects of female and URM (underrepresented minority) "role models" in STEM fields. Having a Black teacher has a positive effect on the educational outcomes of black students (graduation rates and test scores). Furthermore, there are stronger effects of Black male teachers on Black male students and Black female teachers on Black female students (Gershenson et al., 2018), especially in STEM courses (Price, 2010; Carrell et al., 2010). Or, in the context of health decisions, Alsan et al. (2019) found that Black men were more willing to take up preventative care, especially more invasive treatments, after meeting with Black male doctors. Another significant demographic characteristic is political identification, with people feeling close to those with whom they share an ideology. Politically targeted provaccine messaging significantly increased vaccination rates (Larsen et al., 2022).

This insight entails a change in perspective, or relative to prior literature on the effect of social networks. While prior work typically aims to "control for" selection into specific networks, in order to distinguish selection and treatment effects, our model implies that the selection is front and center of their role. People join social networks *because* those spaces are populated by people with a high degree of similarity, whose opinions and experiences they are likely to value. Once social networks are formed, people place higher value on the experiences of those within their networks. Our theory puts makes selection the key ingredient of interest. It is the characteristics that create the community and drive the exchange of information. We can see the positive effects on members of marginalized minority groups. And we can see the negative effects on fostering extremism, which recent events, such as election fraud myths and vaccine misinformation, have brought to the forefront of public debate. Evidence from Kuchler et al. (2020a) and Bailey et al. (2018) illuminates various ways in which online communities influence beliefs and decisions, e.g., regarding homeownership.[3]

To connect this to measurable economic outcomes, Section 5 uses ACS data (2005–2020) to show that young workers' occupational choices are shaped by same-ethnicity and same-gender elders in their local labor market. If an ethnic group is overrepresented in a given occupation within a Public Use Microdata Area (PUMA), younger members of that group are significantly more likely to enter the same occupation.

To help interpret this result, we further find that the observed effect only holds for same-gender occupational choices; there is no positive influence of opposite-gender over-representation in an occupation. We also show that occupational layoff shocks influence the job choices of the younger generation, only if their ethnicity was overrepresented.

While these findings might suggest that job choices and layoffs among socially close community members resonate more. Despite all the control variables, sample segmenting and differencing, it is difficult to rule out alternative explanations. Correlated occupational choices among same ethnicity-gender-PUMA peers might reflect correlated payoffs. Alternatively, this could be about information access. People learn about the desirability or risk of an occupation from ethic peers and lack access to similar information about other occupations. While we choose specifications to address these concerns, the exercise highlights the need for the experimental evidence to convincingly establish resonance and the domains in which it applies.

Taken together, the paper paints a picture of information resonance - what it looks like formally in a model, what it looks like in aggregate data, and what it looks like in a laboratory setting designed to isolate its effects. Each setting requires a change of focus from information access, to make sense of what is going on. This change in focus, could greatly enhance the effectiveness of a broad array of economic and social policies. The conclusion explains how this simple change of

[3]Bailey et al. (2018) also noted significant sociodemographic similarities between Facebook friends implying that people select into homophilic social networks.

focus could also change the way we think of idea diffusion and the process of economic progress.

**Motivating evidence.** Prior literature documenting that people respond more strongly to others with whom they share demographic similarities comes from several strands of literature on medical, educational, and financial decision-making. In addition to the papers discussed above, several papers in finance point to the role of ethnicity and caste in financing decisions. For example, Hegde and Tumlinson (2014) show that VCs are more likely to invest in startups with coethnic founders, even when these startups are of lower quality. Similarly, Pool et al. (2015) find that mutual fund managers in the same neighborhood hold more similar portfolios, especially if they share the same ethnic background. In India, financial cooperatives are found to maintain caste divisions when forming "joint liability groups," in which members act as guarantors for each other's loans, indicating that people feel more trust and responsibility towards members of their own caste (Stuart, 2007).

In a similar phenomenon, abstract information or statistics tend to be weighted significantly less than information gathered from personal experiences or the experiences of others whom we care about, identify with or empathize with (cf. Malmendier, 2021a,b; Hertwig et al., 2018; Hertwig and Wulff, 2021).

Those findings build on a long literature on social learning starting with Bandura (1963). Social learning theories posit that (1) learning occurs when we observe behavior and its consequences (action-payoff pairs in our model), but that (2) it is not pure behavioral imitation (as posited by behaviorism à la Skinner (1938)); rather, the process is affected by the social context (Bandura, 1977). The social context, and how an individual processes the information, depends in turn on characteristics of both the observer and the actor (in the observed behavior or event). In particular, observers are biased towards learning from others they identify with and that are similar to themselves in some way. Similarity has been identified in terms of kinship, familiarity, gender, and language.[4] More broadly, work on how information is transmitted is related to a sociology literature led by Granovetter (1973). A good example of information transmission along ethnic channels arises in a classic case study of Vietnamese nail salons. Workers passed knowledge of how to enter and succeed in this business to others who shared their ethnicity (Federman et al., 2006). Our work proposes a particular modeling structure that provides a lens through which to interpret these studies.

The sociology literature has also anticipated our emphasis on selection effects. Selection effects are important to resonance because they imply that people seek out information from and inter-

[4]For kinship (mother) see Corriveau et al. (2009); for familiarity, or in-group effects, see Learmonth et al. (2005); Corriveau and Harris (2009); Shutts et al. (2009); Buttelmann et al. (2007); Seehagen and Herbert (2012); for gender, see Serbin et al. (2001); Frazier et al. (2012); Shutts et al. (2010); Taylor (2013); for language, see Shutts et al. (2009); Kinzler et al. (2011); Kinzler et al. (2007).

actions with those who are similar to them. The sociology literature has dubbed the tendency to connect with and structure social networks around sociodemographic similarities "homophily," which has been incorporated in both theoretical and empirical economic research (McPherson and Smith-Lovin, 1987; McPherson et al., 2001; Kossinets and Watts, 2009; Smith et al., 2014).

In the development sector, literature on information diffusion, particularly "seeding," has shown that community members are more effective than governments in spreading information (Alatas et al., 2016). Not only are they able to spread the information, but they can actually identify which members of their social network will diffuse the information most efficiently (Banerjee et al., 2019). Banerjee et al. (2019) dub these individuals "gossips" and note that while their ability to circulate information is largely due to their diffusion centrality, their position in the network is not at all based on leadership status, education, or even their physical location in the village. Moreover, the authors emphasize that diffusion centrality fails to fully explain why gossips are good at spreading information and note that there may be unobserved factors used to identify "gossips," including who is most listened to/trusted.[5]

# 1 Belief Updating: From Bayesian Rationality to Resonance

We begin with a minimal mathematical structure that represents what we mean by information resonance.

## 1.1 Bayesian Belief Updating

Consider a standard rational model of belief formation. Each agent $i$ faces uncertainty about the payoff $z_i$ from taking an action $a_i$. Prior to observing any data, the agent holds a normal prior belief:

$$z_i \sim N(\mu_z, \sigma_z^2).$$

The agent then observes the outcomes of other agents' actions. These observations come in the form $z_j + \epsilon_j$, where $\epsilon_j \sim$ iid $N(0, \sigma_\epsilon^2)$ is idiosyncratic noise. Conditional on the agent's information set $\mathcal{I}_i$, Bayesian updating yields the posterior expectation:

$$E[z_i \mid \mathcal{I}_i] = \alpha_i \mu_z + \sum_{j=1}^{N} \beta_{ij}(z_j + \epsilon_j), \tag{1}$$

where $\alpha_i$ is the weight the agent assigns to the prior, and $\beta_{ij}$ is the weight on the observed outcome from agent $j$. This formulation assumes agents process information rationally and symmetrically, assigning weight to others' experiences based only on informational content.

[5] Banerjee et al. (2019) find gossips to be incredibly effective at spreading important information. In their trial, villages that used gossip seeding to spread information about vaccinations had 27% higher immunization rates.

### 1.2 Resonant Belief Updating

We depart from this rational benchmark by introducing a behavioral twist: agents do not process all signals equally. Instead, they apply greater weight to signals from others who are perceived as more similar to themselves. We formalize this idea through a mechanism we call *resonance*.

Under resonant belief updating, the posterior becomes:

$$E[z_i \mid \mathcal{I}_i] = \alpha_i \mu_z + \overline{\omega}_i \sum_{j=1}^{N} \sum_{t'=1}^{t} \omega_{ij} \beta_{ij} (z_j + \epsilon_j), \tag{2}$$

where $\omega_{ij} \in [0, 1]$ captures the resonance between agent $i$ and agent $j$, and $\overline{\omega}_i$ is a normalization factor ensuring that the weighted signal term remains on the same scale as in the Bayesian case.

The resonance weight $\omega_{ij}$ depends on the perceived similarity between agents $i$ and $j$, defined by a vector of observable or perceived characteristics $\theta$. Formally, we specify:

$$\omega_{ij} = 1 - \Phi\left(\|\theta_i - \theta_j\|\right), \tag{3}$$

where $\|\theta_i - \theta_j\|$ is the Euclidean distance between agents' characteristic vectors, and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Signals from agents more dissimilar to $i$ receive less weight, with the down-weighting increasing smoothly in dissimilarity.

### 1.3 What Cognitive Process Is This?

This modification captures a realistic behavioral regularity: individuals are more likely to be influenced by the experiences of others who "resonate" with them—those who are similar in background, preferences, demographics, or other salient traits. By endogenizing the weight placed on others' signals, based on perceived similarity, the model departs from rational expectations updating.

A central feature of resonant belief updating is the notion that agents assign greater weight to signals from others who are similar to them along salient dimensions, captured by the vector $\theta$. But what constitutes similarity? $\theta$ may include demographic characteristics, such as gender, ethnicity, geographic location, or family status. Similarity in life experiences—such as educational background, cultural exposure, or participation in extracurricular activities like sports or music—can also strengthen resonance. Agents may additionally be drawn to others who share core beliefs, including political views, attitudes toward climate change, or interpretations of historical and social events. Even seemingly trivial preferences—such as a shared fondness for dogs over cats, or beaches over mountains—can serve as cues for perceived similarity. Finally, personality traits, such as optimism versus pessimism, extroversion versus introversion, or a tendency toward analytical versus intuitive thinking, may shape the extent to which one agent resonates with another. Each of these dimensions could constitute an entry in the $\theta$ vector, which in turn affects how much influence

others' experiences exert on one's beliefs. The next section explores which of these dimensions of similarity or difference, in fact, trigger resonance.

Modeling resonance as a weight does not fully explain why it arises. The resonance weights $\omega_{ij}$ could represent several cognitive mechanisms. First, resonance may operate through *memory*: people are more likely to remember information that resonates with their own characteristics or experiences. Second, it may reflect *trust*: individuals are more inclined to trust information coming from similar others. Third, resonance may affect *perceived precision*, whereby agents judge resonant information as more accurate or more relevant to their own circumstances. This could operate through a heuristic that is well-adapted to most situations. For many decisions, like what to eat, where to live, the advice of people like you is more relevant. Resonance could be extending that habit of thinking to domains where it should not rationally be applied. Alternatively, resonance may not affect beliefs per se, but rather *actions*: agents may continue to hold Bayesian beliefs about objective truths, yet allow resonant information to disproportionately influence their decisions—a form of action distortion. Another possibility is that resonance reflects *preferences*: individuals may derive utility from acting in ways that align with their social group or "tribe," taking comfort or satisfaction in following the advice or behaviors of similar others. Finally, the resonance mechanism may have a rational basis in some settings, where *similar characteristics genuinely increase the informational relevance* of observed outcomes.

To empirically examine these possibilities, we turn next to an experimental test of resonant updating. It will allow us to distinguish between some, but not all of these mechanisms. The experiment is designed to minimize the roles of memory, trust, objective precision and the part of preferences that involves acting similarly in a group. The possibilities that remain are perceived precision, direct influence over actions, of a preferences for following the advice of similar peers. All of these mechanisms can be modeled with the same formal representation. It is nearly impossible to prove whether weight changes are due to subjective beliefs or preferences because we know from theory that most preference formulations have an observationally equivalent representation with distorted beliefs, and vice-versa.

## 2 Experimental Evidence of Resonance (Ryan's Version)

In order to gather some direct evidence about resonance and learn something about the range, scope and nature of the phenomenon, we ran an experiment using an online subject pool. The goal of the experiment was to measure resonance in a setting free of some confounds that we expect in naturally occurring contexts and to minimize scope for some mechanisms that might generate resonance-like effects like memory limitations (i.e., people may remember information better from people like them). We also designed the experiment to minimize scope for resonance-like behavior

to occur for purely rational reasons (e.g., to minimize the extent to which it is rational to believe similar people are more capable or have systematically more relevant beliefs and preferences).

To do this, we design an experiment in which subjects make several different types of choices including which charity they'd prefer to donate to, which source of information (non-fiction book) they would like to receive, what source of entertainment they would like to experience (fiction book), what answer to a true/question they'd like to (effectively) bet on and what investment (ETF) they would like to link their earnings to. Before making each of these choices, we give subjects recommendations from a previous subject for each option, and tell the participant a single fact about – a personal characteristic of – that recommender. The characterists we reveal to subjects were drawn from a deliberately broad pool and included demographic characteristics, biographical details, personality traits, preferences and beliefs.

We use this design to measure resonance by examining whether subjects are more likely to choose options recommended by people who are like them (share characteristics with them) than options recommended by people who are unlike them.

### 2.1 Details of the Design

The experiment consists of two phases.

**Stage 1.** In Stage 1, we asked 250 subjects to express their preferences or beliefs in a series of 25 questions, presented in a random order. We included five types of questions.

- **Charities**. Choice between pairs of similar charities (e.g., WaterAid vs. Water.org).
- **Entertainiment Demand (Fiction Books)**. Choice between pairs of fiction books (e.g., The Monk vs. vs. Melmoth the Wanderer).
- **Information Demand (Non-Fiction Books)**. Choice between pairs of non-fiction books (e.g., The Joy of Less vs. Essentialism).
- **True/False Questions**. Choice between which of two difficult-to-research factual statements was true (e.g., "Among people who consider themselves vegetarians, the share who filed a tax return in the past 12 months is higher (lower) than 70%").[6]
- **ETFs**. Choice of which of two similar ETFs had a higher return in the past month (e.g., two healthcare funds or two real estate funds).[7]

We call these *Choice Questions* and the decision made in these questions *Choices*. For Charities, Entertainment Demand, Information Demand and ETF questions, Stage 1 subjects were given some

[6] A detailed description of how we constructed the True or False choice questions is given in Online Appendix A.
[7] A detailed description of how we constructed the ETFs choice questions is given in Online Appendix A.

detailed information on each of the two options they were choosing between.[8] Stage 1 subjects' choices were unincentivized but subjects were told that their answers to Choice Questions might be presented to future subjects as recommendations.

After making these 25 choices, we asked each Stage 1 subject to answer a series of 25 questions about themselves. We included five categories of questions.

- **Demographics**. E.g., White vs. Not white or Man vs. Woman.
- **Personality**. E.g., Extrovert vs. Introvert or Optimist vs. Pessimist.
- **Biographical**. E.g., Did (Did) not play a musical for at least 5 years growing up or From the North (South) of the US.
- **Preferences**. E.g., Democrat vs. Republican or Cat person vs. Dog person.
- **Beliefs**. E.g., The assassin of John F. Kennedy acted alone (acted with others) or AI will improve (worsen) society.

We call these *Characteristics Questions* and the answers to these questions *Characteristics*; a full list of the Characteristics Questions we used is listed in Online Appendix A.[9]

**Stage 2.** In Stage 2, we asked a separate group of 300 subjects to complete the same set of Choice Questions and Characteristics Questions but with three main differences.

First, Choice Questions in Stage 2 were incentivized. After the experiment, for 10% of randomly selected subjects we selected a random Choice Question for payment. If the selelected question was a charity, we donated $10 to the charity the subject selected while if it was a book we sent the subject and electronic version of the book they selected. If the question was a True/False question we paid the subject $5 if they made the correct choice and if the question awas an ETF question we paid the subject $5 if they selected the ETF with the higher returns last month.

Second, we gave Stage 2 subjects the same additional information on each Choice in the Choice Questions as in Stage 1, but subjects were required to click on a link to reveal this information. This allowed us to track what additional information subjects sought out to inform their choices in Stage 2.

Finally, and most importantly, for each of the two options in each Choice Question, we assigned a pair of Stage 1 subjects who gave contrasting recomendations and also gave contrasting answers to

[8]For the charities questions, we told participants what the sector the two charities operate in; for the books we told participants whether they were fiction or non-fiction and their genre; for the ETFs we always gave a description of what an ETF is and the sector the ETFs they cover.

[9]Subjects were allowed to choose "Does Not Apply" in these questions. Only about 5% of Stage 1 answers ($299/(25 \times 250)$) and 3% of Stage 2 answers ($250/(25 \times 300)$) were "Does Not Apply." Here 25 is the number of questions participants responded to and 250 and 300 are the respective sample sizes of the Stage 1 and Stage 2 experiments.

a randomly assigned Characteristics Question. For instance for the Choice Question asking subjects to choose between the charities Better Angels and The Bridge Alliance, we might have assigned the a Dog person recommender to Better Angels and a Cat person recommender to The Bridge Alliance. To make these assignments, for each subject we randomly matched Choice Questions to Characteristics Questions and randomly assigned answers to Choice Questions to answers to Characteristics Questions.[10]. We framed this information as a "fun fact" or piece of "trivia" about the recommenders.

**Implementation.** We ran the experiment on Prolific in May of 2025 using custom software programmed in Qualtrics.[11] A total of 550 subjects participated in the experiment (250 in Stage 1 and 300 in Stage 2). Experiments took, on average, 23 minutes and 12 seconds and subjects were paid $9.46/hr for their participants (in addition to additional rewards in Stage 2, discussed above). Experimental instructions and screenshots of the interface are provided in Online Appendix A.

## 2.2 Understanding the Design

We designed the experiment to serve several inferential goals.

First, we wanted to study a broad range of ways humans can be similar to one another to understand the scope of resonance effects and (relatedly) to study characteristics that both might be and are highly unlikely to contain much useful information on the reliability of recommenders. That is we wanted to vary the plausibility of resonance being rational. To do this, we studied an unusually broad range of Characteristics on which to study resonance, studying not only demographics (race, gender) but also opinions, biographical background and personality characteristics. Maybe more importantly we studied not only obviously "important" characteristics (like gender) that might relate to accuracy of judgements and similarity of preferences (making resonance possibly rational), but also seemingly superficial characteristics such as preferences between cats and dogs or Coke and Pepsi in which it becomes highly unlikely that resonance is a rational response.

Second, we wanted to maximize our chances of measuring latent resonance effects in the subject pool by minimizing orthogonal characteristics of choices that might interfere with measurement. To do this, we deliberately selected Choice Questions with relatively similar pairs of options. However, we also selected tasks that produced some meaningful variation in whether and to what degree participants could conduct online research to better distinguish between options. While some ETF options are very difficult to distinguish from one another and our True/False questions are hard to do research on, our charity questions and especially our fiction and non-fiction books questions

[10]We were able to do this without deception because we observed every combination of Characteristics and Choices in the Stage 1 data. Indeed in Online Appendix A, we show that there is little systematic evidence that Characteristics are strongly related to Choices in our data.

[11]We restricted our subject pool to participants located in the US who were fluent in English with an approval rate in the range of 90% or above. We recruited an even balance of males and females for the study.

can be researched online to lessen the ex ante similarity between the options. This give us some natural variation in the degree subjects had easy access to rational substitutes for resonance to inform their choices.

Third, we removed any systematic information in Stage 1 recommendations by uniformly randomizing the match of Characteristics to Choices. This, again, minimized the likelihood that any resonance effects we measure in the experiment are driven by rational considerations.

Finally, we minimized the likelihood that experimenter demand effects (i.e., subjects inferring the purpose of our study and making choices designed to please the experimenter) by only asking Characteristics Questions *after* subjects had made their selections in the Choice Questions, thereby largely cloaking the purpose of the experiment.

### 2.3 Results

When a Stage 2 subject makes a choice (in one of the Choice Questions) that was recommended by a Stage 1 subject with whom she shares a Characteristic, we say the choice was the *hpmophilic choice* (else the non-homophilic choice). To measure resonance, we examine whether subjects are more likely than random (i.e., more than 50% likely) to make homophilic choices. For instance, we examine to what degree subjects who are Democrats disproprtionately (i.e., at a rate greater than 50%) choose a book recommended by a Democrat (rather than the alternative book in the task, recommended by a Republican), and vice versa. To the degree subjects make systematically homophilic choices (i.e., to the degree rates of homophilic choice exceeds 0.5) we say we have evidence of resonance.

To measure this we run LPM regressions on dummy variables, one for each Characteristic Question $c$ giving us estimates $\beta_c$ for each Characteristic Question studied in the experiment (standard errors are clustered at the subject level). In order to avoid spurious evidence of resonance, we use weighted least squares to equally weight subjects with each Characteristic (e.g., to equally weight Republicans and Democrats etc.).[12] Our primary measure of resonance is the *resonance rate* $R_c = \beta_c - 0.5$. Postive values of this statistic are evidence of resonance and negative values evidence of anti-resonance. Given our design, a resonance rate of $R_c$ =0.1 means that subjects are 20 percentage points more likely to follow a recommendation by someone who shares characteristic $c$ with them than someone with the opposite characteristic. In odds term, this subject is 1.5 times

[12]We do this to correct for a potential statistical confound: if (i) *all* subjects (regardless of Characteristic) tend to perceive people with some given Characteristic as better recommenders and (ii) people with that Characteristic are also over-represented in the sample, calculating the rate at which subjects, on average, follow the advice of subjects like them can generate spurious evidence of resonance. For instance, if both coffee drinkers and non-coffee drinkers tend to agree that coffee drinkers give better recommendations, *and* coffee drinkers are dominant in the population, then subjects will on average disproportionately prefer members of their own group even though there is no systematic evidence of homophlic choices. To correct for this, we equally weight people with both Characteristics (i.e., coffee drinkers and non-coffee drinkers) for each Characteristics Question when calculating $\beta_c$.

more likely to follow a recommendation by someone who shares their characteristic than someone who has the opposite characteristic.

The solid dots in Figure 1 plot estimates of $R_c$ for each Characteristic Question, $c$, grouped according to our five categories of Characteristics Questions in our experiment. Green (red) x's plot the raw rates of homophilic choice for the more (less) common answer to each Characteristic Question, and the more common answer in each case is given at the bottom of the graph. We make several observations.[13]

**1. The overall rate of resonance is large**. First, we find significant evidence of resonance, overall, in our data. Averaging across Characteristics, we calculate a rate of 0.147 (SE = 0.006), meaning subjects are overall 1.831 times more likely to follow recommendations by subjects like them than unlike them.

**2. Resonance is widespread across types of homophily**. Second, resonance rates are not only large on average, they are also widespread across the many characteristics we study. We find evidence of resonance emerging in *most* of the Characteristics we study: in 23 out of 25 cases, we find positive evidence of resonance and we can reject the hypothesis of equality with zero in all cases at the 95 percent confidence level.

Notably, we find resonance not only in standard demographic categories like race and gender and in strong markers of identity like political party (Democrat vs. Republican), but also in categories that are not often considered as candidates for effects like these. For instance, political identity generates one of the highest rates of resonance in our sample, but we roughly equally strong resonance effects for soft drink preferences (Coke vs. Pepsi) and pet preferences (Cat vs. Dog). Opinions about history are nearly as resonant as political identity, preferences over things like vacations and coffee are as important as demographic categories like race and gender, as are similarities in relatively superficial aspects of biography. Indeed, in the left hand column of Table 1, we report resonance rates calculated at the Characteristic category level, and find that resonance effects are particularly strong for beliefs and preferences and somewhat weaker for demographics and biography (homophily in personality is weakest of all but still quite strong on average).

We think the breadth of contexts in which resonance arises is interesting in its own right, but we also think it is helpful for understanding the likely psychological mechanisms underlying resonance. One possible reason for resonance effects like ours is that subjects (quasi-)rationally believe the recommendations of people like them are of higher quality than subjects unlike them, either because subjects believe people like them are more capable or that people like them are more likely to have similar preferences. Our experimental design minimizes scope for this kind of effect ex ante because our resonance rates are calculated across many different choice problems for each

[13] In this and other exercises, we remove the 3% of answers to stage 2 Characteristics questions in which subjects declined to provide an answer (i.e., answered "Does not apply" to the question).

Characteristic: while it may be plausible that similarity on any given characteristic signals relevance for *some* Choice Questions, it is less plausible that this would be true across all of our (very different) Choice Questions. We believe the fact that resonance occurs not only for obviously "important" characteristics like gender, race or political identity but often with equal strength for characteristics lke soft drink, pet and vacation preferences casts further doubt on resonance stemming primarily from a rational inference. The fact that resonance doesn't seem systematically weaker for categories in which the informational value of homophily seems severely limited may suggest that resonance is not (or is not only) a rational inference about the relevance of recommenders' advice.

In Online Appendix A we provide additional evidence that resonance is unlikely to primarily represent a rational response to systematic information contained in the Characteristics we reveal to subjects about recommenders. In particular, we show that, overall, characteristics are not very predictive of answers to choice questions and, especially, that the degree of resonance is largely insenstive to the degree to which characteristics *are* predictive of answers of choice questions.

**3. Resonance is widespread across types of behavior**. Third, and relatedly, we find that resonance arises not only for many types of homophily, but also that it arises in all of the varied categories of Choice Questions included in our experiment. The right hand column of Table 1 calculates resonance rates for each of our five categories of Choice Questions and reveals that resonance rates are quite similar resonance rates across categories (rates are slightly lower in True/False questions than for the other four). Thus resonance has strong effects in a wide range of contexts from demand for information to demand for entertainment and from charity preferences to beliefs.

Again, while we think this breadth of application is interesting in its own right, we also think it is diagnostic of the scope and nature of resonance. In particular, our results suggest that reasonance may occur not only in cases in which subjects have limited bases for making a high quality choice, but also settings in which subjects had access to significant informational resources. Our categories of Choice Questions vary not only in the types of behavior they represent (e.g., demand for information, preferences for charities etc.) but also in the availability of additional information to differentiate the options available in the Choices. At one extreme, our ETF tasks do not list the actual name of the two ETFs being compared (in order to prevent subjects from looking up the answer to the question online), but instead gives subjects access to a range of prior monthly returns to compare the two ETFs. But these returns were often quite similar across choices leaving subjects with little basis for differentiating ETFs in some cases. Likewise, our True and False questions are designed to be difficult to research online. At the opposite extreme, the book options provided in our Information Demand and Entertainment Demand questions, while selected to be similar, have different titles and different description and subjects were free to go online and do further research on the books (Charities were similarly easy to research online).

