

INFORMATION RESONANCE

Ulrike Malmendier, Ryan Oprea and Laura Veldkamp*

July 27, 2026

Abstract

We study information resonance – the extra weight people attach to information from people similar to them – and examine its implications. In a new lab experiment designed to isolate resonance, we show that people follow the advice of strangers who share characteristics with them at a much higher rate. We find that this is true not only for shared demographics or political identities, but also for a broad range of beliefs, personality traits and even superficial preferences, and we measure these effects in several types of choices. The design sharply limits rational explanations based on information relevance or advisor quality. To examine the consequences of resonance, we adapt 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 coethnic incumbents, and are less likely to enter those occupations after coethnic incumbents 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 of their lives, 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. One important driver is a form of congruence: research has consistently shown that individuals are more likely to internalize and act upon information when it comes from people they identify with – people who share characteristics with them.¹

In this paper, we examine the hypothesis that such congruence effects are driven, at least in part, by what we call *information resonance*. People, we hypothesize, put more weight on information from individuals who share their characteristics simply because they are psychologically

*Preliminary. We thank Nikki Azerang, Rishi Chandran, Pietro Rizzo and Xiaobo Yu for excellent research assistance. Keywords: information, race, gender, homophily, social networks.

¹For 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).

more receptive to such information, i.e., because at some basic level such information resonates more strongly with them. While congruence effects in information transmission can also be driven by trust (people may trust the veracity of information from similar people more), memory (people may remember information more from people like them) or rational beliefs (people may, for good reason, view information from people like them as more relevant), we consider the possibility that people are simply more receptive to information from people like them. For instance, people might have deep-seated habits of following such information, preferences for following such information, or distorted beliefs in favor of such information. We call this increased receptivity *information resonance*, and we examine its scope and aggregate economic implications.

We take a multi-pronged approach to studying resonance. First, we use a pre-registered experiment to isolate resonance in a setting in which we can distinguish it from competing drivers of congruence effects like information relevance, memory and trust. The experiment teaches us that many types of personal characteristics can trigger information resonance. Perhaps surprisingly, the strongest triggers are not the usual demographic suspects but instead intimate personal details. Then, we provide a simple framework to trace out the logical consequences of such 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 though we also highlight the difficulty of cleanly identifying resonance from non-experimental economic data.

Section 2 presents our lab experiment, which attempts to isolate resonance from other drivers of congruence effects and explores its range and scope. Subjects make incentivized choices between information sources they'd like to receive, entertainment they'd like to experience, charities they'd like to donate to and even small investments they'd like to make, based on the advice of a stranger. Importantly we provide participants one piece of information about that stranger. Sometimes this evidence is demographic information like race or ethnicity, but often it is seemingly more superficial characteristics like whether they are a morning or night person, prefer cats or dogs, or even drink Coke or Pepsi. We measure resonance by studying whether advice of a stranger is more likely to be followed when the participant and her advisor have shared characteristics. This design isolates resonance by removing or sharply limiting the influence of competing drivers of congruence. For instance memory effects are irrelevant here because participants observe information about their advisor at the moment of choice. Likewise the plausibility that congruence effects are driven by beliefs about advisor quality (i.e., the belief that advisors like them are more sophisticated) or relevance (i.e., the belief that advisors like them have more similar preferences) is sharply limited by our inclusion of congruent characteristics (e.g., Coke vs. Pepsi preferences) that are unlikely to carry much substantive information about sophistication or preferences. Randomizing the match between dozens of congruence characteristics and dozens of distinct choices makes such alternative

channels especially unlikely.

We find strong evidence of resonance on average in our data: subjects are overall 1.831 times more likely to follow recommendations by advisors who are like them than unlike them. While previous studies make clear that race and gender produce such effects, 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 give extra weight to a recommendation from another dog lover than to a recommendation from someone who shares their race or gender. We try to push resonance to its breaking point by examining how minute and irrelevant a characteristic can be, and still trigger resonance and find that it occurs even for seemingly superficial characteristics. Nonetheless it is not universal. Resonance is weakest over biographical characteristics and some personality traits. For instance, information from advisors whose parents are divorced, who differ in reasoning style or who are optimistic/pessimistic, does not appear to trigger these effects.

Section 3 analyzes some of the theoretical implications of resonance and discusses some important applications. Our experiment shows that resonance takes the form of additional weight placed on information from congruent sources. We construct a model to trace out the logical consequences of this behavior. Section 3 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.

Our first proposition shows that 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

²In 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).