Arguably, the fact that resonance rates are similar in tasks that likely induce frequent near-

indifference (i.e., ETFs and True or False questions) and tasks in which subjects can more easily form sharper preferences by seeking out further information suggests suggests that resonance may be a relatively robust phenomenon – one that may be relatively insensitive to the availability of resources to make rational choices. This suggests that resonance may occur even in settings in which decision makers have access to alternative resources (e.g., expert advice) to inform their choices.

In Online Appendix A we provide related evidence on the relationship between resonance effects and people's demand for information about their options in Choice Questions. In Stage 2, subjects were allowed to click a link to read additional information about each option for most of the Choice Questions (all but the True/False questions), giving us a measure of demand for information about choices. We find that resonance effects are present both for subjects who do and do not seek out additional information, but that resonance rates are somewhat smaller in the former case. We interpret this as evidence that at least to some extent subjects use homophily as a substitute for other types of information and that resonance therefore may be a boundedly rational way of economizing on the costs of information processing. This may help to explain our finding that rates of resonance are somewhat insensitive to the availability of additional information across our Choice categories.

### 2.4 Summary

To summarize, we find strong (and surprisingly robust) evidence of resonance across a variety of human characteristics and types of behavior. We find that resonance occurs strongly not only for obviously "important" characteristics like race, gender and political identity but also for seemingly superficial characteristics like soft drink preferences and preferences for cats versus dogs. This not only suggests that resonance may apply more widely than previously appreciated, but also that resonance may not represent (or at least may not entirely represent) a rational response to the information contained in the characteristics of others. We also find that resonance occurs across a number of distinct types of decisions – decisions that varied in the availability of further information to inform choice. This insensitivity to availability of information may suggest that resonance is a robust phenomenon that occurs not only in settings in which decision makers have weak bases for choice but also consequential settings in in which auxiliary information is available to inform decision-making.

## 3 Experimental Evidence of Resonance

In the previous section we provided some evidence that resonance (i) can be triggered by many distinct types of humans similarity (including some that are very superficial) and (ii) can influence

many types of behavior. We also find some evidence that resonance can occur for reasons that are not likely to be fully rational and can occur even in decisions in which information to guide rational choice is readily available. Given this, our experiment gives us reasons to believe resonance may be an important driver of behavior in many economic contexts and therefore an important friction to consider when modeling and predicting economic behavior. In this section, motivated by our experiment, we build a model of resonance and explore a small sampling of its many interesting economic implications. Our basic model is deliberately simple and technically familiar – a choice which we hope will facilitate easier and readier integration of resonance into economic modeling. To illustrate some of the (sometimes dramatic) implications of resonance, we highlight three implications of resonance that emphasize how resonant learning differs from more standard informational frictions. We start by exploring the power of role models. Next, we combines information access and resonant learning to explore their interaction. We show how the rise of social networks can undermine the influence of experts. Finally, we extend the model to show how it can be a foundation for labor market discrimination.

Previous papers have shown resonance effects arise in specific financial, educational or medical settings, in response to ethnic or gender similarity. But how does resonance work and how does it generalize? To explore resonance across a broad range of contexts and exclude some memory, trust or information-quality effects, we turn to a controlled experiment. The results that follow are preliminary and subject to change.

In designing our economic experiment, we ensured that several key methodological features were incorporated to enhance the validity and interpretability of our findings. First, subjects made payoff-relevant decisions, ensuring that we are observing economically meaningful choices rather than mere affective expressions of solidarity. Second, we eliminated systematic information in recommendations by fully randomizing the assignment of personal characteristics of recommenders to choices. This was achieved without deception, as all combinations of characteristics and choices occurred naturally in the data. Third, to mitigate demand effects and obscure the purpose of the experiment, subjects reported their own personal characteristics only after making their choices. Fourth, we introduced variation in the choice domains to generalize beyond a single type of decision, incorporating not only charitable giving but also the selection of information sources (non-fiction books), consumables (fiction books), belief-based true or false statements, and investment choices (ETFs). Fifth, we systematically varied the homophily variables, incorporating both demographic traits deemed highly consequential and more arbitrary personal connections (e.g., preference for cats or dogs). This allows us to disentangle whether observed effects stem from certain types of similarities or whether any similarity can trigger resonance. Finally, we go beyond simply measuring the rate at which subjects choose the homophilic option by disaggregating these choices based on the specific homophilic trait. This distinction is critical; for instance, if all subjects perceive

one political group as a better recommender, and that group is overrepresented in the sample, an apparent homophily effect could emerge spuriously. By incorporating these design features, we ensure that our experiment isolates the mechanism underlying resonance and its influence on decision-making.

### 3.1 Experimental Procedure

The experiment is a 'within-subject' design, i.e. every participant is exposed to every level of the independent variable or to every treatment. In each round, we ask participants to select one of two 'choice variables'. The choice variables pairs consist of the names of different charities, movies, and books. We also ask participants to answer a series of True or False statements and make predictions about the returns of a series of ETFs. We design these choices to be payoff-relevant. That is, we tell participants at the beginning of the survey that we will randomly select 10% of respondents to receive a payoff. For each selected participant, we then randomly select one of their choices: if it is a book, we actually send them the book they chose electronically; if it is a charity, we actually donate $10 to the charity they chose on their behalf; if it is a True or False statement *and* they chose the correct statement, we give them a $5 bonus; if it is an Exchange-Traded Fund *and* they correctly predicted which fund had the highest investment return last month, we also give them a $5 bonus.

Below each choice variable pair we display a recommendation of another survey participant. For each recommendation, we reveal a single social characteristic (called a 'fun fact' or 'trivia') of the recommender. We call these pairs of characteristics the 'treatment variables'. There are five types of treatment variables: demographic information (i.e. white v. not white), information about preferences (i.e. prefers working from home or at the office), biographical information (i.e. parents divorced during childhood v. did not), information about their personality (i.e. optimist v. pessimist), or information about beliefs (i.e. for nuclear power v. against it). After having recorded the choices of participants, we ask them to also select which treatment variable best describes them for each treatment variables pair.

The key idea is to examine whether stage 2 subjects tend to follow recommendations from stage 1 subjects who share personal characteristics with them. The resonance hypothesis is that subjects will tend to follow the recommendation of the subject with the common characteristic.

**Data, Sample Size, and Implementation** We implement the experiment in two stages. In stage one we ask a sample of participants to choose between choice variables for each choice variables pair and to pick which treatment variable best describes them for each treatment variables pair. The choice variables pairs in stage one are not recommended and do not have data about the recommenders. These choices are also not randomized. We then use the combination of choice

variables and treatment variables selections from stage one to display the trivia about the recommenders to stage two participants. Stage one is necessary, because IRB does not allow lying to participants in an experiment.

Stage two, which is the main part of the experiment, consists of a different and larger sample of participants. Participants choose a choice variable for each choice variables pair and then are also asked which treatment variable best describes them for each treatment variables pair. In stage two, as described above, the choices are recommended by a recommender and participants can seek additional information about each choice variable. In addition, the order of the choice variables pairs, the side of the screen of each choice variable within a choice variables pair, the match between each choice variables pair and treatment variables pair, and the match of each choice variable and treatment variable within each choice variables pair and treatment variables pair match are randomized.

To minimize noise, this experiment intentionally picks choice variable pairs whose individual items are fairly similar. For instance, for each pair of charities we make an effort so that their geographical scope and mission are similar (i.e. charities that work on water security internationally) and that have very similar names (i.e. Water.org v. WaterAid). Ultimately, participants can seek more information about each choice item. However, resonance rates might decrease in how different the two choice items are.

In the Appendix, Table 5 lists the full set of personal characteristics we explore, each assigned to one of five categories: beliefs, biographic, demographic, personality, or preferences. Table 6 lists the full set of choices we use. Table 7 list summary statistics with the number of observations from each type.

Both stage one and stage two were carried out using Qualtrics software and Prolific participants. Specifically, stage one had a sample size of 250 participants and stage two had a sample size of 300 participants. Participants needed to be American, fluent in English, with a Prolific approval rate (the percentage of a participant's from previous study submissions that are approved by researchers) in the range of 90-100%, and 50% male and 50% female to qualify for either stage one or stage two. In addition, participants who took part in stage one could not participate in stage two. Stage one was carried out on May 22, 2025 and stage two was carried out on May 23, 2025.

### 3.2 Experimental Results

A characteristic triggers information resonance if participants disproportionately pick a choice variable that is suggested by a recommender whose characteristic (treatment variable) is the same as their own.

Figure 1 illustrates the breadth of personal characteristics that trigger resonance. The dots above 0 represent resonance. On average, subjects with these characteristics follow the advice of

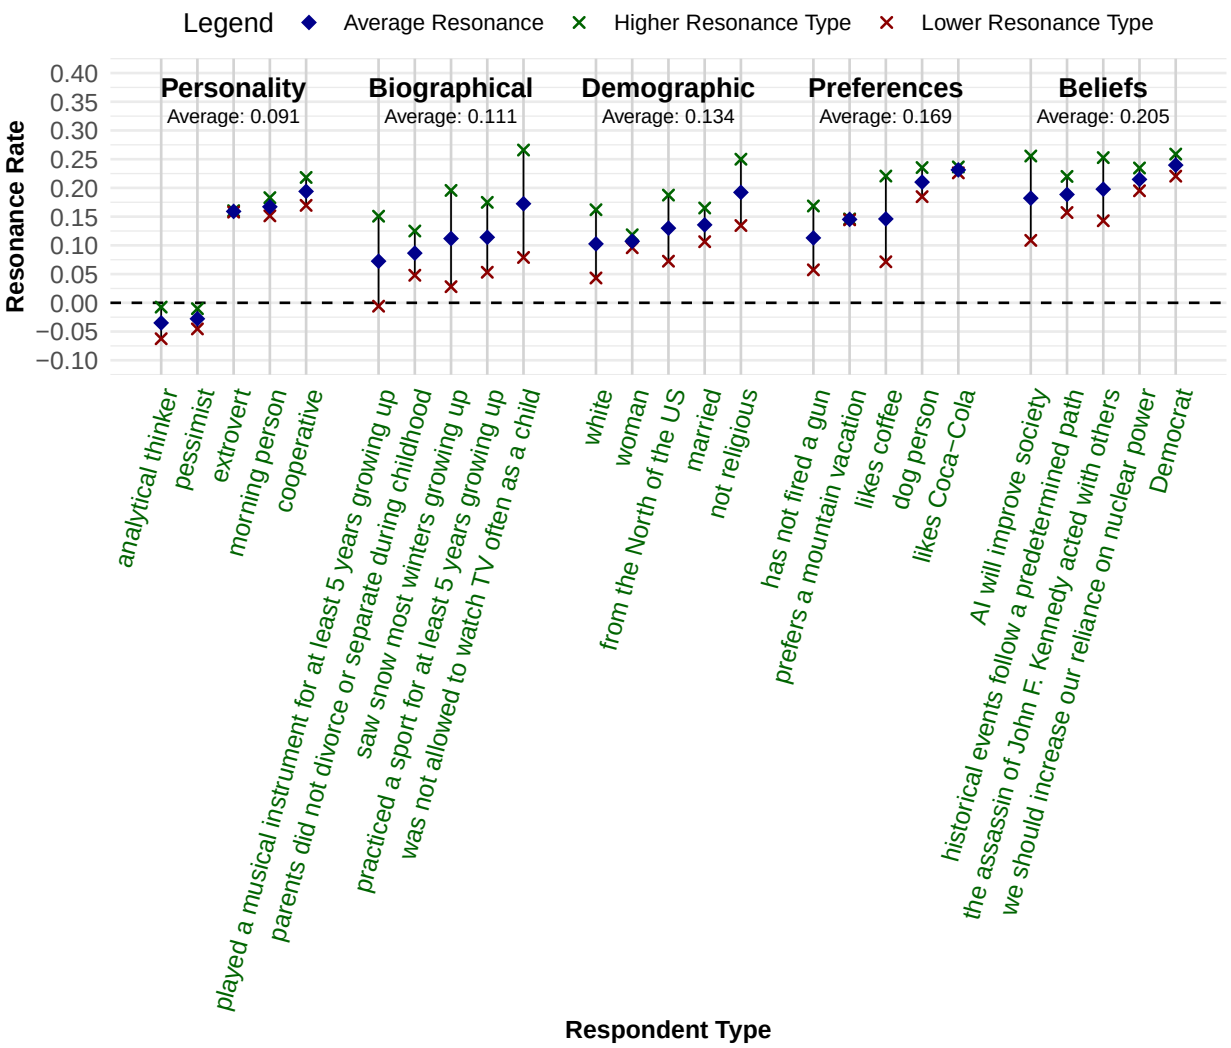



Figure 1: Resonance Estimates: Probability of following advice from others with similar characteristics.

Dots greater than 0 represent a greater probability of choosing the item recommended by another subject with similar personal characteristics. $x$'s are the probability of following own-type advice, estimated for each type separately. Personal characteristics listed below each point estimate.

others who share that characteristic more than others who do not. For example, dog lovers exhibit a higher likelihood of selecting the option recommended by another dog lover than of selecting the cat lover's choice, and vice-versa for those who prefer cats.

An estimated resonance rate of 0.1 means that a person is 10 pp more likely to choose a choice option that is recommended to them than if the choice was random. That means that if a recommender has this characteristic in common, the subject is 20 pp more likely to follow their recommendation than they would be to follow the recommendation of someone with the opposite characteristic.

The $x$'s above and below the dots are estimates of resonance, estimated separately for each type. In the animal preference example, the average effect is an average of cat-lovers response to the advice of other cat-lovers (lower x) and dog-lovers response to other dog-lovers' advice (higher x), where we weight each group to be equally representative.[14] For most characteristics, both types of respondents weight the advice of their own type more heavily. This teaches us that the results do not reflect one type being uniformly regarded are more reliable. One exception is respondents did not play a musical instrument growing up. Both respondents who played or did not play a musical instrument growing up are more likely to follow the advice of respondents who played a musical instrument growing up.

Table 1 shows the resonance rates for each category of personal characteristics (homophily domain) and each choice type. (Tables 8 and 9 in Appendix A report the full results for every single characteristic and for every single choice respectively.) The results reveal systematic differences in the strength of resonance across different categories of characteristics. The types of demographic characteristics that earlier studies emphasized are indeed resonant characteristics. However, beliefs and personality are the categories that give rise to the highest resonance rates. What triggers much less resonance is childhood experiences. People with similar childhood experiences are only slightly more likely to follow the recommendation of others with similar experiences.

While there is a marked difference in resonance rates across homophily domains, we see that the effects are more similar in magnitude across choice types. This shows that resonance is a far reaching and might be present in several aspects of decision making.

These results paint a nuanced picture of what the resonance weights $\omega_{ij}$ look like. We learn that weights on others' information is high when agents $i$ and $j$ share demographic characteristics. That weight is moderate when others share preference or personality characterisitics. The resonance weight is low when there are no similarities or the only similarity is a childhood experience. But overall, they strongly support the hypothesis that resonance exists and affects decisions. The set of factors that trigger resonance is larger than previous studies taught us. Resonance extends

[14] When we compute a weighted average that corrects for the fact that some characteristics are more common than others, we obtain similar results. The magnitude is not statistically different and the ordering of the characteristics that trigger the most resonance is unchanged. Results available on request.

Table 1: Resonance Rates by Homophily Category and Choice Type

| | *Homophily Choice* | | *Homophily Choice* |
|---|---|---|---|
| *Homophily Category* | (1) | *Choice Type* | (2) |
| Beliefs | 0.215 | Books (Fiction) | 0.155 |
| | (0.0126) | | (0.0125) |
| | [5.3e-64] | | [6.6e-35] |
| Biographical | 0.111 | Books (Non-Fiction) | 0.142 |
| | (0.0124) | | (0.0125) |
| | [2.6e-19] | | [1.2e-29] |
| Demographic | 0.137 | Charities | 0.157 |
| | (0.0124) | | (0.0126) |
| | [3.6e-28] | | [1.6e-35] |
| Personality | 0.091 | ETFs | 0.16 |
| | (0.0124) | | (0.0125) |
| | [3.2e-13] | | [5.1e-37] |
| Preferences | 0.182 | True or False | 0.12 |
| | (0.0126) | | (0.0126) |
| | [6.1e-47] | | [2.6e-21] |
| $R^2$ | 0.0944 | $R^2$ | 0.087 |
| Adj.$R^2$ | 0.0938 | Adj.$R^2$ | 0.0864 |
| Sample Size | 7275 | Sample Size | 7275 |

*Note*: OLS estimates of homophily choice (centerned at 0) on homophily categories and choice domains. Individual-level clustered standard errors (HC1) in round brackets. P-values in square brackets.

far beyond gender and race. Human beings find connections in all sorts of similarities. These connections shape the way we make decisions.

### 3.3 Interpreting the Experimental Results

The findings of this experiment suggest that the observed behavior cannot be attributed to conventional explanations such as information frictions, memory limitations, or trust-based considerations. While one could argue that any single pair of shared traits between individuals might be relevant to the decision at hand, it is highly implausible that the same rationale would apply consistently across all observed pairings. This suggests that the mechanism driving these results is not an issue of incomplete or asymmetric information. Furthermore, memory constraints do not play a role, as the relevant information is immediately accessible to participants, eliminating the need for retention or recall.

Likewise, the results do not appear to stem from trust. The recommendations were provided anonymously, with neither the recommender nor the recipient knowing one another. Since there was no potential for a conflict of interest, participants had no rational reason to distrust or devalue the advice based on concerns about the advisor's motivations. In the absence of these conventional explanations, the most plausible interpretations of the observed behavior center on preference or habit.

One possibility is that individuals exhibit an intrinsic preference for acting on information provided by those who share their characteristics. Alternatively, the behavior may reflect a heuristic—a rule of thumb developed from other contexts where information from similar individuals is indeed more relevant. In many real-world settings, shared characteristics can serve as useful signals for preference alignment, such as in product recommendations or lifestyle choices. However, in this experimental setting, those characteristics do not enhance the relevance of the information, yet they still influence decision-making. This suggests a systematic cognitive bias—one that, while often adaptive, can be misapplied in contexts where similarity does not equate to better information. Such a bias could have significant implications at the aggregate level, shaping societal decision-making processes in ways that are not necessarily optimal.

## 4 The Logical Consequences of Resonance

Motivated by our findings that resonance takes the form of greater weight on information from those with similar characteristics, we show how to model that phenomenon in an economic choice setting and trace out its dynamic effects. We merge the representation of resonant updating from Section 1 with information diffusion. The goal is not to innovate technically. The simplicity allows our mechanism to be easily incorporated into other settings. By showing how we can substitute

geographic proximity with type similarity, we can draw on what is already known about geographic diffusion to inform our thinking about how resonance translates to aggregate effects.

While we could derive many possible results from such a model, we choose three that highlight how resonant learning differs from a more standard information access friction. We start by exploring the power of role models. Next, we combines information access and resonant learning to explore their interaction. We show how the rise of social networks can undermine the influence of experts. Finally, we extend the model to show how it can be a foundation for labor market discrimination.

### 4.1 Modeling Resonance as Proximity

**Payoffs.** Agents make decisions between two possible actions: a known baseline action (e.g., maintaining the status quo) and a new action (e.g., adopting a new technology, changing occupations, or receiving a vaccine). The expected payoff from the new action is uncertain, with a persistent component (individual-specific) and a random transitory component.

Each agent $i \in \{1, \ldots, N\}$ chooses one of two actions $a \in \{0, 1\}$. The payoff to action $a = 0$ is known, and we normalize its payoff to be 0. In other words, the payoff to action 1 can be interpreted as its excess payoff over the payoff to action 0. The payoff to action 1 has two components, both of which are unknown to agents: There is a permanent component $z_i$ that agents can learn about and a transitory component $\epsilon_{it}$ is that is i.i.d. over time and thus unlearnable.

Let $z$ be the vector of all agents' payoffs, and $z_i$ its $i$th entry. Agents have priors $z \sim N(\mu_z, \Sigma_z)$, where $\Sigma_z$ is a variance-covariance matrix. If all agents have the same persistent payoff from an action, then $\Sigma_z$ should have every entry with an identical value. If agents have payoffs from action 1 that are entirely independent, and they cannot learn anything from others' experiences, this would be a diagonal $\Sigma_z$ matrix. Agents choose action $a_{it}$ to maximize their expected payoff:

$$U(a_{it}) = a_{it} E[z_a + \epsilon_{it} | \mathcal{I}_{it}]. \tag{4}$$

where $\epsilon_{it} \overset{iid}{\sim} N(0, \sigma_\epsilon^2)$ is the unlearnable random part of payoff, which is time- and individual-specific. Without the unlearnable payoff shock, one observation of an action-payoff pair would reveal exactly what the payoff $z_a$ is for a chosen action. The random payoff $\epsilon_{it}$ is what makes learning a gradual process.

**Information sets.** In order to compare our mechanism with a traditional information friction, we add variation in the access to information, as a feature that we can turn on or off. If agent $i$ takes the action $a_{it} = 1$ at time $t$, then each agent $j$ potentially sees their action and payoff outcome. Define an observation matrix $\Psi_t$, where the $(i, j)$th entry of $\Psi_t$ is 1 if $i$ observes $j$ and $j$ takes action $a_{jt} = 1$.

$\Psi_t$ might have a component that is not time-varying. For example, if $i$ sees only people in her neighborhood or only facebook friends, then $\Psi_t$ is the element-by-element product (a Hadamard product) of the social network matrix – a stable matrix of zero's and ones that indicates social connections – and the matrix $\mathbf{a}_t\mathbf{1}_n'$, where $\mathbf{a}_t$ is a $n \times 1$ vector of every agent's time $t$ action and $\mathbf{1}_n'$ is a $1 \times nh$ vector of 1's. Alternatively, we can consider people who are randomly visible. Or celebrities who are visible with a higher probability than others. This would be captured by a $\Psi_t$ matrix with a random component.

To express our solution, we need a bit of additional notation. Let $K_{it}$ be the number of $a = 1$ actions agent $i$ observes at time $t$. Then, $K_{it}$ is the sum of the entries in the $i$th row of $\Psi_t$. Then define $\psi_{it}$ to be a $K_{it} \times N$ matrix. Each row is a $N$-dimensional vector of zeroes, except for a one in $j$-th column, if $j$ takes action $a = 1$ and $i$ observes $j$'s payoff. This is a re-shaping of the $i$th row of the $\Psi_t$ matrix. This related matrix will be useful to express the model solution. We can express the set of signals that agent $i$ sees about the payoffs to action 1 as $\psi_{it}(z + \epsilon_t)$. This product is the $K_{it} \times 1$ vector containing a subset of the entries of the $N \times 1$ vector of all payoffs that agent $i$ gets to see.

The economic content of these assumptions is simply that everyone who observes agent $j$ take action $a_{jt} = 1$ observes an informative signal about $z_j$, which they can use to infer their own payoff component $z_i$.

**Belief updating** In this framework, $\omega_{ij}$ captures the extent to which agent $j$'s information resonates with agent $i$. A higher $\omega_{ij}$ means that $i$ assigns more weight to $j$'s experience when updating their beliefs, due to stronger perceived social proximity. For agent $i$ taking action $a = 1$, belief updating takes the form:

$$E[z_i|\mathcal{I}_{it}] = \alpha_{iz}\mu_z + \overline{\omega}_{iz}\sum_{j=1}^{N}\sum_{t'=1}^{t}\omega_{ij}\beta_{izj}(z_j + \epsilon_{jt'}), \tag{5}$$

where $\alpha_{iz}$ and $\beta_{izj}$ are the Bayesian weights on priors and signals:

$$\beta_{iz} = (\mathcal{Z}_{it}\Sigma_z\mathcal{Z}_{it}' + \sigma_\epsilon^2 I)^{-1}\mathcal{Z}_{it}\Sigma_z\mathbf{1}_i \;\; = \;\; \frac{Cov(z_i, z + \epsilon)}{Var(z + \epsilon)} \tag{6}$$

$$\alpha_{iz} = 1 - \sum_{j=1}^{N}\beta_{izj}. \tag{7}$$

The indicator variable $\mathbf{1}_i$ is a $N \times 1$ vector of zeros with 1 in row $i$, and $I$ is an identity matrix with size $K_{iz}$. The second equality in (6) shows that the Bayesian weight on signals is simply the

standard OLS estimator derived from linear models.

$\overline{\omega}_{iz}$ is a normalization factor such that:

$$\overline{\omega}_{iz} \sum_{j=1}^{N} \sum_{t'=1}^{t} \omega_{ij} \beta_{izj} = \sum_{j=1}^{N} \sum_{t'=1}^{t} \beta_{izj}. \tag{8}$$

This normalization makes sure that prior belief gets the same weight with or without resonance, and it also allows $\omega$ to be scale-neutral.

he $\omega$ term is the piece that distinguishes our theory from a standard diffusion model. It is the weight someone puts on the observed experience of another person.

In standard Bayesian learning models, the prior and each signal are typically weighted by their precision, the inverse of their variance, and divided by the sum of precisions, to get a weighted average. In our "neurologically realistic" Bayesian model, the conditional expectation is still a weighted average, but the agent adjusts the Bayesian weight for each signal by $\omega$. As we will specify below, this term allows us to account for agents giving less weight to information that comes from sources who are dissimilar. The fourth and final term in the expectation is the signal $z_a + \epsilon_{ajt'}$, the noisy information transmitted about the payoff to action $a$.

**Belief weights and characteristics.** For an observer $i$, different observations have different relevance. When $i$ observes $j$, the weight $i$ places on that observation, $\omega_{ij}$, depends on the distance between the characteristics $\theta_i$ of the observer and the characteristics $\theta_j$ of the person who took the action.[15] Then, if $i$ observes $j$'s actions,

$$\omega_{ij} = \underbrace{(2 - 2\Phi(\chi||\theta_i, \theta_j||))}_{\rho_{ij}} \cdot \sigma_\epsilon^{-2} \tag{9}$$

$$= \qquad \rho_{ij} \qquad \cdot \sigma_\epsilon^{-2}, \tag{10}$$

where $||\theta_i, \theta_j||$ is the Euclidean distance between agents' characteristics $\theta_i$ and $\theta_j$.[16] The multiplicative term $\chi$ allows us to consider different scenarios of agents weighting own versus other experiences more or less, as we discuss in more detail below. The Normal cumulative distribution function $\Phi$ equals 1/2 at 0, and hence the expression $2 - 2\Phi(\cdot)$ takes on a value of 1 at zero and

[15] What $w$ and $\theta$ should not capture is access to information. Some characteristics might, however, also determine access. For example, language, geographic location, and socio-economic status can determine the set of people they can observe. If this is the case, we would capture the informational friction via the information access matrix $\mathcal{Z}_{it}$, which would be related to the entries of $\theta_j$.