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 with unusual or extreme social characteristics. Before social media, many such people 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.

This insight suggests a change in perspective, 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 choose a social network *because* it is 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 suggests that selection is a key ingredient of interest – the characteristics that create the community and drives the exchange of information. This can create positive effects for members of marginalized groups, but also negative effects by 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 4 uses ACS data (2005–2020) to show that young workers’ occupational choices are shaped by same-ethnicity 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 distinguish between genders within a given ethnic group. We 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. As a result, informational constraints between ethnic groups are unlikely to explain the observed patterns. We also show that occupational layoff shocks influence the job choices of the younger generation, only if their ethnicity was overrepresented. The results are robust to a broad range of control variables, sample segmenting and differencing.

Our empirical findings are consistent with the interpretation that job choices and layoffs among

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

socially close community members resonate more. Differently from our controlled laboratory environment, though, the empirical analysis of observational data makes it difficult to rule out alternative explanations. For example, correlated occupational choices among same ethnicity-gender-PUMA peers might reflect correlated payoffs. The ethnicity-gender-location based results do make such interpretations less plausible, though, similarly to interpretations based on differential access to information: If people learn about the desirability or risk of an occupation from ethnic peers, and lack access to similar information about other occupations, such information is unlikely to be affected by gender boundaries within couples. In light of our experimental evidence, information resonance emerges as a possible interpretation – with different policy implications. The need for peers, rather than the provision of information, might be essential in attracting underrepresented groups into different occupations.

Taken together, the paper paints a picture of information resonance as a significant determinant of beliefs and choices. We provide a framework that captures what it looks like formally in a model, laboratory evidence of what it looks like in a setting designed to isolate its effects, and field evidence that suggests what it might look like in aggregate data. Each setting points us toward a shift in focus away from information access and toward understanding what is driving the behavior. This change in focus could greatly enhance the effectiveness of a broad array of economic and social policies.

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 interactions 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).

Economists have been interested in both *why* homophily comes about, *what* its economic effects are, and *how* to measure it. For instance, (Currarini et al., 2009) develop a theoretical model of friendship formation whereby homophily arises because of biases in preferences and meetings: individuals have preferences for forming connections with potential friends whom they share similar characteristics with *and* because they meet potential friends whom they share similar characteristics with disproportionately more frequently. (Golub and Jackson, 2012; Jackson and López-Pintado, 2013; Elliott et al., 2014) have shown that the degree of homophily in networks influences the speed of social learning and of consensus-reaching, the adoption of products, and the spread of financial failures, diseases, ideas, and information. (Goldsmith-Pinkham and Imbens, 2013; Graham, 2017) have proposed random and fixed effects approaches to estimate homophily parameters in network data, while allowing for individuals to have an arbitrary number of links (i.e. degree heterogeneity).

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

⁴For 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).

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 by introducing some minimal mathematical structure that describes 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 $\epsilon_j \sim \text{iid } N(0, \sigma_\epsilon^2)$ is idiosyncratic noise. Conditional on the agent’s information set $\mathcal{I}_i$, Bayesian updating yields the posterior expectation:

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

where α_i is the weight the agent assigns to the prior, and β_{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.

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*.

⁵Banerjee 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.

Under resonant belief updating, the posterior becomes:

$$E[z_i | \mathcal{I}_i] = \frac{1}{\bar{\omega}_i} \left[\alpha_i \mu_z + \sum_{j=1}^N \omega_{ij} \beta_{ij} (z_j + \epsilon_j) \right], \quad (2)$$

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

Let θ_i and θ_j denote the agents' vectors of observable or perceived characteristics. We introduce the pairwise resonance weight ω_{ij} as a decreasing function of the Euclidean distance between those vectors:

$$\omega_{ij} = \omega(\|\theta_i - \theta_j\|), \quad \omega'(\cdot) < 0, \quad (3)$$

The resonance weight is highest when agents have identical characteristic vectors and declines smoothly as perceived distance increases. Section 3 specifies the functional form used in the full diffusion model.

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 θ . But what constitutes similarity? θ 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 θ 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.

1.3 Interpretation

By modeling resonance as we have, we constrain the possible psychological mechanisms driving it. In particular our model of resonance does not capture congruence effects driven by memory or trust. A model of memory would be one in which information is forgotten, perhaps with some probability. Omitting information with a probability is not the same as weighting it less. Omitting information changes the relative weights on priors and other sources of information in a different way than what is represented here. Trust, meaning entertaining the probability that the message sender is lying, involves modeling an incentive to lie. In most such models, the probability of a lie would depend on the message sent. Because our weights on signals are independent of the signal realization, it does not faithfully represent the probability of truth-telling, or trust. For similar reasons, the weights in our model also do not describe congruence effects driven by systematically rational beliefs about the quality or relevance of transmitted information.