[16] One might think that agents should also learn from each others' choices, not just the outcomes. For example, if $i$ sees that $j$ did not choose action 1, $i$ might infer that $j$ has negative information about action 1. That sort of inference does not happen with pure resonance because resonance is not about asymmetric information. Everyone sees the same outcomes. They simply put different weights on them. When we also introduce information asymmetry, this becomes an issue, unless we define communities where if $i$ sees $j$, then $i$ also sees the same community outcomes that $j$ sees. Then, within observation groups, there is not information asymmetry and nothing to be inferred from choices.

then declines quickly towards zero. These terms define the resonance, or relevance, $\rho_{ij}$ that a signal coming from $j$ has for observer $i$.

We scale the resonance weight $\rho$ with the (true) accuracy of the signal, $\sigma_\epsilon^{-2}$, to reflect the idea that agents will not weight signals that have no information content at all. Thus, the overall weight $\omega_{ij}$ is the product of the resonance weight $\rho_{ij}$ and the precision weight $\sigma_\epsilon^{-2}$.

**Equilibrium** Everyone acts simultaneously in every period, knowing the outcomes of the previous periods. Agents maximize expected utility as specified in (4). They update their prior of the unknown payoffs $z_a$ for each action $a$ using Bayes' rule as spelled out in equations (5), with weights $\omega_{ij}$ as defined in (9)-(10).

Note that our agents are myopic. When they choose an action, they are not considering how much information that action will generate to enable better decision making for themselves or for others in the next period. The sort of active experimentation that we rule out here would greatly complicate the analysis, would obscure the main mechanism and would likely have minimal effect on results because of a free rider problem: When everyone can see actions of many others, the incentive is to let others incur the cost of experimentation and learn from their results.

**Embedding related learning mechanisms.** Our model of resonant information can nest other popular versions of information diffusion. At one end of the spectrum, our model allows for people place zero weight on anyone's experience other than their own, namely, when $\chi$ becomes very large. Only personal experiences in the sense of the payoff realizations resulting from own action choices affect beliefs (as in Malmendier and Nagel (2011, 2016)). At the other end of the spectrum, $\chi = 0$ represents the scenario where people learn from everyone's experience collectively, without discounting anyone's information. Theories of belief scarring show that when we learn from our collective experiences, transitory but extreme outcomes can have long-lived effects on aggregate outcomes (Kozlowski et al., 2020).

In between these polar cases, cities or metropolitan areas have beliefs that evolve differently because of their different experiences (Chodorow-Reich et al., 2021). If we think of the characteristic vector $0 < \chi < 1$ including one's location or metro area, the model becomes one of local learning from neighbors, as in (Fogli and Veldkamp, 2010; Burnside et al., 2016).[17]

Dynastic learning (Boerma and Karabarbounis, 2017) is also encapsulated by this framework. The characteristic space would include only whether one is a member of the same family, or not. If $\theta = 1$ for family members and $\theta = 0$ for all non-family members, then for $\chi$ very large, learning is dynastic.

[17]While the neighborhood-based cost functions of Hebert and Woodford (2020) sound similar, their focus is on how difficult it is to tell apart two states of the same world, that might lie in the same neighborhood, in event space. In contrast, we have a standard treatment of states of the world and focus instead on the characteristics of one's information sources.

Information resonance also helps to remedy a perpetual problem with Bayesian updating. In a standard Bayesian updating problem, a small amount of information will correct beliefs very quickly. Since all agents give full weight to all observations, a few realizations can induce quick convergence in posteriors. This is problematic because real people do not learn this efficiently. One advantage of weighting information by its resonance is to make information effectively scarce. People see lots of outcomes; but much of what they see does not resonate, and is therefore largely ignored. That slows changes in community behavior in a realistic way.

*A rational interpretation.* Our representation nests a rational one, where resonant information is simply more statistically relevant information. Characteristics could be loadings on payoff factors such that the distance metric $\omega_{ij}$ reflects the extent to which $i$'s outcomes covary with $j$'s. In that case, our model would capture rational econometric learning, with a known covariance structure of outcomes across the population. The weight on a piece of information about another agent's experience would then be covariance, divided by variance, which is a standard ordinary least squares regression coefficient. The belief formation rule (5) would become an ordinary least squares forecast, given all observed information. The variable $\theta_j$ can be a vector that includes any or all of these characteristics. However, our experiment highlighted settings where characteristics do not affect statistical relevance.

### 4.2 Parameter Choices and Procedure

We choose parameters, without any attempt to match realistic values. These are numerical examples designed not to measure any effect, but simply to illustrate the qualitative mechanics of this mechanism.

The parameters we use are as follows: Action payoffs are 0 for action $a = 0$ and $z_i = 4\ \forall i$, meaning that all agents have the same persistent component of payoffs to action $a = 1$. The individual transitory shocks to payoffs are very noisy, with $\epsilon_{ait} \sim N(0, 100)$ for both actions $a$, independently drawn for actions, agents, and time. Thus, in the absence of taxes or subsidies the total payoff to actions, net of action costs are: Action 0 has payoff $\sim N(0, 100)$. Action 1 has payoff $z_1 + \epsilon_{it} \sim N(4, 100)$. That is, action 1 is the better (higher expected payoff) action. If agents knew the true value of $z_1$, they would maximize their utility, as specified in (4), by choosing action 1. Prior beliefs for action 1 are $z_1 \sim N(-3, 10)$. So agents believe that action 0 is superior, when in truth, it is not. The action costs and subsidies are known to all agents. Agents update their beliefs according to (5).

The probability of agent $j$ observing agent $i$ at time $t$, $\Pr[\psi_{jit} = 1]$ is initially set to 0.2. In other words, all action-outcome pairs have a 20% chance of being observed by any given other agent. Finally, the distance between any two vectors $\theta_i$ and $\theta_j$ is transformed by parameter $\chi = 0.2$.

The simulations below differ only in their initial conditions, including the distribution of charac-

teristics (uniform or normal), the presence of taxes (or subsidies). What they all share in common is the following procedure:

1. For 400 agents, assign each node a geographic location $(x, y)$ on a 20 x 20 grid. Create a matrix where each node has a row and column that represent its physical location on a two-dimensional plane. This implies that the geographic distribution is uniform.

2. Draw social coordinates $(v, w)$ for each of these nodes. To achieve an even distribution of characteristics, Figures 1 and 2 draw $(v, w)$ from the two-dimensional integers $[0, 0]$, $[20, 20]$, without replacement.

3. Given these initial conditions, each agent chooses the action with higher expected payoff.

4. Draw random payoff shocks $(\epsilon_{it})$ for each chosen action.

5. Allow agents to observe the total payoff $z_j + \epsilon_{jt}$ for each $j \epsilon \tilde{\mathcal{Z}}_{it}$. Add this new data to each agent $i$'s information set and update each agent's beliefs, according to (5).

We repeat this procedure for 20 periods, and repeat the entire simulation, starting from the date-0 initial conditions, ten times. For each node at each date, we report (and illustrate graphically) the fraction of simulations in which that node chose action 1. This averaging of results allows to smooth out some of the randomness of a particular simulation.

More details on the simulation procedure for each set of results are in Appendix C.

### 4.3 Experts vs. Role Models

We now illustrate the basic insight emerging from our learning mechanism: the differing influences of "access to information," which is at the core of many traditional learning models, versus "social proximity" (information resonance) to the underlying realizations. We start by describing the model's characteristic space topology. The technical insight is that after re-casting the outcomes in this new topological space, resonance functions just like an information diffusion friction.

We denote someone whom many people feel socially close to a 'role model.' Their actions and experiences resonate with many people. That is, others put significant weight on and thus learn from the information provided by role models. We contrast the role model with a person about whose action and experiences many people are informed about, which we dub a 'celebrity.' The contrast between the influence of role models versus celebrities mirrors the difference between the effect of "social proximity" (information resonance) versus "access to information" that we would like to illustrate. Examples of such local influence on public health abound in the popular press, from the role of rabbis among ultra-orthodox Jews in Israel, Aboriginal health organizations in Australia, to local leaders in the rural areas in the U.S., Union leaders in Harlem, and descendants

of participants in the Tuskegee study in Alabama. More evidence of the power of role models is reported in Appendix D.

Because the leader's characteristic similarity leads others to put more weight on their experience, getting the one local leader to switch actions brings about faster learning and a faster change in actions from the community. This illustrates the power of resonant information and contrasts it with access to abundant information that does not resonate with a given community.

*How many distant examples are equivalent to the example of one local leader?* Turning from the specific numerical example in our simulation to the more general point about the power of local leaders, we consider a theoretical exercise with general parameters. Consider a community where members' characteristics are evenly spaced on an $n \times n$ grid. Suppose all actions are fully visible ($\Psi = 1$). Initially, at time $t = 0$, every agent in the community chooses action 0. Since action 0 has been chosen many times in the past, beliefs about action 0 are precise ($V[z_0|\mathcal{I}_0]$ small). Now consider two alternative scenarios:

At date $t = 1$, either

1. A local leader, located in characteristic space at $(\lfloor n/2 \rceil, \lfloor n/2 \rceil)$ takes action 1, where $\lfloor \cdot \rceil$ represents rounding to the nearest integer; or

2. $\tilde{J}$ adjacent neighbors, each $\gamma > 1$ spaces away from the nearest community member, in characteristic space, all take action 1.

Define the weight that the community members closest to the local leader put on the local leader's experience as $\omega_{ll}$. Similarly, denote the weight that the community members closest to the neighboring community members put on the experience of such neighbor as $\omega_n$.

**Proposition 1.** *In order to have an equal probability of getting the nearest community member to switch to choosing $a = 1$ at date $t = 2$, as one local leader in scenario 1, it requires $\tilde{J}$ neighbors in scenario 2, where $\tilde{J}$ solves*

$$Q_1 \tilde{J} + Q_2 \tilde{J}^{1/2} + Q_3 = 0,$$

*with $Q_1 = \omega_n(E[z_0|\mathcal{I}_0] - z_1)$, $Q_3 = \sigma_1^{-2}(E[z_0|\mathcal{I}_0] - \mu_1)$, and $Q_2 = (\omega_n/\omega_{ll})Q_3 - Q_1$.*

This result teaches us that there are two effects to consider when comparing a single local leader to many distant experts. The first effect is that the local leader resonates more. The resonance of the neighboring nodes can be seen in the $Q_1$ term, where the difference between beliefs about the benefits of action 0 and the average realization of payoff to action 1, $E[z_0|\mathcal{I}_0] - z_1$, is weighted by the information weight of the neighbor. The $\tilde{J}$ signals from the more distant agents each convey, on average, that the payoff to action 1 is higher, by the amount $(z_1 - E[z_0|\mathcal{I}_0])$, weighted with $\omega_n$. The higher resonance of the local leader shows up as a lower value of $\omega_n/\omega_{ll}$.

If the relative weights were the only effect at work, the result would be that $\tilde{J} = \omega_{ll}/\omega_n > 1$. In other words, if a local leader resonated twice as much, then there would need to be twice as many neighboring examples to achieve an equivalent degree of persuasion.

The second effect, and the reason the answer is not that simple, is that having multiple observations from highly visible (albeit socially distant) information providers achieves something like a central limit theorem effect. Each observation of a person's experience has individual-specific noise, with variance $\sigma_\epsilon^2$. The noise of the sum of $\tilde{J}$ neighbor's experiences has standard deviation $\tilde{J}^{1/2}\sigma_\epsilon$. In our adapted information weighting scheme, this shows up as $\tilde{J}^{1/2}(\omega_n/\omega_{ll})$. Thus, the solution incorporates both effects.

Both the theoretical result and the numerical example convey the power that local leaders, with characteristics that are similar to their communities, have to inspire a change of behavior in their communities. But they also show that access to information has a distinct advantage: With more sources of information, an observer can average out most of the individual-specific noise.

### 4.4 Resonance with Information Frictions: Social Networks and the Decline of Expertise

In our last theoretical exploration, we examine the interaction between information access and information resonance. Online social networks have affected a radical change in information access in the last two decades. The model clarifies how these online social network platforms interact with resonance.

Prior to the digital era, access to information was constrained by geography. It was hard to learn news about the community you grew up in once you moved away. For people not part of the local majority group, it was difficult to find resonant information. In the social network era, people can form communities without geographic constraints. At the same time, the influence of experts declined. Resonance offers an explanation for this decline.

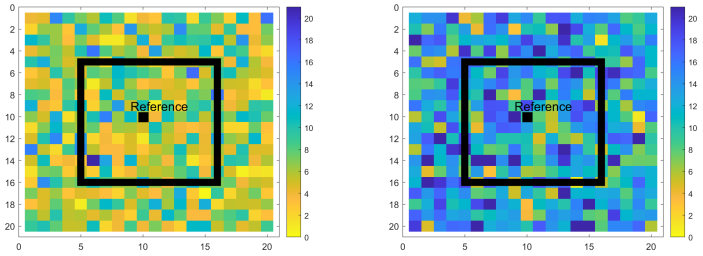
To capture the changing relevance of geography and the endogenous selection into networks in our model, we modify the information access matrix $\mathcal{Z}$. If agents $i$ and $j$ are close enough to each other in location to communicate, we set $\mathcal{Z}_{ij} = 1$ (full visibility). Otherwise, $\mathcal{Z}_{ij} = 0$.

*Simulation details.* The agents are distributed uniformly in geographic space, with one person assigned to each node in a grid. However, while characteristics were assigned uniformly in the previous simulations, they are now normally distributed. The reason for this change is that we want to illustrate the differential impact of social networks on members of the majority community and those in the minority ("in the tails"), such as ethnic or racial minorities or political extremists. Specifically, we draw two characteristics $v$ and $w$ from the distribution $N(10, 5)$. When we plot agents in characteristics space rather than geographic space, we assign each agent to the node with the closest integer value. For example, if an agent draws $(v, w) = (2.7, 9.1)$, they are assigned to

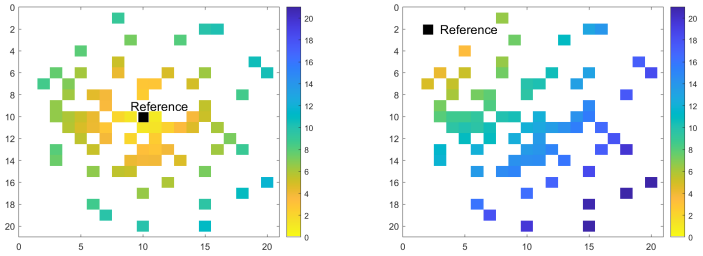
the $(3, 10)$ on the characteristic grid. In many cases, multiple agents are assigned to the same box. The number of agents assigned to a box is not visible on the plot, but it is systematically higher at the center of the grid since many agents draw characteristics closer to the mean. Around the edges of plots in characteristic space, instead, there will be many white (empty) boxes indicating cases where no agent drew characteristics in this space.

An centrist agent (left) and one with extreme characteristics (right) plotted

(a) in geographic space:




(b) in characteristic space, with geographic information constraints:




(c) in characteristic space, with no information constraint:

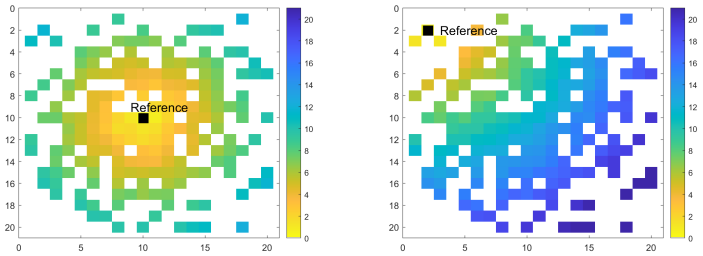



Figure 2: Social proximity in geographic and social space. Extreme agents get far more information flow when geographic constraint is relaxed. Social proximity to a reference person is represented in warmer colors (yellow tones); social distance in cooler colors (blue tones). The black frame in Panel (a) represents the geographic restriction (with $\Psi = 1$ inside the frame, and $\Psi = 0$ outside the frame) prior to social networks. The middle plots in Panel (b) include only geographically nearby agents (those inside the black frame in Panel (a), i. e., with $\Psi = 1$), plotted in characteristics space. The bottom two plots in Panel (c) include both geographically nearby and distant agents.

Figure 2 illustrates the differential effect of social networks on members of the majority and the minority groups. In Panel (a) (top row), we plot community members according to their physical

(geographic) location. One agent lives at each node. Differently from prior figures, we use colors to illustrate social proximity relative to a reference person, shown as a black dot at the center of the grid. (Note that the inital location geographic has no particular significance. We are just examining the 20-by-20 neighborhood around a given agent's location.) Warmer (yellow) tones indicate social proximity, and cooler (blue) tones indicate social distance, and each agent is given a color, according to their social proximity to the reference point. The variegated color pattern on the geographic topology represents the fact that social type and geographic location are not correlated in this example.

The difference between the left and right panels is that the reference agent in the left panel is central not just geographically, but also socially. This person has a low social distance to many other nodes, represented by the high frequency of warmer colors on the top left. The person on the right, instead, has more extreme, or minority characteristics. Because characteristics are normally distributed, the density of other agents in the nearby social space is lower. Extremists see fewer people like them. Most of their geographic neighbors have a high social distance, represented by the cooler colors on the right.

Focusing at first on the agents inside the black frame ("pre-Facebook"), we re-plot all such agents in their respective colors in characteristics space in Panel (b) of Figure 2. These panels show the social proximity of observable community members, i. e., those inside the black frame, both in color and spatially. They are then weighted according to how resonant their information is. The empty (white) boxes are characteristics pairs not represented by any agent to whom the reference person is connected. Note that, in the left graph of Panel (b), the reference person remains centrally located also in characteristics space. This reflects that we picked a socially central person, akin to the local leader in Section 4.3 before. In the right graph, instead, the reference person moves to the upper left corner in characteristics space. This reflects that we picked a socially more isolated, or outlier set of characteristics. Here the reference person has extreme characteristics and is thus located at the fringe of the social space.

Next, consider the effect of widened communities, due to social networks. In Panel (c), we allow the reference node to connect to all agents from Panel (a), whether inside our outside the black frame. The change can be interpreted as representing the ability of social media to connect across large geographic distances. Correspondingly, the bottom two panels show the social location of all nodes from Panel (a), just in different locations, representing their social distance from the reference node. Note that some locations on this plot have multiple nodes, representing multiple agents in that characteristic space location. Other nodes remain white since no agent has the corresponding characteristics. Naturally, we now see fewer white boxes. Relaxing the geographic constraint on information access allows a denser set of agents to be observed. This relaxation of the geographic constraint is meant to represent the way in which social media allows for communication

and community formation, without regard to geographic location.

**Social Networks and Social Outliers** The stronger reaction of outlier (minority) communities to changes in information access is a key prediction that distinguishes our mechanism from information frictions. While relaxing geographic constraints gives all agents access to more information, the effect is much stronger for social outliers. These are agents who have scarce access to trusted information sources, when their community is geographically restricted.

**Proposition 2.** *Expert actions have more influence on others' beliefs when agents are geography-constrained ($\partial E[z_a|\mathcal{I}]/\partial a_{celeb}$ larger).*

The proof is in Appendix A. The idea is that geographic constraints limit agents' data points. With fewer relevant data points, each one gets more weight.

Figure 2 illustrates the logic of this result. This scarce resonant information is the large number of white boxes and few warm-colored boxes around the reference agent in the center-right pane of Figure 2. Giving extremist agents like the reference agent on the right access to information from geographically distant but socially proximate agents greatly increases their resonant-information set (bottom-right panel). For the same reason, it also greatly decreases the weight they place on information from authorities.

The difference between the left and right columns of panels in Figure 2 is that the left column agent is socially central. This is someone in the majority. Since social characteristics are normally distributed, someone near the average has a high density of socially proximate neighbors. In contrast, on the right of Figure 2, the reference agent has outlier social characteristics. Because they are located in the tail of the normal distribution, this agent has fewer socially proximate neighbors to identify with and learn from. When social media relaxes the geographic constraint on information access and allows this agent to observe all nodes, some resonant information sources (yellow boxes) appear. This person with extreme characteristics will now place more weight on those resonant sources of information and less weight on all other information, including reliable and informed authorities.

### 4.5 Resonance as a Foundation for Labor Market Discrimination

Finally, resonance could explain discriminatory outcomes in job promotions. If someone in power responds more to the recommendations of others that share their traits, that bestows the power of influence with these same-trait employees. If all employees' recommendations are are equally good, same-trait employees will create more value by triggering more positive change. Thus, these employees become more valuable to retain and promote because their recommendations have more effect on the organization. Resonance can perpetuate homophily, in a way that looks rationally

justifiable, based on performance metrics. Given resonance, many forms of discrimination can become constrained optimal and self-perpetuate. Appendix E extends our model of resonance to show how it naturally lends itself to this type of labor market effect.

### 4.6 Prediction: Crises Have Type-Dependent Effects

Resonance provides a clear rationale for why there is ethnic bunching across occupations, a phenomenon we will explore next. People follow the advice of those like them, which leads them to select similar careers. However, people also like being around others like them. So geographic and occupational clustering alone do not provide clear support for an information channel. A prediction of the model about what happens when there is a negative shock to an industry helps distinguish these alternatives.

The prediction that ethnic communities will respond more to negative shocks that affect industries in which their ethnic kin are heavily represented. Deriving this prediction simply involves re-interpreting Proposition 1. In this case, the local leader is the same-ethnicity worker. The neighbors are workers of all other ethnicities. If all industries experience a negative shock, negative news is transmitted from all neighbors and local leaders. But the ethnic kin / local leader's news has a disproportionate effect. If those ethnic kin are concentrated in an industry, then the overweighted negative news from workers in that industry will cause the large reaction. We will see this take the form of workers choosing not to enter an industry where their ethnic kin experienced a negative labor market shock. This effect is larger for the ethnic kin than it is for other workers who observe the same information, but weight it less heavily.

## 5 Evidence from Occupational Choice

Having dissected resonance at the individual choice level and modeled how it spans individual choices to aggregate outcomes, we now turn to aggregate evidence. The goal of this section is not to prove that resonance operates - the lab experiment did that. Rather, the goal is to see if aggregate patterns consistent with resonance show up and show how to uncover them. Since our experiment taught us that resonance operates as more weight on information derived from those with similar characteristics, we use ethnic identity and occupational choice. While we did find that other attributes, like pet preferences for example, were stronger resonance triggers than demographic attributes, we use ethnicity because of data availability.

Using the American Community Survey (ACS), we provide suggestive evidence that people place greater weight on information learned from those within their own communities and this can have aggregate effects. Considering observational evidence of ethnic concentration in certain occupations, we explore whether "information resonance" may be a mechanism driving occupational

choice. Here we observe young labor force entrants following their community elders into the same occupations. We see that a young person's choice to enter an occupation is positively correlated with the extent to which people of their same ethnic group are overrepresented in that occupation in their locality. Then, we further test our hypothesis by examining how a negative shock to members of one's community impacts the behavior of the group. Lastly, we test the prediction that the effect on beliefs will be stronger when matching more demographic characteristics. We include an additional measure of social proximity using gender and hypothesize that young members of an ethnic community will enter occupations in which elders of the same ethnicity *and* gender are overrepresented. If men and women of the same ethnic communities have overlapping social networks, these findings can help rule out information access.

This analysis makes progress in isolating resonance from other forces. But it also makes clear the limitations. Aggregate data alone cannot rule out every possible alternative explanation. The setting is too rich in interconnections. While this section paints a picture of how resonance can operate in an aggregate economy, true identification of resonance is the job of the experiment, presented in Section 2.

**ACS Data and Variables** The micro data, obtained via IPUMS, comes from the US Census Bureau's annual American Community Survey (ACS) years 2005 to 2020. The analysis uses variables on age, occupation, education, migration, and ancestry of respondents. We aggregate 338 unique occupations (including "unemployed") into 33 broader occupation groups. Taking out those who are unemployed, we are left with 32 occupation groups. Ethnicity is measured using respondents' self-reported familial ancestry. From these ancestries, we construct 17 ethnic groups based on global geographic regions, splitting North America into 3 categories based on race, North American (primarily white North American), African American, and Indigenous American. We measure geographic proximity using Public Use Microdata Areas (PUMAs). PUMAs are geographic statistical areas of the United States each containing a population of at least 100,000 people. When considering occupation choices, we focus on individuals at the start of their career, so we restrict our sample to 18-22-year-olds. Since we aim to capture whether or not people learn from their local community, we look only at youth who have not moved out of their PUMA in the last year (since prior data about moving is not available). We also exclude 18-22-year-olds who were currently attending school. This restriction ensures that the young people in our analysis are in the occupations they chose, rather than enrolled in degree programs working towards the occupations they would like.

**Ethnic Overrepresentation** We define Ethnic Overrepresentation (Ethnic ORep) as the extent to which an ethnic group is overrepresented in an occupation within a PUMA. It is computed as the difference between the fraction of workers in an occupation from that ethnic group (Ethnic

Ratio) and the overall share of workers in that occupation (Local Ratio).[18]

Let $i$ be an ethnicity, $j$ be an occupation, $k$ be a PUMA, and $t$ be a year, then ethnic overrepresentation is defined as

$$Ethnic\ ORep_{i,j,k,t} = Ethnic\ Ratio_{i,j,k,t} - Local\ Ratio_{j,k,t}.$$

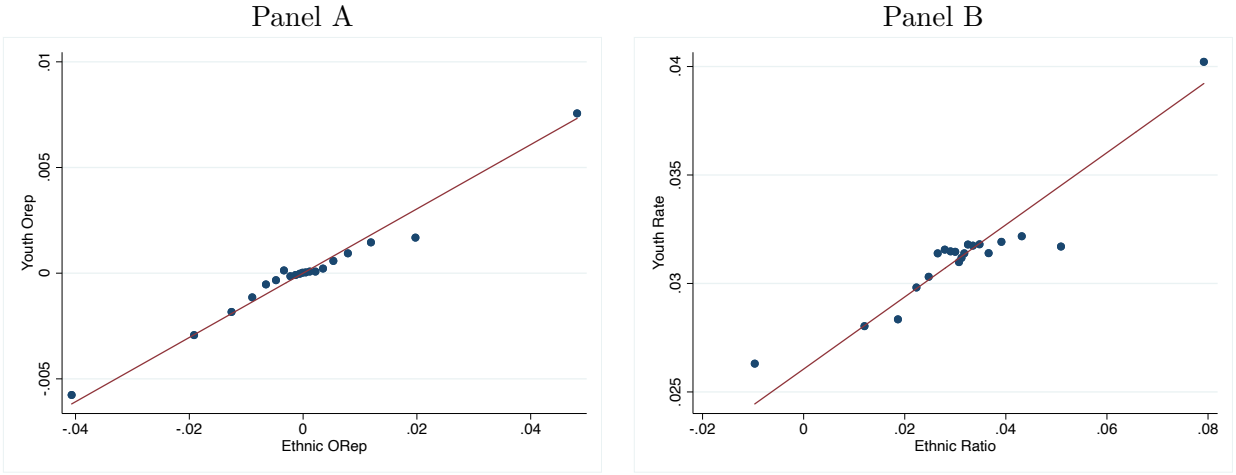
Similarly, we measure the ethnic ratio for the 18–22-year-old cohort, i.e. the fraction of 18-22-year-olds of each ethnic group in each occupation in each PUMA each year (denoted as *Youth Rate*). This variable is defined as

$$Youth\ Rate_{i,j,k,t} = \frac{Occupation\ Size_{i,j,k,t}}{Population\ Size_{i,k,t}}.$$

*Youth ORep* then refers to the extent to which young adults of a given ethnicity are overrepresented in an occupation relative to their peers.

Looking at simple correlations, we can see a pattern emerge. Figure 3 shows that young labor force entrants are following in the path of their same ethnic elders. In Panel A, we can see that as ethnic overrepresentation increases, youth overrepresentation increases, as well. Panel B shows that the same relationship holds for the level of ethnic concentration in an occupation.

Figure 3: Relationship Between Ethnic Overrespresentation and Youth Entry




***Notes***: Figures include occupation fixed effects and regressions are weighted by the number of young people in the ethnic/occupation/PUMA/year bin. Panel B includes *Local Ratio* as an additional control variable.

Taking this analysis further, we run a linear regression of the youths rate of choosing an occupation within a local-ethnic group, the proportion of 18-to-22-year-olds in each occupation in each PUMA-ethnicity pair each year, on ethnic overrepresentation, the percent amount by which the eth-

[18]For a detailed description of how these variables were constructed, see the Appendix-Section F.

nic group is overrepresented in the occupation in the PUMA that year. Our baseline specification is

$$\textit{Youth Rate}_{i,j,k,t} = \gamma_{j,k,t} + \beta(\textit{Ethnic ORep}_{i,j,k,t}) + \varepsilon_{i,j,k,t},$$

where $\gamma_{j,k,t}$ is an occupation-PUMA-year fixed effect and $\varepsilon$ is the error term. $\beta$ is the coefficient of interest on ethnic overrepresentation and has the interpretation that in a given year for a one percent increase in ethnic overrepresentation, youth choice of that occupation in that puma-ethnicity pair will increase by $\beta$ percentage points.

Table 2 displays the result. Column (6) represents our baseline specification. Here, all location, occupation, and time-specific variation is accounted for, meaning that the remaining variation in occupation choice comes from ethnicity. This allows us to interpret the coefficient on *Ethnic ORep* as the influence of ethnicity on occupation choice holding other local factors constant. If Ethnic ORep increases by 1.2% (the IQR of *Ethnic ORep*), *Youth Rate* will increase by 0.2 pp– 6% of the mean. Notice that in column (7) of Table 2, the magnitude of the coefficient drops dramatically. Here, all variation in $\beta$ comes from year-over-year changes in *Ethnic ORep*, not the level of ethnic overrepresentation itself. The drop in the coefficient provides further evidence that it is the ethnic concentration in an occupation that matters for youth occupation choice, not only population changes over time.

**Ruling out Information Frictions** In addition to ethnicity, we consider gender as another measure of social proximity. Including gender allows us to disentangle the role of information frictions in occupation choice. Presumably, husbands and wives, fathers and daughters, share information about their jobs and engage in the same social and professional networks. Following same gender elders in addition to same ethnicity elders into occupations rules out information access as the only mechanism driving the pattern.

The variable *Ethnic ORep (Wrong Gender)* is the the degree to which elders of the *same ethnic group* but *other gender* are overrepresented in an occupation.

The Appendix shows that people are following their same ethnicity elders occupational paths within gender. If a given gender of an ethnic group in a PUMA is overrepresented in an occupation by 1.2 percent, the fraction of young people of the same gender-ethnicity going into the occupation will increase by 0.28 pp. If members of the same ethnicity, but opposite gender are overrepresented by the same amount, we reduce the effect by 0.08 pp, showing that people move towards same-gender occupations, but move away from those popular among the opposite gender.

**Economic Shock** Next, we examine how ethnic communities responded to mass layoffs during the Great Recession. Young workers were significantly less likely to enter occupations where their coethnic elders had faced high job losses. This suggests that negative labor market shocks resonate

most strongly within socially proximate groups, shaping long-term career choices.

To measure the ethnic-specific effect, we calculate the difference in the number of people in the occupation-PUMA-year in 2006 versus in 2010. We then use the reduction in the number of people employed in an occupation from 2006 to 2010 as a proxy for "layoffs." The variable *High Layoffs* is defined to equal 1 if the percent reduction in employment for an occupation in a PUMA from 2006 to 2010 is in the top 30% across all PUMAs and 0 otherwise.[19] We include an indicator variable, *POST*, indicating that the observation is from the year 2010 or later (i.e., post-recession).

Table 4 shows a positive coefficient on the interaction of the high layoffs dummy and pre-recession ethnic overrepresentation, showing that prior to the recession, young people were following their coethnic elders into occupations where they were overrepresented. After the recession, the effect is reversed. Young people are deterred from occupations in which their ethnic community experienced large job losses. Looking at the triple interaction of *High Layoffs* X *POST* X *Ethnic ORep* in column (6), we see that if a job experienced high layoffs during the recession, an ethnic group being overrepresented by 1 percent prior to the recession implies that the fraction of young people entering this occupation will decrease by .2 pp after the recession. Meanwhile, we do not see a significant effect of the recession on occupations that did not experience mass layoffs (the coefficient on *POST* X *Ethnic ORep* is insignificant)– meaning that young people are continuing to follow their same-ethnic elders' career paths in the absence of a shock.

**Identification Challenges** The preceeding empirical results serve as motivating evidence that people place greater weight on information learned from people they are socially close to. However, identification of information resonance poses many questions that this empirical exercise does not address. One challenge of isolating information resonance is that correlated information and selection can mimic each other. People in a community might behave similarly because they have learned from each other and have similar beliefs. Alternatively, they might also select into a community because of similar behavior and preferences, as in models with social norms, or due to their beliefs. Another potential issue is Manski (1993)'s reflection problem: It is difficult to distinguish how an individual affects a group from the group's effect on the individual.[20]

These problems illustrate why it is useful to build a model that captures these compounding effects. They also motivate our use of experimental evidence in Section 2. But even if all these

[19]For this analysis, we fix the ethnic overrepresentation variable to its average of 2005 and 2006. Thus, this variable represents the overrepresentation of the ethnic group in an occupation before the recession. Fixing the variable to a past value allows us to see how young people respond to what they observed *before* making their occupation choice and how those observations impacted their behavior after the recession. Here, we use data from 2005 and 2006 as our "pre-recession" years, and data from 2010 and 11 to study the "post-recession" effects.

[20]The reflection problem is predominantly an issue with linear-in-means models; however, Brock and Durlauf (2001) highlights how in a backward-looking, dynamic linear-in-means model where current behaviors are driven by past beliefs the reflection problem does not apply. Our model relies centrally on past experiences, therefore it is unlikely that the reflection problem would be exhibited.

econometric problems are present, none of them negates the finding that people have social influence on others like them, in ways that affect aggregate outcomes.

# 6 Conclusion

This paper recasts the idea that people listen to people like them, as a systematic information-weighting bias. First, a tightly controlled laboratory experiment rules out memory, relevance and trust, yet uncovers a large premium on advice that resonates through shared beliefs, preferences or other characteristics. Next, we build a model that represents the information weighting found in the experiment. The model shows both how resonance operates like geographic information diffusion, but also highlights the differences and interactions between resonance and information access. It shows how one relatable role model eclipses many distant experts and generates homogenous leadership structures despite neutral promotion rules. Third, U.S. labor data reveal the same kind of bias at scale: young workers sort toward occupations dominated by ethnically similar incumbents and retract when those incumbents suffer shocks.

By formally modeling this phenomenon, we learn about policies that could be more effective in inducing behavioral change. Having the right people convey public health measure might induce better compliance. One well-placed local leader may be as effective as many, many outsiders and expensive advertising. On a darker note, the results also explain why mis-information, conveyed in chat rooms populated by those with similar interests, was so effective.

Finally, information diffusion happens not only in geographic space, but in social characteristic space. This change in focus, from information access to information resonance, could change the way we organize social science data for analysis and the way we enact policy, to improve our collective well-being.

## References


V. Alatas, A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken. Network structure and the aggregation of information: Theory and evidence from indonesia. *American Economic Review*, 106(7):1663–1704, 2016.

M. Alsan, O. Garrick, and G. Graziani. Does diversity matter for health? experimental evidence from oakland. *American Economic Review*, 109(12):4071–4111, 2019.

M. Bailey, R. Cao, T. Kuchler, and J. Stroebel. The Economic Effects of Social Networks: Evidence from the Housing Market. *Journal of Political Economy*, 126(6):2224–2276, Dec. 2018. ISSN 0022-3808. doi: 10.1086/700073.

A. Bandura. *Social learning and personality development*. New York: Holt, Rinehart, and Winston, 1963.

A. Bandura. *Social learning theory*. Englewood Cliffs, NJ: Prentice Hall, 1977.

A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies*, 86(6):2453–2490, 2019.

J. Boerma and L. Karabarbounis. Inferring Inequality with Home Production. Working Paper 24166, National Bureau of Economic Research, Dec. 2017.

W. A. Brock and S. N. Durlauf. Interactions-based models. In *Handbook of econometrics*, volume 5, pages 3297–3380. Elsevier, 2001.

C. Burnside, M. Eichenbaum, and S. Rebelo. Understanding Booms and Busts in Housing Markets. *Journal of Political Economy*, 124(4):1088–1147, 2016.

D. Buttelmann, M. Carpenter, J. Call, and M. Tomasello. Enculturated chimpanzees imitate rationally. *Developmental science*, 10(4):F31–F38, 2007.

S. E. Carrell, M. E. Page, and J. E. West. Sex and science: How professor gender perpetuates the gender gap. *The Quarterly journal of economics*, 125(3):1101–1144, 2010.

G. Chodorow-Reich, A. M. Guren, and T. J. McQuade. The 2000s Housing Cycle With 2020 Hindsight: A Neo-Kindlebergerian View. Working Paper 29140, National Bureau of Economic Research, Aug. 2021.

K. Corriveau and P. L. Harris. Preschoolers continue to trust a more accurate informant 1 week after exposure to accuracy information. *Developmental science*, 12(1):188–193, 2009.

K. H. Corriveau, P. L. Harris, E. Meins, C. Fernyhough, B. Arnott, L. Elliott, B. Liddle, A. Hearn, L. Vittorini, and M. De Rosnay. Young children's trust in their mother's claims: Longitudinal links with attachment security in infancy. *Child development*, 80(3):750–761, 2009.

F. D'Acunto, A. Fuster, and M. Weber. Diverse policy committees can reach underrepresented groups. Technical report, National Bureau of Economic Research, 2021.

M. N. Federman, D. E. Harrington, and K. J. Krynski. Vietnamese manicurists: Are immigrants displacing natives or finding new nails to polish? *Industrial and Labor Relations Review*, 59(2): 302–318, 2006. ISSN 00197939, 2162271X.

A. Fogli and V. Guerrieri. The End of the American Dream? inequality and Segregation in US Cities. Working Paper 26143, National Bureau of Economic Research, Aug. 2019.

A. Fogli and L. Veldkamp. Nature or nurture? learning and the geography female labor force participation. 2010.

B. N. Frazier, S. A. Gelman, N. Kaciroti, J. W. Russell, and J. C. Lumeng. I'll have what she's having: The impact of model characteristics on children's food choices. *Developmental science*, 15(1):87–98, 2012.

S. Gershenson, C. M. Hart, J. Hyman, C. Lindsay, and N. W. Papageorge. The long-run impacts of same-race teachers. Technical report, National Bureau of Economic Research, 2018.

M. S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.

K. M. Harris, C. T. Halpern, E. A. Whitsel, J. M. Hussey, L. A. Killeya-Jones, J. Tabor, and S. C. Dean. Cohort profile: The national longitudinal study of adolescent to adult health (add health). *International Journal of Epidemiology*, 48(5):1415–1415k, Oct. 2019. ISSN 0300-5771. doi: 10.1093/ije/dyz115.

B. Hebert and M. Woodford. Neighborhood-based information costs. 2020.

D. Hegde and J. Tumlinson. Does social proximity enhance business partnerships? theory and evidence from ethnicity's role in us venture capital. *Management Science*, 60(9):2355–2380, 2014.

R. Hertwig and D. U. Wulff. A description–experience framework of the psychology of risk. *Perspectives on psychological science*, page 17456916211026896, 2021.

R. Hertwig, R. M. Hogarth, and T. Lejarraga. Experience and description: Exploring two paths to knowledge. *Current Directions in Psychological Science*, 27(2):123–128, 2018.

F. Kindermann, J. Le Blanc, M. Piazzesi, and M. Schneider. Learning About Housing Cost: Survey Evidence from the German House Price Boom. SSRN Scholarly Paper ID 3886665, Social Science Research Network, Rochester, NY, June 2021.

K. D. Kinzler, E. Dupoux, and E. S. Spelke. The native language of social cognition. *Proceedings of the National Academy of Sciences*, 104(30):12577–12580, 2007.

K. D. Kinzler, K. H. Corriveau, and P. L. Harris. Children's selective trust in native-accented speakers. *Developmental science*, 14(1):106–111, 2011.

G. Kossinets and D. J. Watts. Origins of homophily in an evolving social network. *American journal of sociology*, 115(2):405–450, 2009.

J. Kozlowski, L. Veldkamp, and V. Venkateswaran. The tail that wags the economy: Belief-driven business cycles and persistent stagnation. *Journal of Political Economy*, 128(8):2839–2879, 2020.

T. Kuchler, Y. Li, L. Peng, J. Stroebel, and D. Zhou. Social Proximity to Capital: Implications for Investors and Firms. Working Paper 27299, National Bureau of Economic Research, June 2020a.

T. Kuchler, D. Russel, and J. Stroebel. The Geographic Spread of COVID-19 Correlates with the Structure of Social Networks as Measured by Facebook. Working Paper 26990, National Bureau of Economic Research, Apr. 2020b.

B. Larsen, M. J. Hetherington, S. H. Greene, T. J. Ryan, R. D. Maxwell, and S. Tadelis. Using donald trump's covid-19 vaccine endorsement to give public health a shot in the arm: A large-scale ad experiment. Technical report, National Bureau of Economic Research, 2022.

A. E. Learmonth, R. Lamberth, and C. Rovee-Collier. The social context of imitation in infancy. *Journal of Experimental Child Psychology*, 91(4):297–314, 2005.

U. Malmendier. Experience effects in finance: Foundations, applications, and future directions. *The Review of Finance*, 25:1339–63, 2021a.

U. Malmendier. Exposure, Experience, and Expertise: Why Personal Histories Matter in Economics. *Journal of the European Economic Association*, 19, 2021b.

U. Malmendier and S. Nagel. Depression babies: do macroeconomic experiences affect risk taking? *The quarterly journal of economics*, 126(1):373–416, 2011.

U. Malmendier and S. Nagel. Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87, 2016.

C. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies,*, 60(3):531–542, 1993.

J. M. McPherson and L. Smith-Lovin. Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *American sociological review*, pages 370–379, 1987.

M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

V. Pool, N. Stoffman, and S. Yonker. The people in your neighborhood: Social interactions and mutual fund portfolios. *The Journal of Finance*, 70(6):2679–2732, 2015. doi: https://doi.org/10.1111/jofi.12208. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12208`.

J. Price. The effect of instructor race and gender on student persistence in stem fields. *Economics of Education Review*, 29(6):901–910, 2010.

S. Seehagen and J. S. Herbert. Selective imitation in 6-month-olds: The role of the social and physical context. *Infant Behavior and Development*, 35(3):509–512, 2012.

L. A. Serbin, D. Poulin-Dubois, K. A. Colburne, M. G. Sen, and J. A. Eichstedt. Gender stereotyping in infancy: Visual preferences for and knowledge of gender-stereotyped toys in the second year. *International journal of behavioral development*, 25(1):7–15, 2001.

K. Shutts, K. D. Kinzler, C. B. McKee, and E. S. Spelke. Social information guides infants' selection of foods. *Journal of Cognition and Development*, 10(1-2):1–17, 2009.

K. Shutts, M. R. Banaji, and E. S. Spelke. Social categories guide young children's preferences for novel objects. *Developmental science*, 13(4):599–610, 2010.

B. F. Skinner. *The Behavior of Organisms: An Experimental Analysis*. B.F. Skinner Foundation, Cambridge, Massachusetts, 1938.

J. A. Smith, M. McPherson, and L. Smith-Lovin. Social distance in the united states: Sex, race, religion, age, and education homophily among confidants, 1985 to 2004. *American Sociological Review*, 79(3):432–456, 2014.

O. Stolper and A. Walter. Birds of a feather: The impact of homophily on the propensity to follow financial advice. *The Review of Financial Studies*, 32(2):524–563, 2019.

G. Stuart. Institutional change and embeddedness: Caste and gender in financial cooperatives in rural india. *International Public Management Journal*, 10(4):415–438, 2007.

M. G. Taylor. Gender influences on children's selective trust of adult testimony. *Journal of Experimental Child Psychology*, 115(4):672–690, 2013.

Table 2: Occupation Choice: Ethnicity

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| Ethnic ORep | 0.571*** | 0.176*** | 0.175*** | 0.234*** | 0.170*** | 0.180*** | 0.0393*** |
| | (0.0436) | (0.0137) | (0.0137) | (0.0133) | (0.00874) | (0.0106) | (0.0115) |
| Constant | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** |
| | (0.000449) | (0.000124) | (0.000125) | (0.000102) | (0.0000652) | (0.0000813) | (8.24e-20) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 1203328 | 1203328 | 1203328 | 1203328 | 1202976 | 1050624 | 1196736 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Occupation Choice: Ethnicity, Gender

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
|---|---|---|---|---|---|---|---|---|
| Ethnic ORep (Same Gender) | 0.406*** | 0.221*** | 0.221*** | 0.253*** | 0.218*** | 0.230*** | 0.159*** | 0.333*** |
| | (0.0258) | (0.0131) | (0.0131) | (0.0132) | (0.0121) | (0.0139) | (0.0127) | (0.0262) |
| Ethnic ORep (Wrong Gender) | 0.157*** | -0.0636*** | -0.0639*** | -0.0225* | -0.0634*** | -0.0642*** | -0.121*** | 0 |
| | (0.0200) | (0.0129) | (0.0129) | (0.0127) | (0.0117) | (0.0136) | (0.0124) | (.) |
| Constant | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** |
| | (0.000449) | (0.000125) | (0.000125) | (0.000102) | (0.0000653) | (0.0000702) | (5.57e-19) | (5.73e-19) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No | Yes |
| Occupation X Year FE | No | No | Yes | No | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy X Year FE | No | No | No | No | No | No | No | Yes |
| Observations | 2108064 | 2108064 | 2108064 | 2108064 | 2107904 | 2056448 | 2104672 | 1809472 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Impact of a Shock on Occupation Choice (2005/06 and 2010/11 with Ethnic $ORep_{2005}$)

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| Ethnic ORep (2005) | 0.614*** | 0.139*** | 0.153*** | 0.323*** | 0.147*** | 0.170*** | 0 |
| | (0.0755) | (0.0335) | (0.0338) | (0.0347) | (0.0291) | (0.0316) | (.) |
| POST=1 | 0.000798** | 0.000830** | 0 | 0.000788** | 0.000774** | 0 | 0.000798** |
| | (0.000393) | (0.000346) | (.) | (0.000346) | (0.000359) | (.) | (0.000346) |
| POST=1 X Ethnic ORep (2005) | -0.000996 | -0.0182 | -0.0468 | 0.00929 | 0.00523 | -0.0456 | 0.00557 |
| | (0.0437) | (0.0391) | (0.0396) | (0.0390) | (0.0403) | (0.0414) | (0.0391) |
| High Layoffs=1 | -0.0179*** | 0.00163*** | -0.00000514 | 0.00143*** | 0 | 0 | 0 |
| | (0.000819) | (0.000434) | (0.000425) | (0.000416) | (.) | (.) | (.) |
| High Layoffs=1 X Ethnic ORep (2005) | -0.0360 | 0.249*** | 0.225*** | 0.153** | 0.262*** | 0.229*** | 0 |
| | (0.115) | (0.0664) | (0.0661) | (0.0637) | (0.0559) | (0.0610) | (.) |
| POST=1 X High Layoffs=1 | -0.00354*** | -0.00327*** | 0.000127 | -0.00310*** | -0.00307*** | 0 | -0.00317*** |
| | (0.000590) | (0.000539) | (0.000503) | (0.000537) | (0.000557) | (.) | (0.000538) |
| POST=1 X High Layoffs=1 X Ethnic ORep (2005) | -0.178** | -0.241*** | -0.184** | -0.252*** | -0.259*** | -0.204*** | -0.259*** |
| | (0.0812) | (0.0763) | (0.0743) | (0.0752) | (0.0759) | (0.0789) | (0.0719) |
| Constant | 0.0372*** | 0.0322*** | 0.0326*** | 0.0322*** | 0.0326*** | 0.0323*** | 0.0326*** |
| | (0.000620) | (0.000262) | (0.000194) | (0.000248) | (0.000182) | (0.000136) | (0.000134) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 275427 | 275427 | 275427 | 275427 | 274986 | 242493 | 273211 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

# Online Appendix

## A Experimental Evidence

### A.1 Experiment Design and Instructions

We begin by reporting the specific instructions of the Stage 1 and Stage 2 experiments as well as information on the full set of choices and characteristics used in the experiment.

Figures 4 and 5 show the instructions of the Stage 2 and Stage 1 experiments. We split the Stage 2 experiment instructions into three pages: in the first page we describe the types of choices participants will make, in the second page we communicate the incentives associated with randomly selected choices, in the third page we make explicit the fact that each choice will be recommended by a Stage 1 participant and that we will display a 'trivia' about these recommenders. In the third page we also let participants know they can click on a link at the bottom of each choice option to display more information about these choices (except for the 'True or False' choice questions). In the Stage 1 instructions we write that the respondents choices might be used as recommendations or advice to participants in future surveys. Figure 6 shows a sample Stage 2 choice question and Figure 7 shows a sample Stage 1 choice question. Finally, Figure 8 shows a sample characteristics question (these questions are not randomized and are the same in Stage 1 and Stage 2).

Table 5 lists the characteristics of experiment participants that we elicited (also called recommender characteristics or homophily types) to determine if such characteristics evoke resonance. They are grouped into five categories (Beliefs, Biographical, Demographics, Personality, and Preferences). We chose these categories and the characteristics pairs aiming to span the set of possible characteristics that generate resonance.

Next, Table 6 lists the complete set of choices each participant was presented with. The choices are also grouped in five domains (Fiction Books, Non-Fiction Books, Charities, ETFs, and True or False).

#### A.1.1 Book and Charity Choice Questions

As already mentioned in Section 2, we chose pairs of books and charities to minimize orthogonal characteristics of choices that might interfere with our detection of resonance. In other words, we aimed to pick books and charities which are similar alongside several dimensions (comparable setting, topic, time period, title for books, and comparable field and name for charities). We then asked participants to choose between the two books or the two charities. For each choice question regarding a pair of books or charities, we showed the books or charities names and allowed participants to click on a link to access a 180-200 words description of the charity or book. The following is an example of one of the descriptions participants could access (referring to the charity called 'Plan International'):

> *Plan International is a global development and humanitarian organization that was founded in 1937 with the mission to promote children's rights and achieve gender equality. The organization places a strong emphasis on empowering girls, believing that gender equality is essential to improving the well-being of all children. Operating in over 75 countries worldwide, Plan International addresses the root causes of issues like poverty, violence, exclusion, and discrimination that disproportionately affect children, especially*

*girls. By partnering with local communities, governments, and other organizations, Plan International implements long-term, sustainable solutions that help children learn, lead, decide, and thrive. The organization focuses on critical areas such as education, health, child protection, and youth economic empowerment to foster environments where children can flourish. In addition to these efforts, Plan International provides emergency humanitarian response during crises like natural disasters and conflicts, ensuring that children receive the protection and care they need. The organization also engages in advocacy work, aiming to influence policies and practices at local, national, and global levels to create systemic change. Ultimately, Plan International's goal is to create a lasting, positive impact for all children, with a particular focus on advancing the rights and opportunities of girls.*

### A.1.2 True or False Choice Questions

We constructed 'True or False' choice questions as follows:

1. We used data from the The National Longitudinal Study of Adolescent to Adult Health (Add Health)[21]. Specifically, we used the third wave of the study, which contains follow-up interviews from 15,197 original Wave I respondents. These interviews were conducted between April 2001 and April 2002 and all participants were between 18 and 26 years old at the moment of the interviews.

2. We restricted our focus to questions from the 'general health and diet', 'propensity for risk', 'civic participation and citizenship', 'gambling', 'daily activities', and 'bio-specimen participation' Wave III interview modules.

3. We picked pairs of questions (for instance question H3GM1 'Have you ever bought lottery tickets, such as daily, scratch-offs, or lotto?' and question H3GH19 'On how many of the past seven days did you eat breakfast—that is, a meal within an hour of getting up?') and formulated the 'True or False' choice questions to be about conditional population shares (for instance we asked participants 'Please guess which of the following is true **among people who have ever bought a lottery ticket**: the share of individuals who have eaten breakfast in the past 7 days is higher (lower) than 50%').

4. We then obtained the 'true' answers to the 'True or False' choice questions by calculating the corresponding properly weighted conditional population shares in the data.

5. Finally, we used the conditional population shares to assign truth values to the correct choice option within each choice question pair.