On the other hand, our model is consistent with congruence effects occurring due to i) a preference for weighting information from others with close characteristics, ii) a distorted belief that such information is more accurate than it is, or iii) a habit of weighting congruent information more, which might be well-adapted to other circumstances where information from similar others might truly be more relevant. Indeed, these possible mechanisms are all consistent with resonance as we will use the term throughout the paper. In the next section, we report an experiment designed to isolate resonance, so defined, and measure it in a context in which alternative drivers of such effects (like memory and trust) are shut down by design.

2 Experimental Evidence of Resonance

In order to gather some direct evidence about resonance and learn something about the range, scope and nature of the phenomenon, we ran a pre-registered 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 due to trust or 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/false 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 characteristics 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 stages.

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).
- **Entertainment Demand (Fiction Books).** Choice between pairs of fiction books (e.g., The Monk 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 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.

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

⁷A detailed description of how we constructed the ETFs choice questions is given in Online Appendix A.

⁸For 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.

- **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 selected question was a charity, we donated \$10 to the charity the subject selected while if it was a book we sent the subject an 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 was 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 recommendations and also gave contrasting answers to 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 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

⁹Subjects 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.

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

the recommenders.

Implementation. We pre-registered our experiment at aspredicted.org (#230,193). 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/hour for their participation (in addition to rewards in Stage 2, discussed above). Experimental instructions and screenshots of the interface are provided in Online Appendix A.

2.2 Understanding the Design

The experiment is a ‘within-subject’ design, i.e. every participant is exposed to every level of the independent variable or to every treatment. 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 relevance or 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. 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 can be researched online to lessen the ex ante similarity between the options. This gives 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

¹¹We 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.

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

A Characteristic triggers information resonance if participants disproportionately pick a choice variable that is suggested by a recommender whose Characteristic is the same as their own. 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 *congruent choice* (else the non-congruent choice). To measure resonance, we examine whether subjects are more likely than random (i.e., more than 50% likely) to make congruent choices. For instance, we examine to what degree subjects who are Democrats disproportionately (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 congruent choices (i.e., to the degree rates of congruent choice exceed 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 β_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* $R_c = \beta_c - 0.5$. Positive 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 terms, this subject is 1.5 times more likely to follow a recommendation by someone who shares their characteristic than someone who has the opposite characteristic.¹³

The solid dots in Figure 1 plot estimates of R_c for each Characteristic Question, c , grouped

¹²We 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 congruent 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 β_c .

¹³In 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).