6. In the Stage 2 experiment instructions (see Figure 4) we wrote that we would randomly select one of the choice questions participants answered to (for a randomly 10% of participants). If

[21] Add Health is a longitudinal study of a nationally representative sample of over 20,000 adolescents who were in grades 7-12 during the 1994-95 school year, and have been followed for five waves to date, most recently in 2016-18 (Harris et al., 2019). Over the years, Add Health has collected rich demographic, social, familial, behavioral, psychosocial, cognitive, and health survey data from participants and their parents; a vast array of contextual data from participants' schools, neighborhoods, and geographies of residence; and in-home physical and biological data from participants, including genetic markers, blood-based assays, anthropometric measures, and medications (Harris et al., 2019).

the randomly selected choice was a True or False choice question and the participant chose correct choice option, we compensated them with a $5 bonus.

The choice of the survey modules in (2) and the construction of the choice matches in (3) were deliberately carried out to make it hard for experiment respondents to easily find out which of the two choice options was actually 'true'.

### A.1.3 ETFs Choice Questions

We constructed 'ETFs' choice questions as follows:

1. We obtained monthly price data on pairs of ETFs which cover the same sector (for instance two ETFs tracking healthcare companies, gold, etc.) from Yahoo Finance[22]. We specifically price data in the range January 1st 2020 - November 1st 2024.

2. We calculated monthly rates of return (using the 'adjusted close' price, i.e. the closing price after adjustments for all applicable splits and dividend distributions).

3. For each pair of ETFs, we randomly selected 14 months of returns among those months whereby the difference in the rates of return of the two ETFs was less than or equal to 1pp in absolute value.

4. We assigned generic names to the ETFs within each pair (for instance 'Healthcare Fund A' and 'Healthcare Fund B') and asked participants to pick which of the two ETFs had a higher rate of return in the month preceding the experiment. In the table where I show the choice questions I only report the ETFs names we used in the experiment. Should we also report the actual names of the ETFs we used somewhere in this section?

5. We gave participants the possibility to click on a link to display the 14 randomly selected months of returns for each ETF within a given ETFs pair. For both ETFs within a given ETF pair we displayed the returns into a table composed of two columns each counting seven rows. Each row entry corresponded to the return realized during a specific month, which we imposed to be the same for *both* ETFs. This insured that the returns of the two ETFs were similar row-wise. To avoid any systematic relationship between the display order of the rates of return and future predicted rates, we randomized the position of the months in the rates of return display tables.

6. In the Stage 2 experiment instructions (see Figure 4) we wrote that we would randomly select one of the choice questions participants answered to (for a randomly 10% of participants). If the randomly selected choice was an ETFs choice question and the participant chose the ETF which had the highest rate of return in the month preceding the survey, we compensated them with a $5 bonus.

The choice of the two ETFs for each pair in (1) and the sampling of the rates of return months in (2) were deliberately aimed at making it hard for participants to ascertain which ETF had a higher rate of return. Nevertheless, as will be more thoroughly discussed in subsection A.3, for some ETFs participants disproportionately chose a given option over the other. This suggests that despite our

[22] https://finance.yahoo.com/

efforts, participants are not completely indifferent between the ETFs of every ETFs pair. Here's an example of a possible ordering of the 14 monthly returns of 'Healthcare Fund A':

*Some recent monthly returns (from the past six years):*

- -4.53%
- 2.4%
- 7.12%
- 4.32%
- -0.88%
- -5.25%
- 1.91%
- 5.61%
- -4.11%
- -1.63%
- 3.55%
- 3.12%
- 4.27%
- -0.04%

Table 5: Categorized Homophily Variables

| **Categories** | **Homophily Variables** |
|---|---|
| Beliefs | AI will improve (worsen) society |
| Beliefs | Democrat (Republican) |
| Beliefs | Historical events are mostly random (follow a predetermined path) |
| Beliefs | The assassin of John F. Kennedy acted alone (with others) |
| Beliefs | We should (should not) increase our reliance on nuclear power |
| Biographical | Parents divorced or separated (did not divorce or separate) during childhood |
| Biographical | Played (did not play) a musical instrument for at least 5 years growing up |
| Biographical | Practiced (did not practice) a sport for at least 5 years growing up |
| Biographical | Saw (didn't see) snow most winters growing up |
| Biographical | Was (was not) allowed to watch TV often as a child |
| Demographics | From the North (South) of the US |
| Demographics | Man (Woman) |
| Demographics | Married (not married) |
| Demographics | Religious (not religious) |
| Demographics | White (not white) |
| Personality | Analytical (creative) thinker |
| Personality | Competitive (cooperative) |
| Personality | Introvert (extrovert) |
| Personality | Morning person (night owl) |
| Personality | Optimist (pessimist) |
| Preferences | Cat (dog) person |
| Preferences | Has (has not) fired a gun |
| Preferences | Likes (does not like) coffee |
| Preferences | Likes Coca-Cola (Pepsi) |
| Preferences | Prefers a beach (mountain) vacation |

*Note*: List of homophily variable pairs used in the experiment.

Table 6: Categorized Choice Variables

| **Domain** | **Choices** |
|---|---|
| Books (Fiction) | Suldrun's Garden (The Riddle-Master of Hed) |
| Books (Fiction) | The Futurological Congress (The Shockwave Rider) |
| Books (Fiction) | The Game of Kings (The King's General) |
| Books (Fiction) | The Leavenworth Case (The Mystery of a Hansom Cab) |
| Books (Fiction) | The Monk (Melmoth the Wanderer) |
| Books (Non-Fiction) | Deep Work (The Power of Habit) |
| Books (Non-Fiction) | Emotional Agility (Insight) |
| Books (Non-Fiction) | Learn Python the Hard Way (Python Crash Course) |
| Books (Non-Fiction) | The Joy of Less (The Power of Purpose) |
| Books (Non-Fiction) | The Simple Path to Wealth (Your Money or Your Life) |
| Charities | Better Angels (The Bridge Alliance) |
| Charities | ChildFund International (Plan International) |
| Charities | Disability Rights Fund (Disability Rights International) |
| Charities | Jump$tart Coalition for Financial Literacy (National Endowment for Financial Education) |
| Charities | Water.org (WaterAid) |
| ETFs | Emerging Markets Fund A (B) |
| ETFs | Gold Fund A (B) |
| ETFs | Healthcare Fund A (B) |
| ETFs | Japan Exposure Fund A (B) |
| ETFs | Real Estate Fund A (B) |
| True or False | Among people who like to drink alcohol or smoke marijuana, share of registered voters is higher (lower) than 60% |
| True or False | Among people who consider themselves vegetarians share who filed a tax return in the past 12 months is higher (lower) than 70% |
| True or False | Among people who think that something is wrong with people who need liquor to feel good, share who have done housework is higher (lower) than 80% |
| True or False | Among people who have ever bought a lottery ticket, share who have eaten breakfast in the past 7 days is higher (lower) than 50% |
| True or False | Among people who bike, skateboard, dance, hike, hunt, or do yard work at least once a week, share who wake up before 5:00 AM is higher (lower) than 20% |

*Note*: List of choice variable pairs used in the experiment.

**Instructions: Making Choices**

We will present you with three kinds of questions:

1. We will present you with a series of pairs of charities or books. We are interested in knowing which charities or books **you would prefer** from each of these pairs.
2. We will also present you with a series of pairs of True or False statements. For this set of questions we would like you to choose which statement **you think is correct**.
3. Finally, we will present you with a series of pairs of Exchange-Traded Funds or ETFs (investment funds that hold multiple underlying assets and can be bought and sold on stock exchanges). For this set of choices, we would like you to **predict which fund had the higher investment return last month**.

→

**Instructions: Effects of Your Choices**

After you make your choices we will randomly select one of these questions and

- if it is a book choice, actually send you the book you chose electronically;
- if it is a charity choice, actually donate $10 to the charity you chose on your behalf;
- if it is a True or False statement *and* you chose the correct statement, we will give you a $5 bonus;
- if it is an Exchange-Traded Fund *and* you correctly predicted which fund had the highest investment return last month, we will give you a $5 bonus.

We will do this for 10% of randomly selected participants. We will only contact you if you are selected. If selected we will contact you within 30 days of the day you took the survey.

→

**Instructions: Recommendations**

- Previously, we had a number of participants make the **same choices** based on information about each option in each case.
- To help you make your choices, we will tell you which option one of these past participants recommended, or which guess they submitted.
- Our Institutional Review Board (IRB) does not allow us to give personal information about participants in our studies. Instead, we will tell you a piece of trivia about each of these recommenders, based on their answers to a questionnaire we administered.
- For book, charity and ETF choices you will have the option to learn more information about the items by clicking links at the bottom of the screen.

→

Figure 4: Stage 2 instructions.

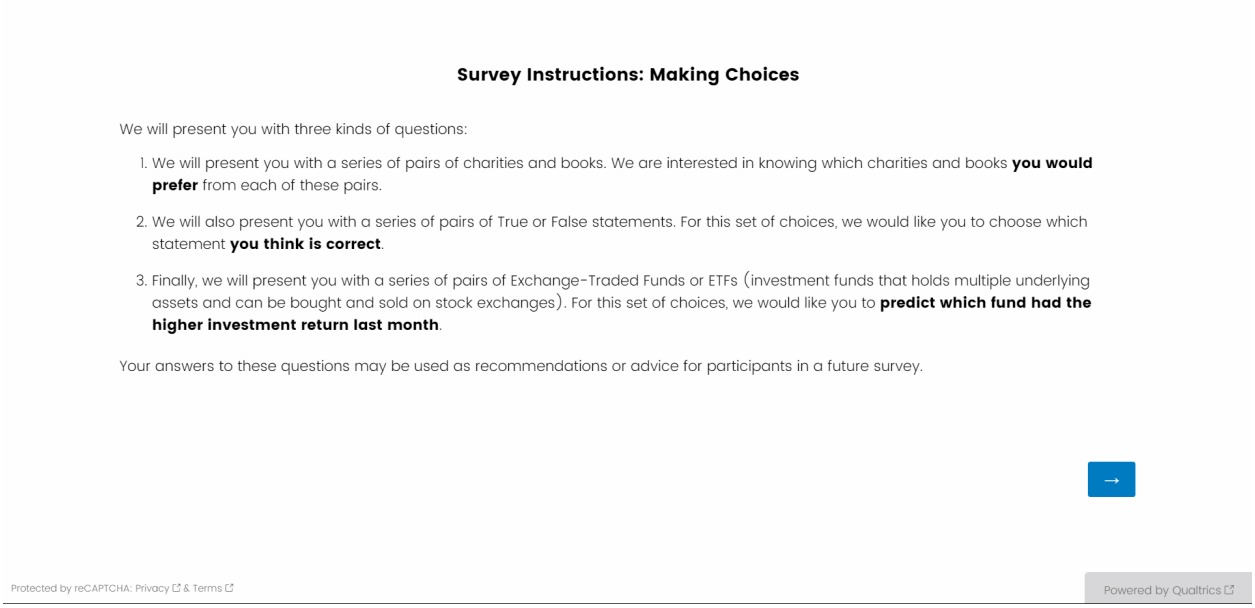



Figure 5: Stage 1 instructions.

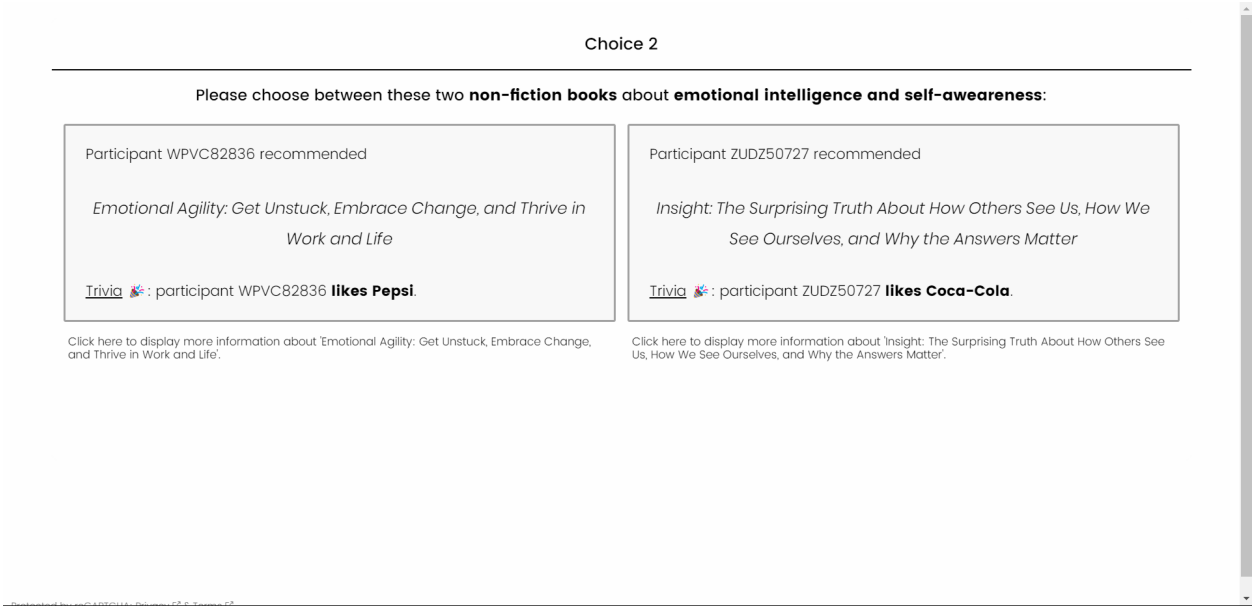



Figure 6: Example of a Stage 2 choice question.

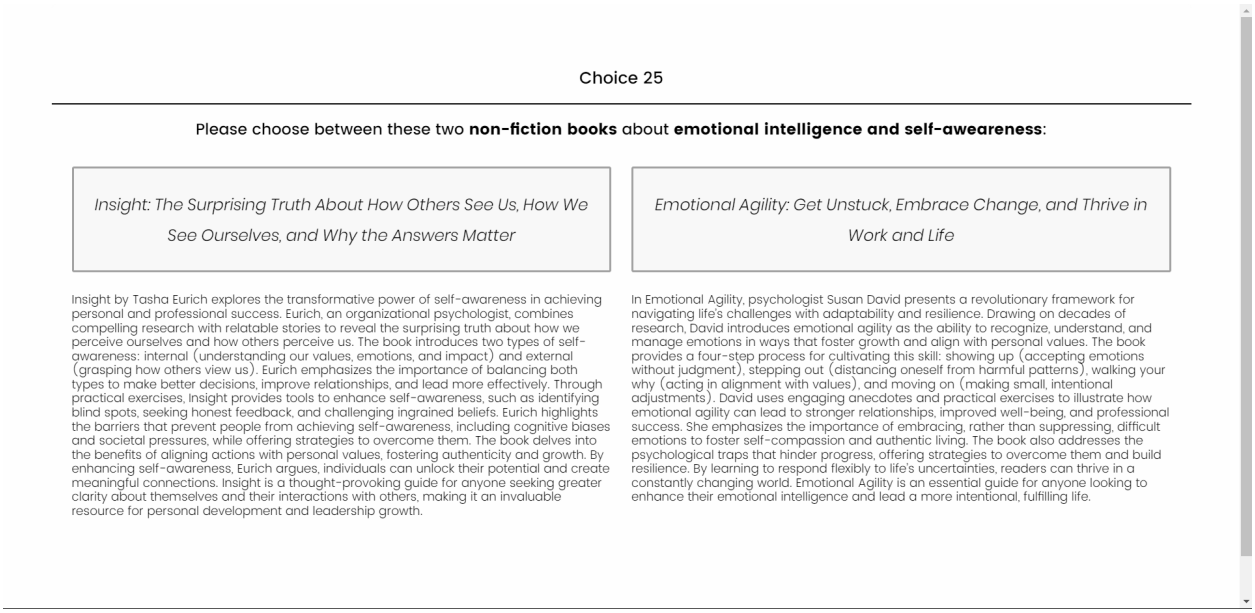



Figure 7: Example of a Stage 1 choice question.

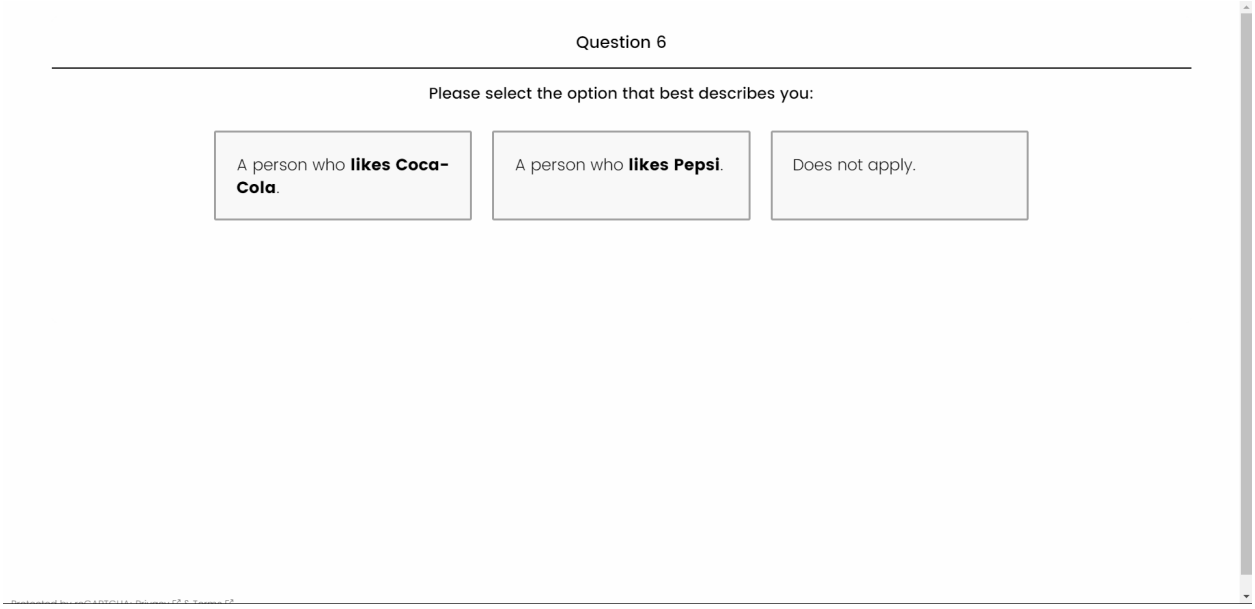



Figure 8: Example of a Stage 1 or Stage 2 characteristics question.

## A.2 Resonance Estimates are Robust to the Distribution of Characteristics

In this section we show that resonance rates are robust to the distribution of characteristics in the sample of respondents.

Table 7 reports how many Stage 2 participants are best described by each of these characteristics. As we can see, the characteristics are mostly balanced. Nevertheless there still are characteristics pairs which are more skewed towards a specific type. For instance, there are disproportionately more participants who like coffee (247) than who do not (49) or there are disproportionately more participants who consider themselves as optimists (242) as compared to pessimists (49).

However, as we show in Tables 8 and 9 resonance rates are stable whether we account for these imbalances in respondent types or not. That is, resonance rates from unweighted or weighted specifications by the inverse of the respondent type counts are mostly comparable both when we breakdown resonance rates by recommender characteristics and by choice questions. We also show the difference between the unweighted and weighted resonance rates graphically for the breakdown in therms of characteristics in Figure 9. Am I forgetting some robustness analyses we want to show in this section? Should we make this whole section a general 'Robustness' section?

Table 7: Homophily Type Counts

| Characteristic (1) | Count (2) | Total (3) |
|---|---|---|
| Cat person | 92 | 281 |
| Dog person | 189 | 281 |
| Does not like coffee | 49 | 296 |
| Likes coffee | 247 | 296 |
| Extrovert | 106 | 296 |
| Introvert | 190 | 296 |
| Morning person | 183 | 292 |
| Night owl | 109 | 292 |
| Prefers a beach vacation | 180 | 293 |
| Prefers a mountain vacation | 113 | 293 |
| Likes Coca-Cola | 186 | 270 |
| Likes Pepsi | 84 | 270 |
| Democrat | 112 | 273 |
| Republican | 161 | 273 |
| From the North of the US | 144 | 282 |
| From the South of the US | 138 | 282 |
| Not religious | 96 | 293 |
| Religious | 197 | 293 |
| Man | 146 | 298 |
| Woman | 152 | 298 |
| Not white | 81 | 297 |
| White | 216 | 297 |
| Did not play a musical instrument for at least 5 years growing up | 170 | 296 |
| Played a musical instrument for at least 5 years growing up | 126 | 296 |
| Was allowed to watch TV often as a child | 247 | 294 |
| Was not allowed to watch TV often as a child | 47 | 294 |
| Married | 203 | 297 |
| Not married | 94 | 297 |
| Optimist | 242 | 291 |
| Pessimist | 49 | 291 |
| Competitive | 112 | 293 |
| Cooperative | 181 | 293 |
| Did not see snow most winters growing up | 89 | 296 |
| Saw snow most winters growing up | 207 | 296 |
| Analytical thinker | 134 | 294 |
| Creative thinker | 160 | 294 |
| We should decrease our reliance on nuclear power | 141 | 288 |
| We should increase our reliance on nuclear power | 147 | 288 |
| AI will improve society | 225 | 294 |
| AI will worsen society | 69 | 294 |
| Parents did not divorce or separate during childhood | 224 | 297 |
| Parents divorced or separated during childhood | 73 | 297 |
| Has fired a gun | 122 | 297 |
| Has not fired a gun | 175 | 297 |
| Did not practice a sport for at least 5 years growing up | 94 | 297 |
| Practiced a sport for at least 5 years growing up | 203 | 297 |
| The assassin of John F. Kennedy acted alone | 98 | 280 |
| The assassin of John F. Kennedy acted with others | 182 | 280 |
| Historical events are mostly random | 140 | 290 |
| Historical events follow a predetermined path | 150 | 290 |

*Note*: Frequency of characteristics (homophily types) in the Stage 2 sample. The entries of column (3) do not add to 300 because participants are allowed to opt out of responding the homophily questions. Thus, the entries of the 'Total' column represent the total number of participants who did not respond 'Does not apply' for each characteristics question.

Table 8: Resonance Rates by Recommender Characteristics Pair

| | | *Homophily Choice* | |
|---|---|---|---|
| *Category* | *Characteristic* | Unweighted (1) | Weighted (2) |
| Beliefs | AI will improve (worsen) society | 0.221 | 0.182 |
| | | (0.0276) | (0.0279) |
| | | [1.4e-15] | [6.8e-11] |
| Beliefs | Democrat (Republican) | 0.236 | 0.24 |
| | | (0.0287) | (0.0279) |
| | | [2e-16] | [9.5e-18] |
| Beliefs | Historical events are mostly random (follow a predetermined path) | 0.19 | 0.189 |
| | | (0.0278) | (0.0279) |
| | | [9.8e-12] | [1.4e-11] |
| Beliefs | The assassin of JFK acted alone (with others) | 0.214 | 0.198 |
| | | (0.0283) | (0.0279) |
| | | [4.1e-14] | [1.4e-12] |
| Beliefs | We shuold increase (decrease) reliance on nuclear power | 0.215 | 0.215 |
| | | (0.0279) | (0.0279) |
| | | [1.4e-14] | [1.4e-14] |
| Biographical | Parents divorced or separated (did not divorce or separate) during childhood | 0.106 | 0.086 |
| | | (0.0275) | (0.0279) |
| | | [0.00011] | [0.0019] |
| Biographical | Played (did not play) a musical instrument for at least 5 years growing up | 0.061 | 0.072 |
| | | (0.0275) | (0.0279) |
| | | [0.027] | [0.0093] |
| Biographical | Practiced (did not practice) a sport for at least 5 years growing up | 0.136 | 0.114 |
| | | (0.0275) | (0.0279) |
| | | [7.1e-07] | [4.3e-05] |
| Biographical | Saw (didn't see) most winters growing up | 0.145 | 0.112 |
| | | (0.0275) | (0.0279) |
| | | [1.3e-07] | [6e-05] |
| Biographical | Was (was not) allowed to watch TV often as a child | 0.109 | 0.172 |
| | | (0.0276) | (0.0279) |
| | | [8.2e-05] | [6.4e-10] |
| Demographic | From the North (South) of the US | 0.131 | 0.13 |
| | | (0.0282) | (0.0279) |
| | | [3.3e-06] | [3.2e-06] |
| Demographic | Man (Woman) | 0.107 | 0.107 |
| | | (0.0274) | (0.0279) |
| | | [9.1e-05] | [0.00012] |
| Demographic | Married (not married) | 0.146 | 0.136 |
| | | (0.0275) | (0.0279) |
| | | [1e-07] | [1.1e-06] |
| Demographic | Religious (not religious) | 0.172 | 0.192 |
| | | (0.0277) | (0.0279) |
| | | [4.9e-10] | [5.7e-12] |
| Demographic | White (not white) | 0.13 | 0.103 |
| | | (0.0275) | (0.0279) |
| | | [2.4e-06] | [0.00023] |
| Personality | Analytical (creative) thinker | -0.037 | -0.035 |
| | | (0.0276) | (0.0279) |
| | | [0.18] | [0.21] |
| Personality | Cooperative (competitive) | 0.2 | 0.194 |
| | | (0.0277) | (0.0279) |
| | | [5.8e-13] | [3.7e-12] |
| Personality | Extrovert (introvert) | 0.159 | 0.159 |
| | | (0.0275) | (0.0279) |
| | | [8.3e-09] | [1.2e-08] |
| Personality | Morning person (night owl) | 0.171 | 0.167 |
| | | (0.0277) | (0.0279) |
| | | [6.8e-10] | [2.1e-09] |
| Personality | Optimist (pessimist) | -0.04 | -0.028 |
| | | (0.0278) | (0.0279) |
| | | [0.15] | [0.32] |
| Preferences | Cat (dog) person | 0.219 | 0.21 |
| | | (0.0282) | (0.0279) |
| | | [1.1e-14] | [5.3e-14] |
| Preferences | Has (has not) fired a gun | 0.123 | 0.113 |
| | | (0.0275) | (0.0279) |
| | | [7.8e-06] | [5.1e-05] |
| Preferences | Likes (does not like) coffee | 0.196 | 0.146 |
| | | (0.0275) | (0.0279) |
| | | [1.2e-12] | [1.6e-07] |
| Preferences | Likes Coca-Cola (Pepsi) | 0.233 | 0.231 |
| | | (0.0288) | (0.0279) |
| | | [6.6e-16] | [1.2e-16] |
| Preferences | Prefers a beach (mountain) vacation | 0.145 | 0.145 |
| | | (0.0277) | (0.0279) |
| | | [1.6e-07] | [1.9e-07] |
| $R^2$ | | 0.1061 | 0.0991 |
| Adj.$R^2$ | | 0.1031 | 0.096 |
| Sample Size | | 7275 | 7275 |

*Note*: OLS estimates of homophily choice (centerned at 0) on homophily characteristics. Individual-level clustered standard errors (HC1) in round brackets. P-values in square brackets. Column (2) adjust observations so that each homophily type contributes equally. The weights are constructed as follows: (i) we select respondent's choices with the same recommender's homophily characteristics, (ii) we count how many respondents are of either homophily type, (iii) we assign weights equal to the inverse of the number of respondents of the same type, (iv) we repeat the former for every recommender characteristic and choice set.