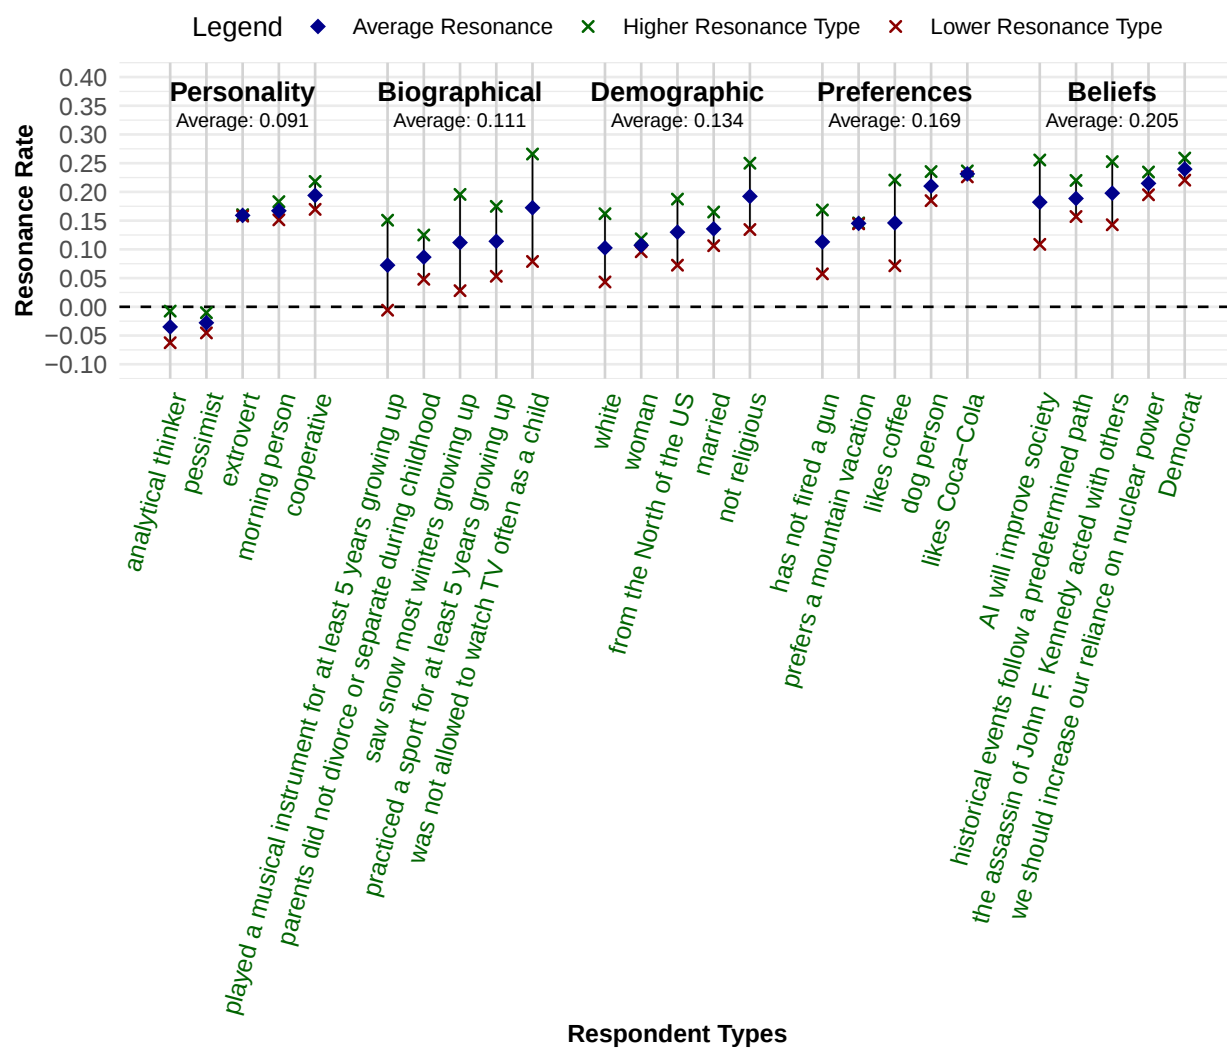


Figure 1: Resonance Estimates: Probability of following advice from others with similar characteristics.

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

according to our five categories of Characteristics Questions in our experiment. 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 Green (red) x 's plot the raw rates of congruent 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. 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 as more reliable. One exception is respondents who did not play a musical instrument growing up. Both respondents who played and did not play a musical instrument growing up are more likely to follow the advice of respondents who played a musical instrument growing up.

2.4 What the Results Teach Us.

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 congruence. 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 find 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

¹⁴When 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.

and biography (congruence in personality is weakest of all but still quite strong on average).

The breadth of contexts in which resonance arises is interesting in its own right. It is also helpful for understanding the likely psychological mechanisms underlying resonance. One possible reason for resonance effects like ours is that subjects 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 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 like 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 congruence seems severely limited may suggest that resonance is not (or is not only) a rational inference about the relevance of recommenders’ advice.

3. Resonance is not driven by systemic information contained in recommenders’ characteristics. Third, our findings cast further doubt on the rational explanation that the respondent type is informative. If this were true, resonance should mimic a covariance between personal characteristics and choices. In other words, if dog-lovers resonate more strongly when they recommend charities because they are more informed about charities, then dog-lovers should have more consistent preferences between the two charities than would recommendations with other characteristics.

In Online Appendix A we provide additional evidence about information contained in personal characteristics. This evidence takes two forms. First, most options are chosen just as frequently by one type of person as by the other type. Statistically, personal characteristics are not predictive of answers to Choice Questions. Second, the degree of resonance is insensitive to the degree to which characteristics *are actually predictive* of answers to Choice questions. If Stage 2 participants disproportionately followed advice from recommenders when those recommenders were more knowledgeable, we would expect a positive relationship between resonance rates and the rate of choice by that type. No such predictive relationship exists. That makes our results difficult to rationalize with a rational Bayesian inference model.

4. Resonance is widespread across types of behavior. Fourth, and relatedly, we find that resonance arises not only for many types of congruence, 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 resonance 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).