Table 9: Resonance Rates by Choice Pair

| | | *Homophily Choice* | |
|---|---|---|---|
| *Domain* | *Choice* | Unweighted (1) | Weighted (2) |
| Books (Fiction) | Melmoth the Wanderer (The Monk) | 0.104 | 0.112 |
| | | (0.0279) | (0.0284) |
| | | [0.00019] | [8e-05] |
| Books (Fiction) | Suldrun's Garden (The Riddle-Master of Hed) | 0.163 | 0.171 |
| | | (0.0279) | (0.0281) |
| | | [4.9e-09] | [1.3e-09] |
| Books (Fiction) | The Futurological Congress (The Shockwave Rider) | 0.144 | 0.128 |
| | | (0.0278) | (0.0281) |
| | | [2.3e-07] | [5.2e-06] |
| Books (Fiction) | The Game of Kings (The King's General) | 0.198 | 0.187 |
| | | (0.028) | (0.028) |
| | | [1.9e-12] | [2.5e-11] |
| Books (Fiction) | The Leavenworth Case (The Mystery of a Hansom Cab) | 0.166 | 0.142 |
| | | (0.0281) | (0.028) |
| | | [3.8e-09] | [4.1e-07] |
| Books (Non-Fiction) | Deep Work (Power of Habit) | 0.114 | 0.116 |
| | | (0.0281) | (0.0282) |
| | | [5e-05] | [3.8e-05] |
| Books (Non-Fiction) | Emotional Agility (Insight) | 0.159 | 0.156 |
| | | (0.0279) | (0.0279) |
| | | [1.4e-08] | [2.5e-08] |
| Books (Non-Fiction) | Essentialism (Joy of Less) | 0.175 | 0.194 |
| | | (0.028) | (0.028) |
| | | [4.4e-10] | [4.8e-12] |
| Books (Non-Fiction) | Learn Python the Hard Way (Python Crash Course) | 0.131 | 0.136 |
| | | (0.0279) | (0.0278) |
| | | [2.6e-06] | [1.1e-06] |
| Books (Non-Fiction) | The Simple Path to Wealth (Your Money or Your Life) | 0.131 | 0.121 |
| | | (0.0279) | (0.0273) |
| | | [2.6e-06] | [9.2e-06] |
| Charities | Better Angels (Bridge Alliance) | 0.129 | 0.129 |
| | | (0.028) | (0.0281) |
| | | [4.3e-06] | [4.2e-06] |
| Charities | ChildFund (Plan) International | 0.14 | 0.143 |
| | | (0.0281) | (0.028) |
| | | [6.3e-07] | [3.4e-07] |
| Charities | Disability Rights Fund (International) | 0.112 | 0.116 |
| | | (0.0283) | (0.0284) |
| | | [7.6e-05] | [4.7e-05] |
| Charities | Jump$tart Coalition (National Endowment) for Financial Literacy | 0.178 | 0.185 |
| | | (0.0281) | (0.0278) |
| | | [2.4e-10] | [2.9e-11] |
| Charities | Water.org (WaterAid) | 0.224 | 0.212 |
| | | (0.0279) | (0.0281) |
| | | [9.1e-16] | [5.1e-14] |
| ETFs | Emerging Markets Fund A (B) | 0.19 | 0.211 |
| | | (0.0282) | (0.0282) |
| | | [1.8e-11] | [7.5e-14] |
| ETFs | Gold Fund A (B) | 0.137 | 0.134 |
| | | (0.0278) | (0.0278) |
| | | [8.2e-07] | [1.4e-06] |
| ETFs | Healthcare Fund A (B) | 0.176 | 0.144 |
| | | (0.0279) | (0.0287) |
| | | [3.2e-10] | [5.1e-07] |
| ETFs | Japan Exposure Fund A (B) | 0.147 | 0.114 |
| | | (0.0278) | (0.028) |
| | | [1.2e-07] | [5e-05] |
| ETFs | Real Estate Fund A (B) | 0.149 | 0.122 |
| | | (0.028) | (0.0276) |
| | | [9.7e-08] | [1.1e-05] |
| True or False | Share who filed a tax return is higher (lower) than 70% | 0.153 | 0.134 |
| | | (0.028) | (0.0284) |
| | | [4.9e-08] | [2.5e-06] |
| True or False | Share who did housework is higher (lower) than 80% | 0.127 | 0.142 |
| | | (0.0282) | (0.0283) |
| | | [6.6e-06] | [5.1e-07] |
| True or False | Share who ate breakfast is higher (lower) than 50% | 0.096 | 0.084 |
| | | (0.0282) | (0.0287) |
| | | [0.00068] | [0.0033] |
| True or False | Share who wake up before 5AM is higher (lower) than 20% | 0.123 | 0.108 |
| | | (0.0281) | (0.0281) |
| | | [1.3e-05] | [0.00012] |
| True or False | Share of registered voters is higher (lower) than 60% | 0.099 | 0.106 |
| | | (0.0282) | (0.028) |
| | | [0.00043] | [0.00015] |
| $R^2$ | | 0.0901 | 0.0851 |
| Adj.$R^2$ | | 0.0869 | 0.0819 |
| Sample Size | | 7275 | 7275 |

*Note*: OLS estimates of homophily choice (centerned at 0) on homophily choices. Individual-level clustered standard errors (HC1) in round brackets. P-values in square brackets. Column (2) adjust observations so that each homophily type contributes equally. The weights are constructed as follows: (i) we select respondent's choices with the same recommender's homophily characteristics, (ii) we count how many respondents are of either homophily type, (iii) we assign weights equal to the inverse of the number of respondents of the same type, (iv) we repeat the former for every recommender characteristic and choice set.

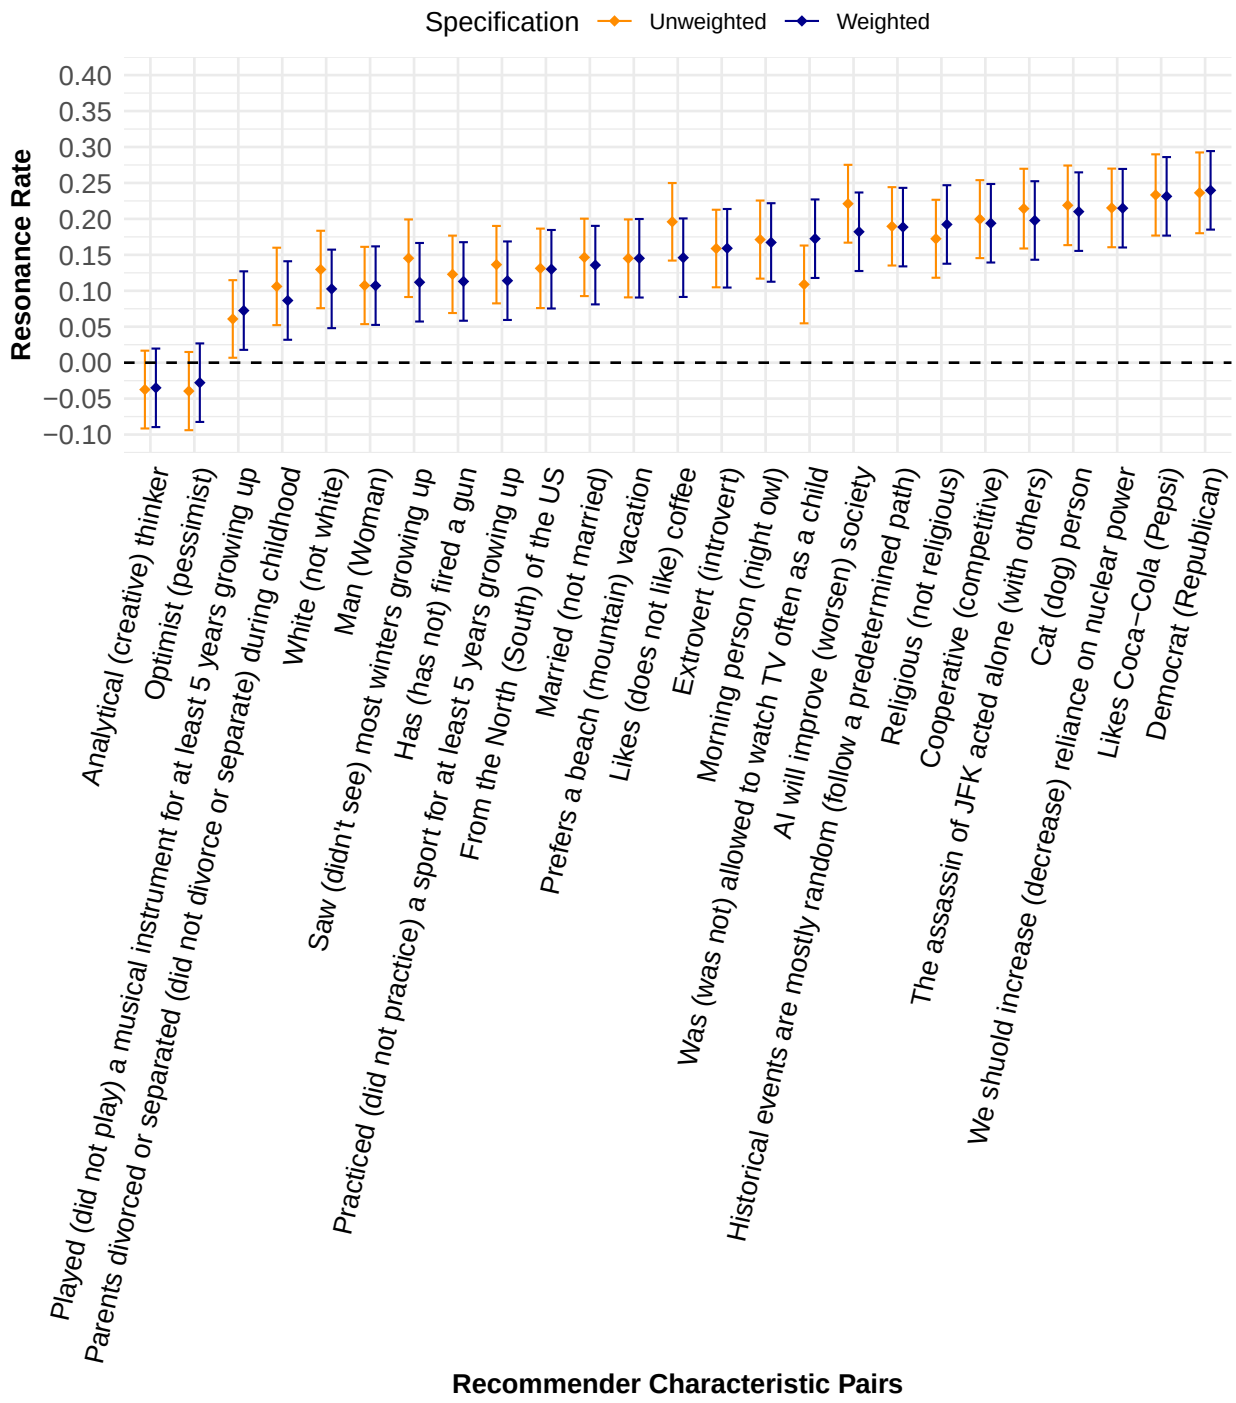



Figure 9: OLS estimates of homophily choice (centerned at 0) on homophily characteristics using individual-level clustered standard errors (HC1) to construct asymptotic confidence intervals at the 95% confidence level. The weights are constructed as follows: (i) we select respondent's choices with the same recommender's homophily characteristics, (ii) we count how many respondents are of either homophily type, (iii) we assign weights equal to the inverse of the number of respondents of the same type, (iv) we repeat the former for every recommender characteristic and choice set.

## A.3 Information Content of Characteristics

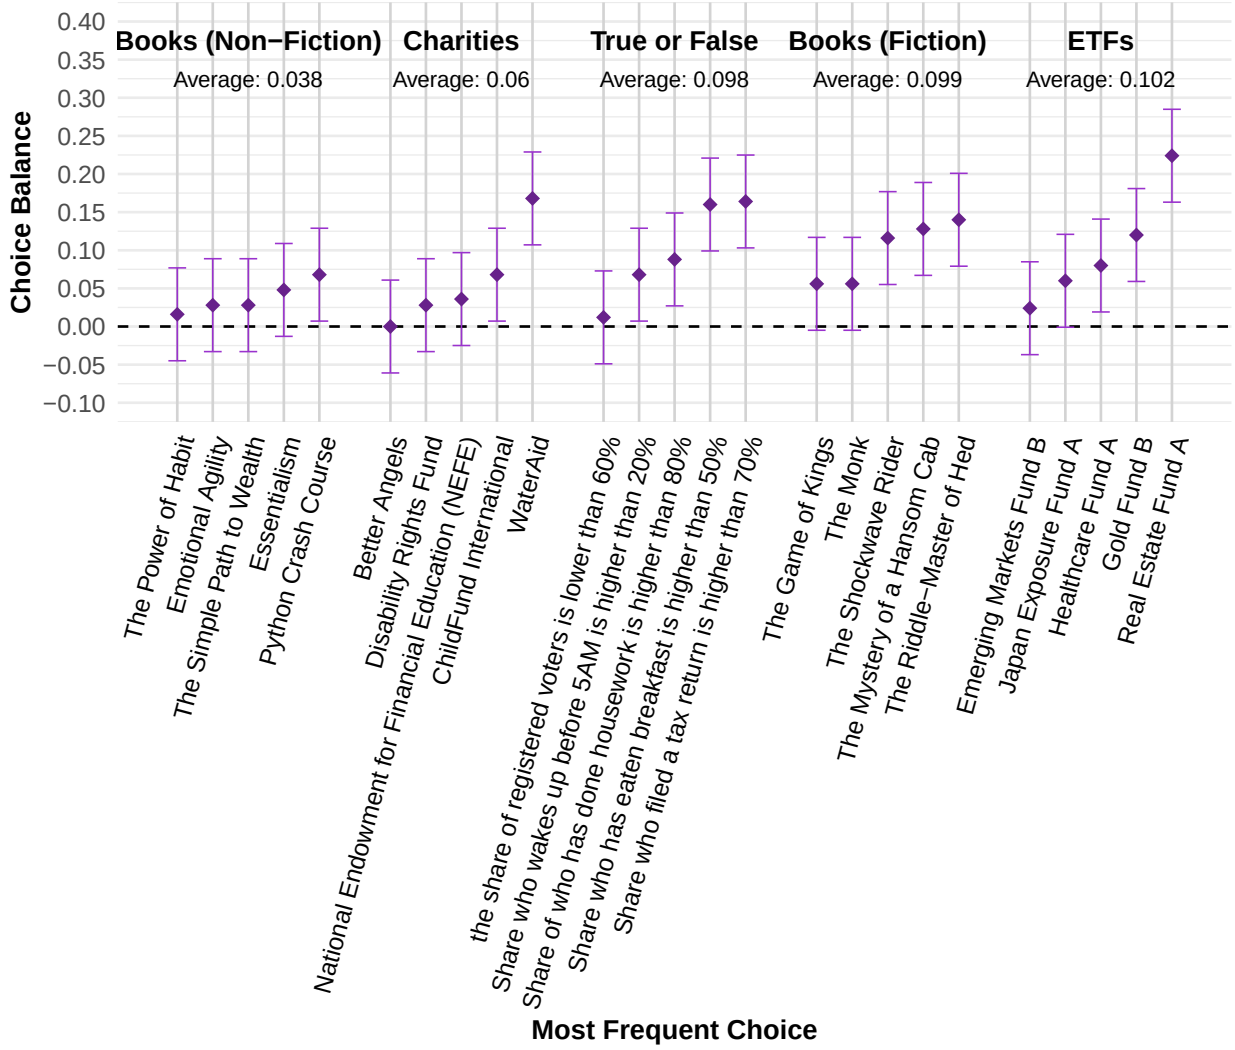



Figure 10: Testing whether Stage 1 participants unconditionally prefer one of the two choice options for each choice question. Showing OLS estimates of a dummy (centered at 0) equal to 1 if the most frequently chosen option among a given choice question gets selected on dummies for the choice questions. The x-axis labels show the most frequently chosen option among a given choice question. Using individual-level clustered standard errors (HC1) to construct asymptotic confidence intervals at the 95% confidence level.

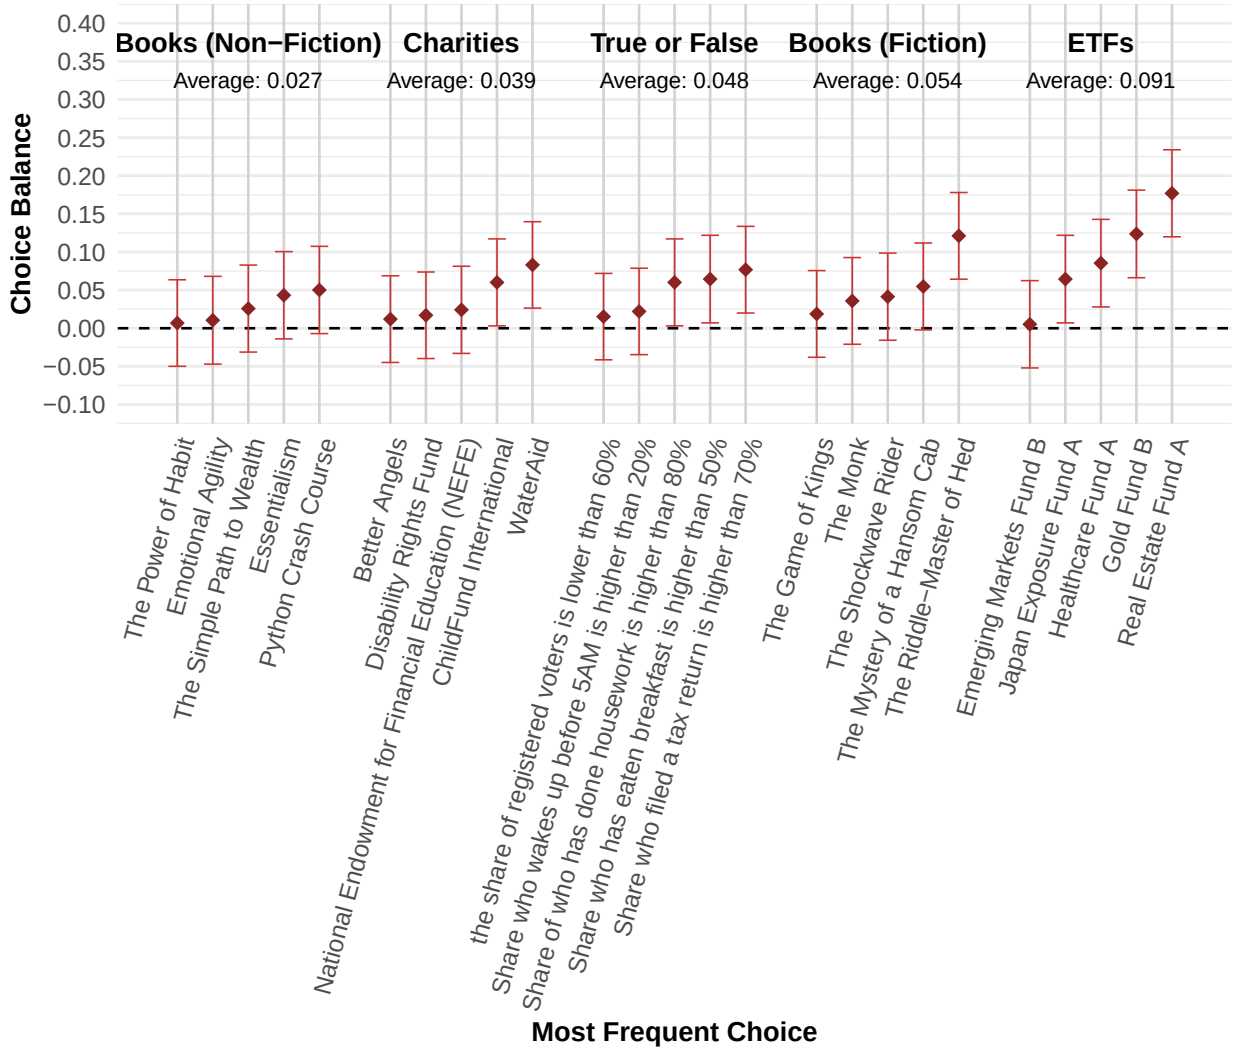



Figure 11: Testing whether Stage 2 participants unconditionally prefer one of the two choice options for each choice question. Showing OLS estimates of a dummy (centered at 0) equal to 1 if the most frequently chosen option among a given choice question gets selected on dummies for the choice questions. The x-axis labels show the most frequently chosen option among a given choice question. Using individual-level clustered standard errors (HC1) to construct asymptotic confidence intervals at the 95% confidence level.

**Background for Figure 12**

Motivation:

- Investigate whether higher **Stage 1 Choice Imbalance** is associated with lower **Stage 2 Resonance Rates**
- If yes, evidence that as a choice option becomes more preferable, the strength of resonance decreases
- If no, evidence that resonance acts above and beyond the informational content of each choice option

How are the Stage 1 (and Stage 2) imbalance rates calculated?

- Use Stage 1 data / recommender data
- Condition on choice pair (i.e. Real Estate Fund A v. Real Estate Fund B) and recommender type (i.e. Democrat)
- Observe which choice option is most frequent (for instance Democrats might choose Real Estate Fund B disproportionately)
- Calculate the probability that the most frequent item is chosen (i.e. the probability that Real Estate Fund B is chosen by Democrats)
- Report the probability as a deviation from 0.5

How are the Stage 2 resonance rates calculated?

- Use Stage 2 data / respondent data
- Condition on
  - choice pair (i.e. Real Estate Fund A v. Real Estate Fund B),
  - Stage 1 recommender characteristics pair (i.e. Democrat v. Republican),
  - and Stage 2 respondent type (i.e. Democrat)
- Calculate resonance rates

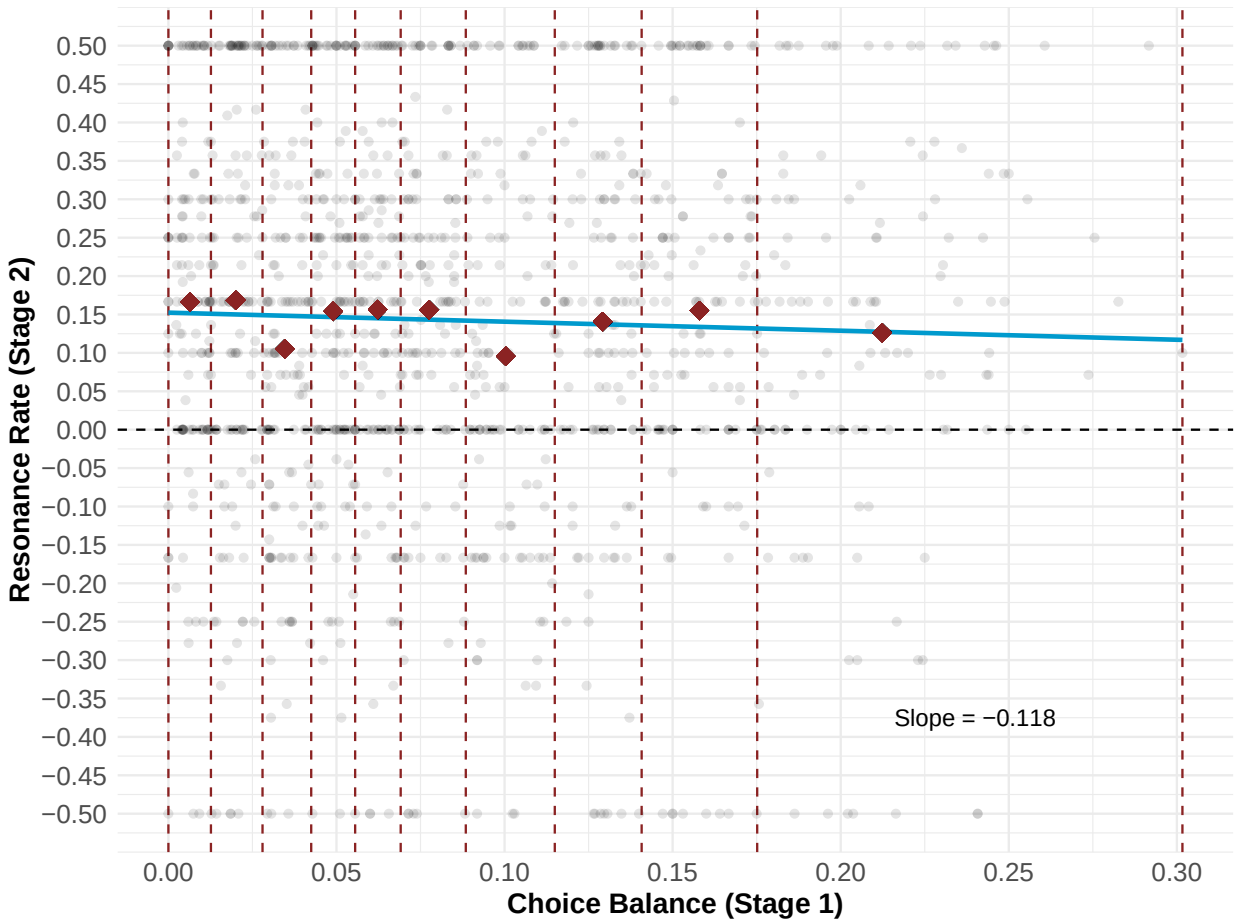



Figure 12: Binned scatter plot of Stage 2 resonance rates on Stage 1 choice imbalance rates at the choice question (i.e. Real Estate Fund A v. Real Estate Fund B) and recommender type (i.e. Democrat) level. Stage 1 imbalance rates are constructed as follows: condition on a choice question and recommender type, observe which choice option is most frequent, calculate the probability that the most frequent item is chosen. Stage 2 resonance rates are constructed as follows: condition on a choice question, Stage 1 recommender characteristics pair, and Stage 2 respondent type, calculate resonance rates.

## A.4 Resonance and Demand for Information

- Resonance rates are on average **lower** for the respondent $i$ choices $c$ sample where participants sought more information on at least one of the two choice options $\Rightarrow$ **selection story**.
- Clicking on the additional information buttons reveals participants who where already more skeptical about the information content of the recommender characteristics.
- However, some variation in the direction of the effect for some recommender characteristics possibly due to statistical noise.

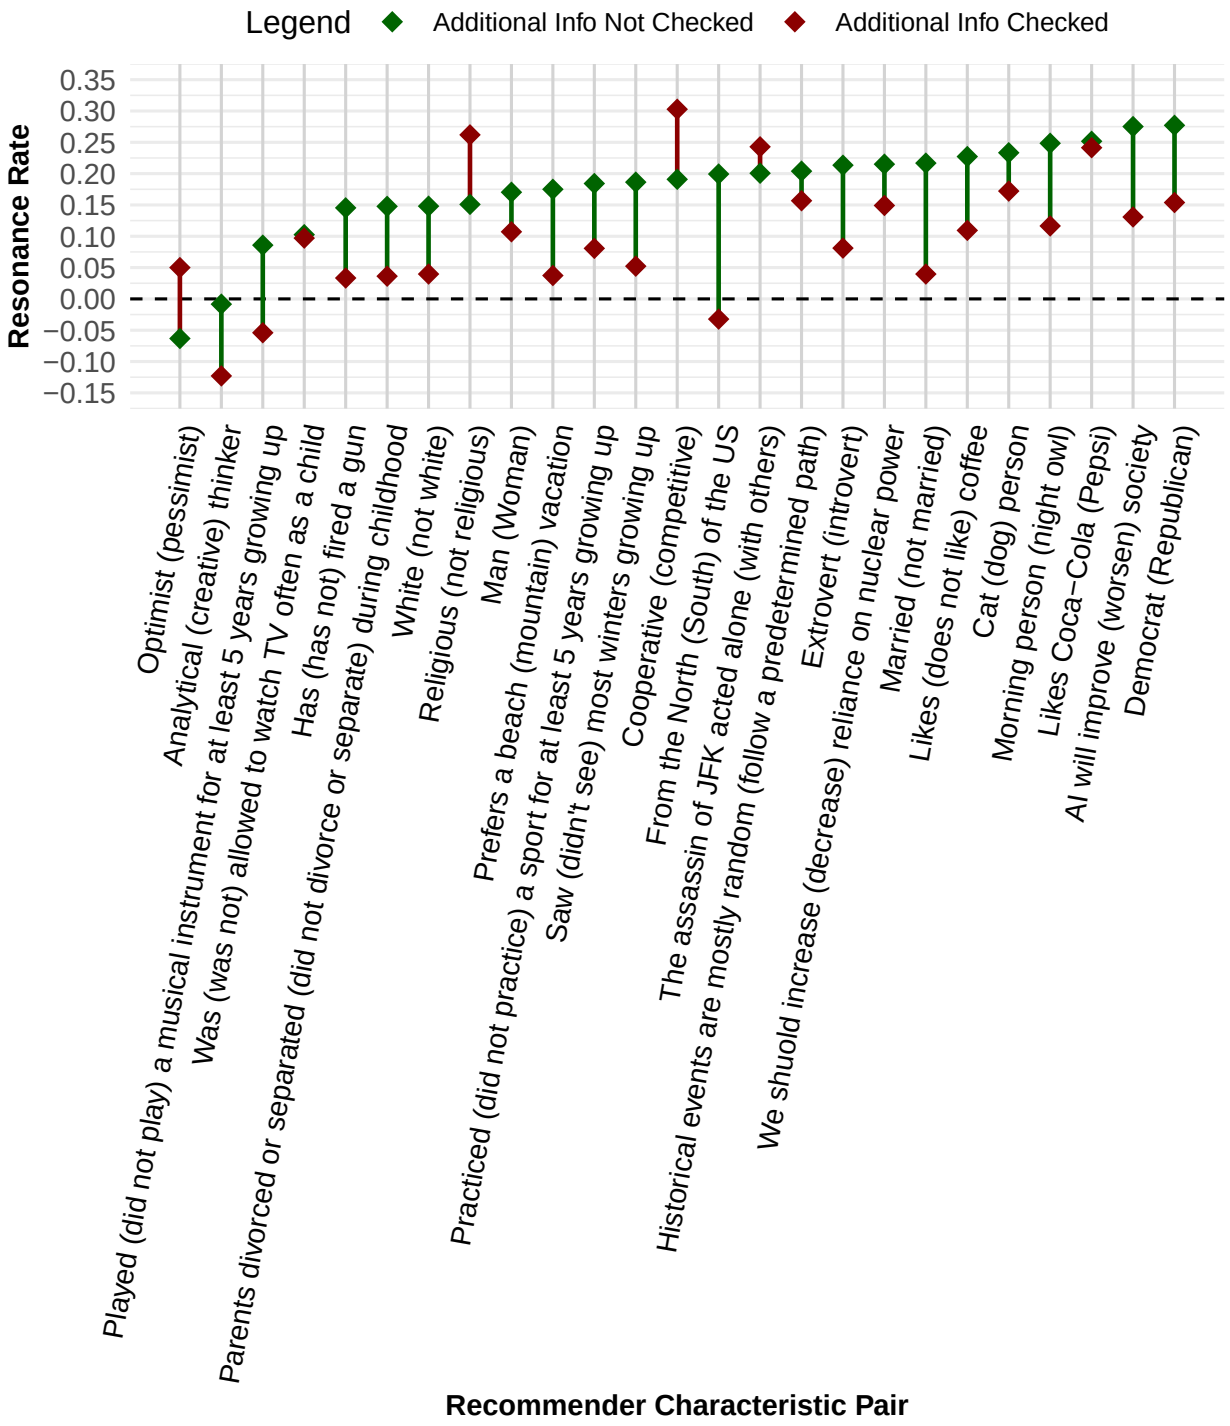



Figure 13: Resonance rates split by information seeking. Conditioning on books, charities, and ETFs choice questions (since the possibility to check additional information was not present for the True or False chocie questions). Red diamonds show resonance rates conditioning on those choices where participants clicked on at least one of the two 'more information' links at the bottom the choice items. Green diamonds show resonance rates calculated on the remaining choices.

# B Appendix: Proofs

## B.1 Proof of Proposition 1

**Proposition 1.** It requires $\tilde{J}$ neighbors in scenario 2, to have an equal probability of getting the nearest community member to switch to choosing $a_1$ at date $t = 2$, as one local leader in scenario 1, where $\tilde{J}$ solves

$$Q_1\tilde{J} + Q_2\tilde{J}^{1/2} + Q_3$$

where $Q_1 = \omega_n(E[z_0|\mathcal{I}_0] - z_1)$, $Q_3 = \sigma_1^{-2}(E[z_0|\mathcal{I}_0] - \mu_1)$, and $Q_2 = \omega_n/\omega_{ll} \cdot Q_3 - Q_1$.

**Proof of Proposition 1**:
*Probability of switching action in scenario 1:* The community members closest to the local leader have distance 1 on the grid. Therefore, the weight these nearest neighbors put on the local leader's experience is

$$\omega_{ll} = 2 - 2\Phi\left(\chi \cdot 1\right).$$

An agent chooses action 1 at time 2 if $E[z_1|\mathcal{I}_1] > E[z_0|\mathcal{I}_1]$. Using the formula for expectations (5), we can rewrite this inequality. In scenario 1, where the closest community member $i$ gets only one signal about $z_1$ from the local leader, the beliefs about $z_1$ at time 1 are the left side of

$$V[z_1|\mathcal{I}_{i1}]\left(\sigma_1^{-2}\mu_1 + \omega_{ll}(z_1 + \epsilon_{1,ll})\right) > E[z_0|\mathcal{I}_1].$$

According to our belief updating rule, the inverse of the conditional variance is the precision of the prior plus the resonance weight (which includes the precision) on the new information: $V[z_1|\mathcal{I}_{i1}]^{-1} = \sigma_1^{-2} + \omega_{ll}$. We can rearrange the inequality as

$$\omega_{ll}(z_1 + \epsilon_{1,ll}) > (\sigma_1^{-2} + \omega_{ll})E[z_0|\mathcal{I}_1] - \sigma_1^{-2}\mu_1.$$

Since beliefs about $z_0$ are precise, we can approximate $E[z_0|\mathcal{I}_1] \approx E[z_0|\mathcal{I}_0]$. This approximation allows us to ignore the random component in beliefs about the well-known action $z_0$.

The only random variable in this expression is the signal noise $\epsilon_{1,ll} \sim N(0, \sigma_\epsilon^2)$. Therefore, the probability that the expression holds is given is the normal cumulative density (cdf)

$$Pr[a_i = 1] = 1 - \Phi\left(\frac{1}{\omega_{ll}\sigma_\epsilon}\left[(\sigma_1^{-2} + \omega_{ll})E[z_0|\mathcal{I}_1] - \sigma_1^{-2}\mu_1 - \omega_{ll}z_1\right]\right).$$

*Probability of switching action in scenario 2:* The community member closest to the the neighboring community members have a distance $\gamma$ on the grid. Therefore, the weight this nearest neighbor puts on the experience of each neighbor is

$$\omega_n = 2 - 2\Phi\left(\chi\gamma\right).$$

In scenario 2, the closest community member gets $\tilde{J}$ signals about action 1, each with weight $\omega_n$. Note that these neighbors must be in an arc, a semicircle or a circle to be equally distant to the community member. This arrangement is not the same as the linearly spaced locations illustrated in the figures. In this scenario, the closest agent chooses action 1 if the belief about $z_1$ on the left side is greater than the right:

$$V[z_1|\mathcal{I}'_{i1}]\left(\sigma_1^{-2}\mu_1 + \omega_n\sum_{j=1}^{\tilde{J}}(z_1 + \epsilon_{1,j})\right) > E[z_0|\mathcal{I}_1].$$

According to our belief updating rule, the inverse of the conditional variance is the precision of the prior plus the sum of resonance weights on all the new information: $V[z_1|\mathcal{I}_{i1}]^{-1} = \sigma_1^{-2} + \tilde{J}\omega_n$. We can rearrange this as

$$\omega_n \sum_{j=1}^{\tilde{J}} (z_1 + \epsilon_{1,j}) > (\sigma_1^{-2} + \tilde{J}\omega_n) E[z_0|\mathcal{I}_1] - \sigma_1^{-2}\mu_1.$$

The only random variables in this expression are the signal noise $\epsilon_{1,j} \sim N(0, \sigma_\epsilon^2)$. What matters is the sum of these independent normal variables. That sum has standard deviation $\tilde{J}^{1/2}\sigma_\epsilon$. We use the same approximation of $E[z_0|\mathcal{I}_1] \approx E[z_0|\mathcal{I}_0]$ as before. Therefore, the probability that the expression holds is given my the normal cumulative density (cdf):

$$Pr[a_i = 1] = 1 - \Phi\left( \frac{1}{\omega_n \tilde{J}^{1/2}\sigma_\epsilon} \left[ (\sigma_1^{-2} + \tilde{J}\omega_n) E[z_0|\mathcal{I}_0] - \sigma_1^{-2}\mu_1 - \tilde{J}\omega_n z_1 \right] \right).$$

*Finding the number of neighboring signals $\tilde{J}$ that equates the probabilities of switching actions.* The two probabilities are equal if the expressions inside the cdf's are equal. Thus, $\tilde{J}$ solves

$$\frac{1}{\omega_{ll}\sigma_\epsilon} \left[ (\sigma_1^{-2} + \omega_{ll}) E[z_0|\mathcal{I}_0] - \sigma_1^{-2}\mu_1 - \omega_{ll} z_1 \right] = \frac{1}{\omega_n \tilde{J}^{1/2}\sigma_\epsilon} \left[ (\sigma_1^{-2} + \tilde{J}\omega_n) E[z_0|\mathcal{I}_0] - \sigma_1^{-2}\mu_1 - \tilde{J}\omega_n z_1 \right]$$

Simplifying yields the expression in proposition 1.

## B.2 Proof of Proposition 2

**Proposition 2.** Celebrity outcomes have more influence on others' beliefs when agents are geography -constrainted ($\partial E[z_a|\mathcal{I}]/\partial z_{celeb}$ larger).

**Proof of Proposition 2**:
*Step 1:* Derive $\partial E[z_a|\mathcal{I}]/\partial a_{celeb}$ without the geographic constraint.

To form conditional expectations, we use our belief-formation rule (5). When an agent as access to the geographically-unconstrained information set $\mathcal{I}_{it}^U$, with $N_U$ regular observations and 1 observed celebrity action, their beliefs are

$$\begin{aligned} E[z_k|\mathcal{I}_{it}^U] \quad &= Var[z_k|\mathcal{I}_{it}^U] \Big\{ \sigma_k^{-2}\mu_k + \tilde{\psi}_{i,celeb} 1_{a_{celeb,t}=a_k} \omega_{i,celeb} (z_{celeb} + \epsilon_{celeb,t}) \\ &\qquad + \textstyle\sum_{j=1}^{N_U} \sum_{t'=1}^{t} \tilde{\psi}_{ij} 1_{a_{jt}=a_k} \omega_{ij} (z_k + \epsilon_{kit}) \Big\}. \end{aligned} \tag{11}$$

Notice that this is a linear function of the celebrity action payoff, $(z_{celeb} + \epsilon_{celeb,t})$. The partial derivative $\partial E[z_a|\mathcal{I}^U]/\partial a_{celeb}$ is simply the terms that multiply that payoff, namely, $Var[z_k|\mathcal{I}_{it}^U] 1_{a_{celeb,t}=a_k} \omega_{i,celeb}$.

Importantly, the conditional variance term at the start of the partial derivative is given by (**??**). With $N_U$ regular observations and 1 observed celebrity action, the conditional variance of agent $i$'s belief is

$$Var[z_k|\mathcal{I}_{it}^U] = \left( \sigma_k^{-2} + 1_{a_{celeb,t}=a_k} \omega_{i,celeb} \sum_{j=1}^{N_U} \sum_{t'=1}^{t} 1_{a_{jt}=a_k} \omega_{ij} \right)^{-1}$$

*Step 2:* Derive $\partial E[z_a|\mathcal{I}]/\partial a_{celeb}$ with the geographic constraint.

Using the same belief updating rule for an agent who has a geographically-constrained information set $\mathcal{I}_{it}^C$, with $N_C < N_U$ regular observations and 1 observed celebrity action, the belief is still Notice that this is a linear function of the celebrity action payoff, $(z_{celeb} + \epsilon_{celeb,t})$. The partial derivative $\partial E[z_a|\mathcal{I}^C]/\partial a_{celeb}$ is simply the terms that multiply that payoff. In this case, those terms are $Var[z_k|\mathcal{I}_{it}^C]1_{a_{celeb,t}=a_k}\omega_{i,celeb}$. Notice that the latter terms are identical. The only difference in the influence of celebrity actions comes from the conditional variance. For the geographically-constrained agent, that variance is

$$Var[z_k|\mathcal{I}_{it}^C] = \left(\sigma_k^{-2} + 1_{a_{celeb,t}=a_k}\omega_{i,celeb}\sum_{j=1}^{N_C}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}\omega_{ij}\right)^{-1} \tag{12}$$

*Step 3:* Compare $\partial E[z_a|\mathcal{I}]/\partial a_{celeb}$ in both cases.

The partial derivatives in the two cases differ only by their conditional variances. To facilitate comparison, consider the conditional precisions, the inverse of the variances:

$$Var[z_k|\mathcal{I}_{it}^U]^{-1} = \sigma_k^{-2} + 1_{a_{celeb,t}=a_k}\omega_{i,celeb}\sum_{j=1}^{N_U}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}\omega_{ij} \tag{13}$$

We can break the sum of $N^U$ nodes into the $N^C$ nodes that are inside the geographic constraint and the $N^U - N^C$ nodes that are not:

$$Var[z_k|\mathcal{I}_{it}^U]^{-1} = \sigma_k^{-2} + 1_{a_{celeb,t}=a_k}\omega_{i,celeb}\sum_{j=1}^{N_C}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}\omega_{ij} + \sum_{j=N_C+1}^{N_U}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}\omega_{ij} \tag{14}$$

$$Var[z_k|\mathcal{I}_{it}^U]^{-1} = Var[z_k|\mathcal{I}_{it}^C]^{-1} + \sum_{j=N_C+1}^{N_U}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}\omega_{ij} \tag{15}$$

Since the indicator function is non-negative and the weight $\omega_{ij} > 0$, we have $Var[z_k|\mathcal{I}_{it}^U]^{-1} \geq Var[z_k|\mathcal{I}_{it}^C]^{-1}$, with a strict inequality, whenever $\sum_{j=N_C+1}^{N_U}\sum_{t'=1}^{t} 1_{a_{jt}=a_k}$, which means that at least one agent outside the geographic constraint zone took action $k$.

Since $Var[z_k|\mathcal{I}_{it}^U]^{-1} \geq Var[z_k|\mathcal{I}_{it}^C]^{-1}$, $Var[z_k|\mathcal{I}_{it}^U] \leq Var[z_k|\mathcal{I}_{it}^C]$. Since this was the only term that differed between the derivatives and the other terms were positive, $\partial E[z_a|\mathcal{I}^C]/\partial a_{celeb} > \partial E[z_a|\mathcal{I}^U]/\partial a_{celeb}$.

### B.3 Proof: Labor Market Discrimination

The boss updates their belief about the unknown state $\theta$ using signals from workers:

$$s_i = \theta + \xi_i, \quad \xi_i \sim N(0, \sigma_i^2).$$

The boss's posterior belief about $\theta$, given the signals received, follows a weighted average:

$$E[\theta \mid I_j] = \frac{\sum_i w_i s_i}{\sum_i w_i},$$

where the weight $w_i$ is inversely proportional to the worker's signal variance:

$$w_i = \frac{1}{\sigma_i^2}.$$

Under the resonance mechanism, workers whose past signals had a greater impact on the firm's decisions receive increased weight in future decision-making. That is, if worker $i$ has been historically influential, their effective weight is amplified:

$$w_i' = \rho w_i,$$

where $\rho > 1$ is a resonance factor that depends on past impact.

**Impact on the Chosen Action $X$.** The boss chooses an action to minimize expected squared loss:

$$X^* = E[\theta \mid I_j] = \frac{\sum_i w_i' s_i}{\sum_i w_i'}.$$

The impact of worker $i$'s signal on $X^*$ is:

$$\frac{\partial X^*}{\partial s_i} = \frac{w_i'}{\sum_j w_j'}.$$

A worker is **more likely to be promoted** if their signal has the greatest impact, i.e., if $\frac{\partial X^*}{\partial s_i}$ is maximized.

**Type Dependence in Weights.** Workers belong to two types, $a$ and $b$, where:

- **Type** $a$ workers have signal precision $\sigma_a^2$.
- **Type** $b$ workers have signal precision $\sigma_b^2$.
- Initially, suppose $\sigma_a^2 = \sigma_b^2$, so both types have equal baseline weights.

Under the resonance mechanism, suppose type $b$ workers historically received greater weight (due to small early differences or stochastic noise). Then, in the next period:

$$w_b' = \rho_b w_b, \quad w_a' = \rho_a w_a.$$

Since weight accumulation is self-reinforcing ($\rho_b > \rho_a$), type $b$ workers' signals have a larger effect on $X^*$, leading to:

$$\frac{\partial X^*}{\partial s_b} > \frac{\partial X^*}{\partial s_a}.$$

**Promotion Probability.** Since the boss promotes the worker whose signal had the **largest effect on** $X^*$, the probability that a type $b$ worker is promoted is:

$$P(\text{Promotion} \mid \tau_i = b) = P\left(\frac{\partial X^*}{\partial s_b} > \frac{\partial X^*}{\partial s_a}\right).$$

Given that $w_b' > w_a'$, we conclude that the probability of promotion of type $b$ is higher:

$$P(\text{Promotion} \mid \tau_i = b) > P(\text{Promotion} \mid \tau_i = a).$$

# C Simulation Details

There are three sets of simulations in the results section.

The true distributions of the action payoffs are $z_0 \sim N(5, 100)$ with cost 0 for action 0, and $z_1 \sim N(10, 100)$ with cost 1 for action 1. The individual payoff shocks are distributed $\epsilon_{ait} \sim N(0, 100)$ for both actions, independently drawn for actions, agents, and time. Prior beliefs about the payoffs are $z_0 \sim N(3, 10)$ for action 0, meaning that an agent has prior belief that $E[z_0] = 2.5$, with prior variance 10; and $z_1 \sim N(1.5, 10)$ for action 1.

This simulation is run and plotted in characteristic space, not geographic space. The words location and distance here refer to the location in characteristic space and the distance between such characteristic locations.

**Panel a**: Every node on the (20 x 20) grid has a location and a randomly drawn type. Note that this implies a uniform distribution of characteristics at each location.

**Panel b**: Social distance is the Euclidean distance between nodes/agents on the characteristic grid.

**Digital Social Networks and the Decline of Experts (Figure 2)** In this exercise, there are no choices or payoffs, only agents with geographic locations and characteristics.

Every node should have a 4-tuple of characteristics: 2 geographical coordinates (x and y-axis) and 2 social coordinates, call those coodinates v and w.

Start by placing one node at each vertex in a 20x20 lattice. There should be 400 nodes with x and y coordinates evenly spaced between 1 and 20. For each node, assign a random value to v and w, drawn from a normal distribution: $v \sim N(10, 5)$ independent of $w \sim N(10, 5)$, for each of the 400 nodes. These draws are also independent across nodes.

Then choose one reference node at the center: $(10, 10)$ location means $x = y = 10$. Set this node's social values to $v = w = 10$. That makes the geographically central node socially central as well.

The geographic figure is a color box plot where the color of each node is the euclidean distance between (10,10) and (v,w). So, location is governed by geographic (x,y) and color is determined by social distance v,w.

The 10x10 box in the center of this plot (around boxes 6-15 on both axes) represent the set of community members that an individual can communicate with when they are geographically constrained. The becomes relevant for the results with a digital community.

Next, we take the subset of nodes in that inner 10x10 box and plot them on the 20x20 $(v, w)$ axis. As before, this is a color box grid. There will be lots of empty nodes on this grid. They show up as white space.

A third color box plot, again on the v,w axis, includes all the 400 original nodes, not just the ones in the 10x10 box.

Finally, we repeat the same exercise, but make the central node ($(x, y) = (10, 10)$) have outlier social characteristics. Set $v = w = 2$. This places that node in the minority of the social spectrum.

We construct three more color box plots, following the same procedure as above. Colors now reference the social distance of nodes from the new $(v, w) = (2, 2)$ reference point.

# D Evidence on Role Models

Role models, those other feel close to or have resonance with, play a key role in public health campaigns. For example, we saw during the COVID pandemic that famous people recommending vaccines did not work well. The fact that most others had no adverse effects also seemed to have little influence. However, having a local leader, like a priest or community doctor, recommend the same action was more effective. NPR reported about the efforts of to "keep an open line of communication with recalcitrant rabbis ... knowing they would have the most influence to convince their followers to get vaccinated," and the ultimate success of this approach: "The rabbis agreed to be vaccinated — and their adherents followed suit" (April 22, 2021). The Australian ABC Network describes how "the knowledge and trust Aboriginal health organizations had with local communities was another reason for Victoria's success" in convincing the Aboriginal population, in particular by "re-framing government messaging for local communities" (August 19, 2021). As for the U.S., the PR Newswire explains, "With the disproportionate impact of the pandemic in rural areas across the country, it's critically important that Americans in rural communities hear these powerful stories from each other, their neighbors" (October 19, 2021). The Villager report that "Union leaders receive COVID-19 vaccine in Harlem, hoping to inspire others to do the same" (February 27, 2021). And the New York Times attributes the narrowed vaccination gap between Black and other racial groups to "decisions made in many states to send familiar faces to knock on doors and dispel myths about the vaccines' effectiveness," such as the case of a store owner and county commissioner in Panola, a small rural town near the Mississippi border, who "led the effort to vaccinate nearly all of her majority Black community," and pro-vaccination campaigns including descendants of participants in the Tuskegee study (October 13, 2021). The campaign successfully eased vaccine hesitancy derived from medical mistrust by positioning community members who had directly suffered from the betrayals of the healthcare industry as trusted informants. The Ad Council noted that "[the descendants] are so perfect to be the storytellers because if they can move past it and see the lessons and drive towards getting people vaccinated, it feels like everyone should be able to do the same."

Story after story spells out how, when local leaders provide resonant information and model pro-vaccination behavior, this influences a segment of the population that had heard, but did not internalize other pro-vaccine information. The effectiveness of information originating from community members is so widely recognized that it shows up as a key strategy in the Center for Disease Control's official Vaccination Field Guide. `www.ccbh.net/wp-content/uploads/2021/09/Field-Guide_CDC_Final_2021.pdf`

# E Resonance as a Foundation for Labor Market Discrimination

We now extend the model, by introducing a labor market in which firms promote workers based on their impact on the organization. It generates systematic differences in promotion probabilities across worker types, even with a type-neutral promotion criterion and without preference-based bias.

Consider a labor market with a single firm, consisting of one boss and $n$ employees. The firm's objective is to take an action $X$ that maximizes expected payoffs, given an unknown state of the world $\theta \in \mathbb{R}$. The firm receives signals from employees and updates its beliefs using a resonance-based learning process.

**Information Structure.** Each employee $i$ observes a private signal about a state $z$ for which all agents have flat priors:

$$s_i = z + \xi_i, \quad \xi_i \sim N(0, \sigma_i^2), \tag{16}$$

where $\sigma_i^2$ represents the precision of worker $i$'s information. A richer theory might have a boss learn about the worker's knowledge. But we assume $\sigma_i^2$ is common knowledge for simplicity. Employees truthfully communicate their signals to the boss. A worker with a lower $\sigma_i^2$ transmits more precise, more useful input to their boss.

Each employee belongs to one of two groups, indexed by $\tau_i \in \{a, b\}$. We assume that the distribution of signal precisions within each group is identical.

**Boss's Decision Rule.** The boss seeks to choose an action $x$ that minimizes expected quadratic loss:

$$\mathbb{E}\left[(x - z)^2 \mid I_B\right], \tag{17}$$

where $I_B$ is the boss's information set. Given the signals $s_1, \ldots, s_n$, the boss updates beliefs using a weighted average:

$$\mathbb{E}[z \mid I_B] = \frac{\sum_i w_i s_i}{\sum_i w_i}, \tag{18}$$

where $w_i = \rho_{ib}\sigma_i^{-2}$ represents the resonance-adjusted weight assigned to worker $i$'s signal. As before, that weight depends on information precision $\sigma_i^{-2}$, but also on the social distance through $\rho_{ib}$. For $b$ types, the social distance with the boss is zero and $\rho_{ib} = 1$ if $\tau_i = b$. For $a$ types, there is distance in type space and $\rho_{ib} < 1$ if $\tau_i = a$.