Differences in resonance. Table 1 shows the resonance rates for each category of personal characteristics (congruence domain) and each choice type. (Tables A4 and A5 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 preferences are the categories that give rise to the highest resonance rates. What triggers much less resonance is personality and childhood experiences.

The role of similar choices or expert advice. Resonance rates are similar in tasks that likely induce near-indifference (i.e., ETFs and True or False questions) and in tasks in which subjects can form sharper preferences by seeking out further information. Resonance appears to 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.

Additional information seeking. 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

Table 1: Resonance Rates by Congruence Domain and Choice Type

<i>Congruence Domain</i>	<i>Congruent Choice</i>	<i>Choice Type</i>	<i>Congruent Choice</i>
	(1)		(2)
Beliefs	0.215 (0.0126) [0.0000]	Books (Fiction)	0.155 (0.0125) [0.0000]
Biographical	0.111 (0.0124) [0.0000]	Books (Non-Fiction)	0.142 (0.0125) [0.0000]
Demographic	0.137 (0.0124) [0.0000]	Charities	0.157 (0.0126) [0.0000]
Personality	0.091 (0.0124) [0.0000]	ETFs	0.160 (0.0125) [0.0000]
Preferences	0.182 (0.0126) [0.0000]	True or False	0.120 (0.0126) [0.0000]
R^2	0.094	R^2	0.087
Adj. R^2	0.094	Adj. R^2	0.086
Sample Size	7275	Sample Size	7275

Note: OLS estimates of congruence choice (centered at 0) on congruence domains and choice types. Individual-level clustered standard errors (HC1) in round brackets. P-values in square brackets.

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 congruence 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.

2.5 Summary

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 in settings in which auxiliary information is available to inform decision-making.

These results paint a nuanced picture of what the model’s resonance weights ω_{ij} look like. We learn that weights on others’ information is high when agents i and j share demographic characteristics. The implied weight on others’ information is highest when agents i and j share beliefs, followed by shared preferences and then shared demographic characteristics. The weight is smaller for shared biographical characteristics and lowest for shared personality traits. But overall, they strongly support the hypothesis that resonance exists and affects decisions. The set of factors that trigger resonance is larger than we might have inferred from previous evidence, extending far beyond demographic categories like gender and race. Human beings find connections in all sorts of similarities and these connections shape the way we make decisions.

3 The Logical Consequences of Resonance

We use theory to map our individual-level choice findings to macro predictions. Motivated by our findings that resonance takes the form of greater weight on information from those with similar characteristics, we show how to trace out its dynamic aggregate effects. We merge the representa-

tion of resonant updating from Section 1 with information diffusion. The goal is not to innovate technically. Showing how we can model resonant information diffusion like geographic diffusion allows us to draw on what is already known about geographic diffusion to inform our thinking. We then employ standard diffusion logic to trace out predictions about how resonance translates to aggregate effects.

While we could derive many possible results from such a model, we choose a few that highlight how resonant learning differs from a more standard information friction, which limits access to information. We start by exploring the power of role models. Next, we combine 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, homophily in career choice and predictions that we later test with labor market data.

3.1 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 $i \in \{1, \dots, N\}$ chooses one of two actions $a_{it} \in \{0, 1\}$. 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 that agents can learn about and a transitory component ϵ_{it} that is i.i.d. over time and thus unlearnable. Agents have prior beliefs over the action 1 payoff $z \sim N(\mu_z, \sigma_z^2)$, for a finite σ_z^2 .¹⁵

Agents choose an action a_{it} to maximize their expected payoff:

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

Without the time- and individual-specific payoff shock $\epsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2)$, for a finite σ_ϵ^2 , one observation of the payoff from action 1 by any agent would reveal exactly the payoff z for all agents. The random payoff ϵ_{it} is what makes learning a gradual process.

Information sets. If any agent j takes the action $a_{jt} = 1$ at time t , then every agent i sees their action and payoff outcome at the end of period t . There is no information asymmetry here. All

¹⁵A richer model could have correlated z 's across agents. This would introduce variations in statistical relevance, as well as resonance. To isolate the resonance effect, we consider the simple case of identical payoffs z .

agents observe the same set of informative signals about z . How agents use those signals is dictated by resonance.

Belief updating 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 ω . This term allows us to account for agents giving less weight to information that comes from sources who are dissimilar.