**Promotion Mechanism.** After choosing $x$, the boss promotes one employee to replace them in the next period. Promotion is based on impact, defined as the worker whose signal had the greatest effect on the chosen action:

$$i^* = \arg\max_i \left| \frac{\partial x}{\partial s_i} \right|. \tag{19}$$

Impact criteria are used widely in promotion policies, including academic promotions, which typically consider impact in the form of citation counts and explicit requests to comment on impact in promotion letters.

**Implications for Labor Market Discrimination** Resonance here implies that the boss assigns additional weight to messages that come from workers who share characteristics with the boss. For those similar workers, their messages receive more weight in belief updating. Since resonance increases the weight placed on past influential messages, workers whose signals were previously more impactful are more likely to be promoted. Over time, this dynamic reinforces the influence of certain types of employees. This happens, even though the promotion policy is type-neutral and no agents have a type preference for promotion.

A key result of the model is that type $b$ workers have a higher probability of promotion than type $a$ workers, even when the two groups have identical signal distributions. Proof in Appendix B.3.

**Proposition 3.** ***Workers with leaders' types get promoted more frequently.*** $P(Promotion \mid \tau_i = b) > P(Promotion \mid \tau_i = a)$.

Thus, even in the absence of explicit bias or productivity differences, resonance can generate systematic disparities in promotion rates across worker types by making same-type workers more influential. This suggests a novel mechanism through which structural disadvantage can emerge in the labor market.

This model provides a foundation for understanding how resonance-based updating can create persistent labor market disparities. Unlike traditional models of discrimination, which rely on taste-based preferences or statistical differences, this framework shows that endogenous dynamics in information transmission can lead to persistent inequality.

# F Labor Market Analysis: Data, Variable Construction and Additional Statistics

## F.1 ACS Data and Variables

Micro data was obtained via IPUMS and comes from the US Census Bureau's annual American Community Survey (ACS) years 2005 to 2020. The analysis uses IPUMS variables on age, occupation, education, migration, and ancestry.

IPUMS was formerly an acronym for Integrated Public Use Microdata Series prior to a 2016 change in branding that uses IPUMS name to describe additional projects.

IPUMS variable OCC1990 provides consistent occupation codes across years based on the occupation categorizations from the 1990 decennial census. This variable categorizes 338 unique occupations (including "unemployed"). Using IPUMS's broader categorizations of these occupations, the variable is aggregated into 33 larger occupation groups. Taking out those who are unemployed (since our goal is to measure what occupation a person chooses), we are left with 32 occupation groups on which to do the analysis.

IPUMS variable ANCESTR1 provides the respondent's self-reported ethnic origin and contains 249 unique categories. Removing unknown and other uninterpretable responses,[23] remaining ancestries were grouped into 21 ethnic groups based on global geographic regions as they are defined by the UN Statistical Division and separated North America into three categories based on race, North American (primarily white North American), African American, and Indigenous American. Smaller groups were then combined to create 17 final ethnic groups; South and Central Asia were combined into one region, West Asia and North Africa were combined into "Middle East/North Africa," and Melanesia, Micronesia, and Polynesia were combined into "Pacific Islander."

We include categories "Sub-Saharan African" and "North African" for respondents who identify with African ancestry, however, many Black Americans do not identify with these categories since they've experienced no recent migration in their family. Since ancestry was self-reported, we get the most accurate measure of which community respondents feel most connected to. For instance, even if a Black American has parents from Africa, they may identify more with the Black community in the US than with the African immigrant community. In this case, the respondent would likely choose to identify as "African American" in the survey. This serves our analysis because we want to isolate how people react to the information they receive from groups they identify with.

IPUMS constructed variable CPUMA0010 provides consistent codes for PUMAs across years by matching PUMAs from the 2010 census to corresponding PUMAs from the 2000 census with a 1% mismatch tolerance. Using this variable we were able to identify the PUMA of residence for each individual in the sample across all 15 years of data.

IPUMS variable PERWT provides individual population weights for each person in the sample. This variable describes how many people in the population each particular observation in the sample represents.

IPUMS variable SCHOOL indicates if the respondent was currently attending school. We use this variable to construct a younger cohort of 18–22-year-olds who are not currently attending

[23] Uninterpretable responses include categories that IPUMS has dubbed as "residual," including "mixture," "uncodable," "other," and "not reported," as well as categories that are not descriptive enough to make conculsions about the ethnicity of the respondent ("Eurasian," "Hispanic," "Asian," "Spanish American" and "European, nec)." Removing these observations reduces the number of 26-65 year olds in the sample from 16,516,288 to 16,060,861 and the number of 18-22 year olds in the sample from 417,837 to 405,810.

school. We include this exclusion in order to ensure that the young people in our analysis are in the occupations they chose, rather than enrolled in degree programs working towards the occupations they would like.

IPUMS variable MIGRATE1D provides information on whether and where respondents have moved in the year prior to the survey. We use this variable to maximize the likelihood that the young people in our sample currently live in the place where they grew up. This allows us to capture whether they are learning from those in their communities. Therefore, we look only at 18-22-year-olds who have not moved out of their PUMA in the last year.

### F.2 Constructed Variables

Using a combination of these variables, we measure the fraction of each ethnic group in each occupation in each PUMA each year and construct the variable *Ethnic Ratio*. We do so by taking the number of people of primary working age (which we defined as 26-65) in each of these occupation/ethnic/PUMA/year groups and dividing it by the number of people of the ethnic group in the PUMA in the same year. We weight the estimates appropriately using nationally representative weights provided by IPUMS.

Let $i$ be an ethnic group, $j$ be an occupation, $k$ be a PUMA, and $t$ be a year. Then the number of people in an ethnic/occupation/PUMA/year group is $Occupation\ Size_{i,j,k,t}$ and the total number of people of an ethnic group in a PUMA each year is $Population\ Size_{i,j,k,t}$. Therefore, we have:

$$Ethnic\ Ratio_{i,j,k,t} = \frac{Occupation\ Size_{i,j,k,t}}{Population\ Size_{i,k,t}}$$

All estimates are weighted using the PERWT variable to get the most accurate population estimates of these fractions. Similarly, we construct the variable *Local Ratio* as the fraction of all people in the PUMA in each occupation each year.

$$Local\ Ratio_{j,k,t} = \frac{Occupation\ Size_{j,k,t}}{Population\ Size_{k,t}}$$

*Ethnic ORep* is the difference between *Ethnic Ratio* and *Local Ratio* and it measures the extent to which the ethnic group is overrepresented in an occupation in each PUMA each year.

$$Ethnic\ ORep_{i,j,k,t} = Ethnic\ Ratio_{i,j,k,t} - Local\ Ratio_{j,k,t}$$

Next, we construct the variable *Youth Rate* similarly to *Ethnic Ratio* but for the 18–22-year-old cohort. *Youth Rate* measures the fraction of 18-22-year-olds of each ethnic group in each occupation in each PUMA each year.

$$Youth\ Rate_{i,j,k,t} = \frac{Occupation\ Size_{i,j,k,t}}{Population\ Size_{i,k,t}}$$

We add gender re-estimate the variables *Ethnic Ratio* and *Ethnic ORep* to be gender-specific. Here, *Ethnic Ratio (Same Gender)* measures the fraction of people aged 26-65 of ethnicity $i$ and gender $l$ in occupation $j$ and PUMA $k$ during year $t$.

$$Ethnic\ Ratio\ (Same\ Gender)_{i,l,j,k,t} = \frac{Occupation\ Size_{i,l,j,k,t}}{Population\ Size_{i,l,k,t}}$$

*Ethnic ORep (Same Gender)* now measures the overrepresentation of an ethnic-gender group in an occupation in their PUMA. For instance, how many more Vietnamese women are nail technicians compared to all nail technicians in New York.

$$\textit{Ethnic ORep (Same Gender)}_{i,l,j,k,t} = \textit{Ethnic Ratio (Same Gender)}_{i,l,j,k,t} - \textit{Local Ratio}_{j,l,k,t}$$

We then create *Ethnic ORep (Wrong Gender)*, which is the the degree to which elders of the *same ethnic group* but *opposite gender* are overrepresented in an occupation.

$$\textit{Ethnic ORep (Wrong Gender)}_{i,-l,j,k,t} = \textit{Ethnic Ratio (Wrong Gender)}_{i,-l,j,k,t} - \textit{Local Ratio}_{j,-l,k,t}$$

By including this variable in the regression, we show that young people are not just following the lead of elders of their same ethnicity, but that gender plays an important role as well. This also provides evidence that information access is not the only mechanism at play in determining occupation choices. It is highly likely that people recieve information from community members of both genders, however, we find that they respond to the actions of their same-gender elders.

We construct additional variables for the purpose of analyzing the impact of an economic shock. Using the number of people in each occupation in each PUMA each year, we calculate the difference in the number of people in the occupation in 2006 versus the number of people in the occupation in 2010. We use the reduction in the number of people employed in an occupation from 2006 to 2010 as a proxy for "layoffs." We then construct the variable *High Layoffs*. We take the change in the number of people in the occupation/PUMA from 2006 to 2010 and divide it by the number of people in the occupation/PUMA in 2006, giving the percent reduction in employment for that occupation.

$$\textit{Job Loss}_{j,k} = \textit{Occupation Size}_{j,k,2006} - \textit{Occupation Size}_{j,k,2010}$$

$$\textit{Percent Layoffs}_{j,k} = \frac{\textit{Job Loss}_{j,k}}{\textit{Occupation Size}_{j,k,2006}}$$

The variable *High Layoffs* is then defined to equal 1 if $\textit{Percent Layoffs}_{j,k}$ is in the top 30% and 0 otherwise.

Next, we fix the *Ethnic ORep* variable to $\textit{Ethnic ORep}_{2005}$, which is the average of the *Ethnic ORep* in 2005 and 2006. This variable represents the overrepresentation of the ethnic group in an occupation before the recession. Fixing the variable to a past value allows us to see how young people respond to what they observed *before* making their occupation choice. Using the fixed values, we also run the regression separately for ethnic groups that were overrepresented ($\textit{Ethnic ORep}_{2005}$ was positive) and ethnic groups that were underrepresented ($\textit{Ethnic ORep}_{2005}$ was negative). We see similar results among overrepresented groups, with those who were overrepresented in 2005 turning away from occupations that experienced severe layoffs after the Recession.

### F.3 Labor Market Evidence: Summary Statistics

To construct the baseline estimates of the *Ethnic Ratio* and *Local Ratio* variables, we use the ACS data from 2005 to 2020 including all employed people ages 26 to 65 for which we had occupation and ethnicity information. This data totals 16,060,861 observations. In order to ensure reliable estimates, we required that at least 100 observations of an ethnic group be present in a PUMA in order to include it in the data. With this restriction, one ethnic group, Australians/New Zealanders, drops out of the dataset and we are left with 16 unique ethnicities. We observe a total 1,054 PUMAs

in the data. We observe all 32 occupations for each ethnic group.[24] Table A1 shows the average number of observations of each ethnicity in each occupation/PUMA/year group. Table A2 shows the average number of people in each ethnicity represented by the data. These are the weighted estimates of the number of people in each occupation/PUMA/year group. Tables A3 and A4 provide summary statistics of the constructed variables. Table A3 contains values for the variables dependent on ethnicity only, while Table A4 provides the gender-specific values.

Table A1: Number of Observations of Each Ethnic Group ($Occ\ Size_{i,j,k,t}$ (unweighted))

| Ethnicity | Average | Standard Deviation | 25th Percentile | Median | 75th Percentile | Maximum |
|---|---|---|---|---|---|---|
| Eastern European | 6.11 | 11.15 | 1 | 3 | 7 | 304 |
| Northern European | 9.63 | 25.14 | 1 | 3 | 9 | 1536 |
| Western European | 9.72 | 22.52 | 1 | 4 | 10 | 773 |
| Southern European | 6.82 | 16.69 | 1 | 3 | 7 | 816 |
| Pacific Islander | 6.41 | 10.88 | 0 | 2 | 8 | 81 |
| South Asian | 5.95 | 11.069 | 0 | 2 | 7 | 139 |
| East Asian | 6.33 | 12.74 | 0 | 2 | 7 | 263 |
| Southeast Asian | 6.46 | 12.89 | 0 | 2 | 7 | 233 |
| Middle Eastern/ North African | 5.46 | 8.06 | 0 | 3 | 7 | 56 |
| Latin American/ Caribbean | 7.32 | 19.36 | 0 | 2 | 6 | 392 |
| Central American | 12.67 | 33.50 | 0 | 2 | 11 | 584 |
| South American | 6.41 | 11.78 | 0 | 2 | 7 | 127 |
| North American | 8.33 | 19.40 | 1 | 3 | 8 | 606 |
| Sub-Saharan African | 4.96 | 7.56 | 0 | 2 | 6 | 56 |
| African American | 9.27 | 24.61 | 0 | 2 | 8 | 574 |
| Indigenous American | 6.79 | 13.13 | 0 | 2 | 7 | 186 |

[24] At least one ethnicity is represented in each occupation, however, some ethnicities do not have any members in some occupations. For such ethnicity/occupation pairs the fraction of people in the occupation is equal to zero.

Table A2: Number of People in Each Ethnic Group ($Occ\ Size_{i,j,k,t}$ (weighted))

| Ethnicity | Average | Standard Deviation | 25th Percentile | Median | 75th Percentile | Maximum |
|---|---|---|---|---|---|---|
| Eastern European | 583.41 | 1098.45 | 50 | 236 | 659 | 26336 |
| Northern European | 909.93 | 2394.16 | 77 | 315 | 889 | 145700 |
| Western European | 903.52 | 2063.67 | 75 | 317 | 914 | 69517 |
| Southern European | 689.62 | 1697.21 | 57 | 246 | 719 | 80291 |
| Pacific Islander | 574.54 | 1009.54 | 0 | 146 | 653 | 8289 |
| South Asian | 673.74 | 1273.79 | 0 | 193 | 744 | 17078 |
| East Asian | 639.25 | 1296.60 | 0 | 216 | 681 | 22937 |
| Southeast Asian | 688.45 | 1381.21 | 0 | 240 | 730 | 24692 |
| Middle Eastern/ North African | 619.10 | 944.75 | 0 | 266 | 759.25 | 7147 |
| Latin American/ Caribbean | 989.91 | 2589.06 | 0 | 245 | 863 | 49535 |
| Central American | 1663.67 | 4534.18 | 0 | 287 | 1390 | 89952 |
| South American | 874.75 | 1591.111 | 0 | 246 | 951.25 | 15722 |
| North American | 773.17 | 1823.82 | 38 | 243 | 778 | 56980 |
| Sub-Saharan African | 732.84 | 1194.12 | 0 | 282 | 853 | 11622 |
| African American | 1212.63 | 3218.18 | 0 | 297 | 1057 | 77562 |
| Indigenous American | 594.63 | 1117.05 | 0 | 195.5 | 659.25 | 18394 |

Table A3: Summary of Constructed Variables (Ethnicity)

| | N | Mean | SD | 25th | Median | 75th | Max |
|---|---|---|---|---|---|---|---|
| $Youth\ Rate_{i,j,k,t}$ | 1,203,328 | 0.031 | 0.118 | 0.000 | 0.000 | 0.000 | 1.000 |
| $Youth\ ORep_{i,j,k,t}$ | 1,203,328 | 0.000 | 0.090 | 0.000 | 0.000 | 0.000 | 0.999 |
| $Ethnic\ Ratio_{i,j,k,t}$ | 1,203,328 | 0.031 | 0.046 | 0.003 | 0.014 | 0.039 | 0.566 |
| $Ethnic\ ORep_{i,j,k,t}$ | 1,203,328 | -0.000 | 0.019 | -0.007 | -0.001 | 0.005 | 0.381 |
| $Ethnic\ ORep_{i,j,k,t-1}$ | 1,122,624 | -0.000 | 0.019 | -0.007 | -0.001 | 0.005 | 0.417 |
| $Local\ Ratio_{j,k,t}$ | 1,203,328 | 0.031 | 0.042 | 0.006 | 0.015 | 0.039 | 0.400 |
| $Ethnic\ ORep_{2005}$ | 1,203,328 | -0.000 | 0.017 | -0.005 | -0.000 | 0.004 | 0.337 |
| *Percent Layoffs* | 1,035,800 | -0.172 | 1.176 | -0.277 | -0.004 | 0.220 | 1.000 |
| *Job Loss* | 1,098,528 | -63.336 | 1779.922 | -489.000 | -3.000 | 403.000 | 17714.000 |

Table A4: Summary of Constructed Variables (Ethnicity, Gender)

| | N | Mean | SD | 25th | Median | 75th | Max |
|---|---|---|---|---|---|---|---|
| *Youth Rate*$_{i,l,j,k,t}$ | 2,108,064 | 0.031 | 0.136 | 0.000 | 0.000 | 0.000 | 1.000 |
| *Youth ORep*$_{i,l,j,k,t}$ | 2,108,064 | 0.000 | 0.103 | 0.000 | 0.000 | 0.000 | 0.999 |
| *Ethnic Ratio*$_{i,l,j,k,t}$ | 2,108,064 | 0.031 | 0.054 | 0.000 | 0.009 | 0.037 | 0.652 |
| *Ethnic ORep (Same Gender)*$_{i,l,j,k,t}$ | 2,108,064 | -0.000 | 0.024 | -0.007 | -0.000 | 0.005 | 0.441 |
| *Ethnic ORep (Same Gender)*$_{i,l,j,k,t-1}$ | 1,964,544 | -0.000 | 0.024 | -0.007 | -0.000 | 0.005 | 0.464 |
| *Local Ratio (Same Gender)*$_{i,l,j,k,t}$ | 2,108,064 | 0.031 | 0.048 | 0.003 | 0.012 | 0.038 | 0.560 |
| *Ethnic ORep (Same Gender)*$_{2005}$ | 2,108,064 | -0.000 | 0.020 | -0.005 | -0.000 | 0.004 | 0.389 |
| *Percent Layoffs* | 1,833,188 | -0.170 | 1.149 | -0.274 | -0.004 | 0.218 | 1.000 |
| *Job Loss* | 1,941,440 | -64.108 | 1803.196 | -494.000 | -3.000 | 407.000 | 17714.000 |

Table A5: Occupation Choice: Ethnicity, 1-year Lag

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| Ethnic ORep (Lag) | 0.533*** | 0.152*** | 0.151*** | 0.207*** | 0.148*** | 0.156*** | -0.00914 |
| | (0.0445) | (0.0142) | (0.0141) | (0.0139) | (0.00905) | (0.0110) | (0.0123) |
| Constant | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** | 0.0312*** |
| | (0.000454) | (0.000126) | (0.000127) | (0.000104) | (0.0000673) | (0.0000839) | (6.14e-20) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 1122624 | 1122624 | 1122624 | 1122624 | 1122208 | 978592 | 1115904 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

### F.4 Gender-Specific Results

Figure 14 shows the relationship between ethnic-gender overrepresentation and youth entry into occupations. Panel B shows, however, that the relationship is stronger for men, likely due to the weaker attachment of women to the labor force and other cultural factors.

Table A6: Occupation Choice: Youth Rate on Ethnic Ratio

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| Ethnic Ratio | 0.339*** | 0.166*** | 0.166*** | 0.277*** | 0.171*** | 0.180*** | 0.0435*** |
| | (0.0251) | (0.0123) | (0.0123) | (0.0119) | (0.00872) | (0.0106) | (0.0115) |
| Local Ratio | 1.041*** | 0.307*** | 0.280*** | 0.216*** | 0.0969*** | 0 | 0.223*** |
| | (0.0295) | (0.0184) | (0.0187) | (0.0156) | (0.0209) | (.) | (0.0221) |
| Constant | -0.0119*** | 0.0165*** | 0.0173*** | 0.0158*** | 0.0229*** | 0.0256*** | 0.0229*** |
| | (0.000228) | (0.000463) | (0.000481) | (0.000427) | (0.000622) | (0.000332) | (0.000612) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 1203328 | 1203328 | 1203328 | 1203328 | 1202976 | 1050624 | 1196736 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A7: Occupation Choice: Youth ORep on Ethnic ORep

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| Ethnic ORep | 0.148*** | 0.152*** | 0.152*** | 0.133*** | 0.160*** | 0.180*** | 0.0505*** |
| | (0.00751) | (0.00752) | (0.00753) | (0.00790) | (0.00808) | (0.0106) | (0.00911) |
| Constant | 1.14e-19 | 1.10e-19 | 1.10e-19 | 1.10e-19 | 1.10e-19 | 1.29e-19 | 1.12e-19*** |
| | (0.0000657) | (0.0000656) | (0.0000656) | (0.0000587) | (0.0000634) | (0.0000813) | (1.23e-20) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 1203328 | 1203328 | 1203328 | 1203328 | 1202976 | 1050624 | 1196736 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

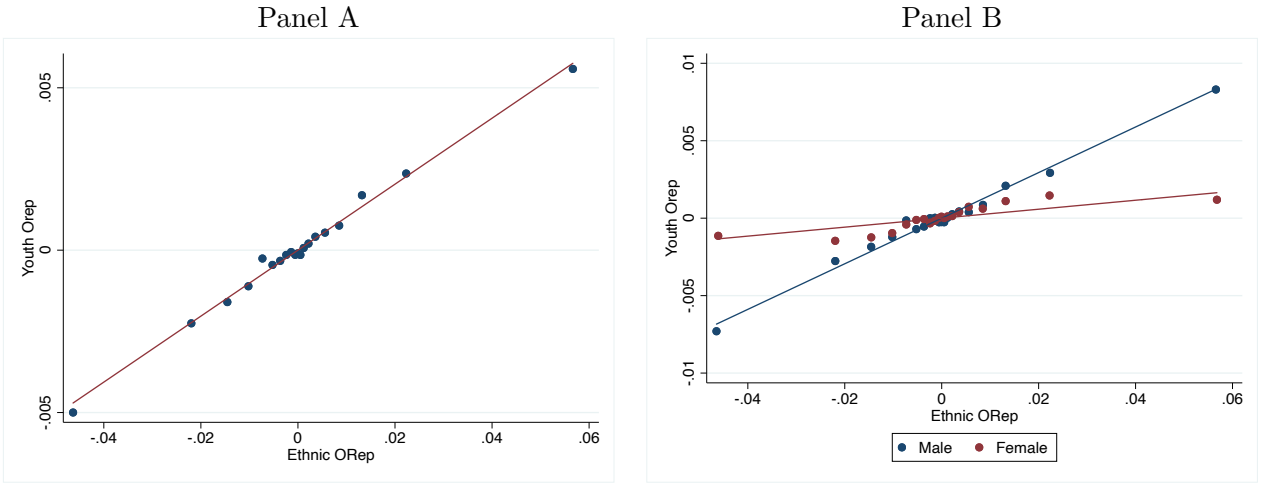
Table A8: Impact of a Shock on Occupation Choice (2005/06 and 2010-20)

| | (1) Youth Rate | (2) Youth Rate | (3) Youth Rate | (4) Youth Rate | (5) Youth Rate | (6) Youth Rate | (7) Youth Rate |
|---|---|---|---|---|---|---|---|
| Ethnic ORep (2005) | 0.614*** | 0.143*** | 0.153*** | 0.292*** | 0.146*** | 0.170*** | 0 |
| | (0.0755) | (0.0339) | (0.0338) | (0.0338) | (0.0298) | (0.0315) | (.) |
| POST=1 | 0.000460 | 0.000617** | 0 | 0.000556** | 0.000560** | 0 | 0.000558** |
| | (0.000339) | (0.000275) | (.) | (0.000276) | (0.000277) | (.) | (0.000275) |
| POST=1 X Ethnic ORep (2005) | 0.0204 | 0.0183 | 0.00391 | 0.0480 | 0.0396 | 0.000839 | 0.0433 |
| | (0.0403) | (0.0325) | (0.0328) | (0.0326) | (0.0324) | (0.0337) | (0.0323) |
| High Layoffs=1 | -0.0179*** | 0.00190*** | -0.00000514 | 0.00160*** | 0 | 0 | 0 |
| | (0.000819) | (0.000443) | (0.000425) | (0.000427) | (.) | (.) | (.) |
| High Layoffs=1 X Ethnic ORep (2005) | -0.0360 | 0.249*** | 0.225*** | 0.175*** | 0.270*** | 0.229*** | 0 |
| | (0.115) | (0.0666) | (0.0661) | (0.0643) | (0.0570) | (0.0608) | (.) |
| POST=1 X High Layoffs=1 | -0.00293*** | -0.00246*** | -0.000154 | -0.00221*** | -0.00222*** | 0 | -0.00222*** |
| | (0.000498) | (0.000437) | (0.000413) | (0.000437) | (0.000439) | (.) | (0.000435) |
| POST=1 X High Layoffs=1 X Ethnic ORep (2005) | -0.141* | -0.219*** | -0.185*** | -0.225*** | -0.231*** | -0.178*** | -0.229*** |
| | (0.0741) | (0.0657) | (0.0657) | (0.0647) | (0.0620) | (0.0644) | (0.0618) |
| Constant | 0.0372*** | 0.0321*** | 0.0326*** | 0.0321*** | 0.0325*** | 0.0322*** | 0.0325*** |
| | (0.000620) | (0.000271) | (0.000152) | (0.000261) | (0.000199) | (0.0000937) | (0.000185) |
| Occupation FE | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Year FE | No | No | Yes | No | No | Yes | No |
| Occupation X Year FE | No | No | Yes | No | No | No | No |
| Ethnicity FE | No | No | No | Yes | No | No | Yes |
| Occupation X Ethnicity FE | No | No | No | Yes | No | No | No |
| PUMA FE | No | No | No | No | Yes | Yes | Yes |
| Occupation X PUMA FE | No | No | No | No | Yes | No | No |
| Occupation X PUMA X Year FE | No | No | No | No | No | Yes | No |
| Occupation X PUMA X Ethniciy FE | No | No | No | No | No | No | Yes |
| Observations | 829690 | 829690 | 829690 | 829690 | 829580 | 717246 | 829220 |

Clustered standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 14: Relationship Between Ethnic Overrespresentation, Youth Entry, and Gender




***Notes***: Figures include occupation fixed effects and and *Ethnic ORep*. Regressions are weighted by the number of young people in the gender/ethnic/occupation/PUMA/year bin.