A higher ω_{ij} means that i assigns more weight to j ’s experience when updating their beliefs, due to stronger perceived social proximity. $\bar{\omega}_{it}$ is a constant that ensures beliefs are on the right scale. For agent i , belief updating takes the form

$$E^r[z_i|\mathcal{I}_{it}] = \frac{1}{\bar{\omega}_{it}} \left[\alpha_t \mu_z + \sum_{j=1}^N \sum_{t'=1}^t \omega_{ij} \beta_t a_{jt'} (z + \epsilon_{jt'}) \right], \quad (5)$$

where α_t and β_t are the Bayesian weights on priors and signals. This is the dynamic analogue of equation (2): under the common-payoff assumption in this section, the static signal $z_j + \epsilon_j$ is replaced by the person-date payoff $z + \epsilon_{jt'}$ observed from agent j in period $t' \leq t$, and the term $a_{jt'}(z + \epsilon_{jt'})$ ensures that only informative actions with $a_{jt'} = 1$ enter the resonant signal sum.

First define the size of the data set at date t to be $m_t := \sum_{j=1}^N \sum_{t'=1}^t a_{jt'}$. This is the number of person-date observations everyone has of payoffs to action 1 by date t . As data accumulates over time, m_t grows. The Bayesian weights on new payoff signals are

$$\beta_t = \frac{\sigma_\epsilon^{-2}}{\sigma_z^{-2} + m_t \sigma_\epsilon^{-2}}. \quad (6)$$

The Bayesian weight β weights each signal by the precision of the signal σ_ϵ^{-2} , divided by the sum of precisions of all the signals plus the prior precision. The Bayesian weight on the prior belief about the payoff z is

$$\alpha_t = 1 - m_t \beta_t. \quad (7)$$

The more abundant or precise the new observations are, the less any agent relies on their prior. For a linear model with normally-distributed signal noise, the Bayesian weight on signals is simply the standard OLS estimator. The β is the covariance of the signals and the true z , divided by the signal variance, where the true variances and covariances are known.

The scaling factor $1/\bar{\omega}_{it}$ ensures the weights on priors and signals add up to one:

$$\bar{\omega}_{it} := \alpha_t + \sum_{j=1}^N \sum_{t'=1}^t \omega_{ij} \beta_{t'} a_{jt'} \quad (8)$$

The ω 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.

Belief weights and characteristics. For an observer i , different observations have different resonance. When i observes j , the weight i places on that observation, ω_{ij} , depends on the distance between the characteristics θ_i of the observer and the characteristics θ_j of the person who took the action.¹⁶

$$\omega_{ij} = 2 - 2\Phi(\chi \|\theta_i - \theta_j\|) \quad (9)$$

where $\|\theta_i - \theta_j\|$ is the Euclidean distance between agents' characteristics θ_i and θ_j .¹⁷ The multiplicative term χ allows us to consider different scenarios of agents weighting own versus other experiences more or less, as we discuss in more detail below. The reason for the 2's is that Normal cumulative distribution function Φ equals 1/2 at 0, and hence the expression $2 - 2\Phi(\cdot)$ takes on a value of 1 at zero. This means that among agents who resonate perfectly, information is treated in a Bayesian fashion. For agents who resonate less, the weight on their information declines towards zero. These terms define the resonance that a signal coming from j has for observer i .

We multiply the resonance weight with the (true) accuracy of the signal, σ_ϵ^{-2} , reflected in β_t , to reflect the idea that agents will not weight signals that have no information content. That would be a model of the propagation of misinformation, which is also realistic, but not the subject of this study. Thus, the overall weight is the product of the resonance weight ω_{ij} and the Bayesian weight β_t .

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

¹⁶What ω and θ 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 an information access matrix Θ_{it} , which could be related to the entries of θ_j . See proposition 2 for such a setting.

¹⁷One 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.

unknown payoffs z , using Bayes' rule as spelled out in equations (5), with weights ω_{ij} as defined in (9).

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.

Our analysis restricts attention to a set of parameters that generate interesting dynamics. In what follows, assume that the unknown action ($a = 1$) has a prior expected payoff that is lower than the payoff of the known action payoff of 0 ($\mu_z < 0$) and a true value above the known payoff ($z > 0$). These assumptions focus us on the subset of environments where learning leads to change. The $\mu_z < 0$ assumption corresponds to the assumption that the action $a = 1$ has unknown payoffs. The reason its payoffs are unknown is that no one in the community chooses to take the action. They don't choose it because its expected payoff is lower than the payoff of the action they do take. If this assumption were violated, we would simply switch which action was the known payoff action. There could be many cases where $z < 0$. In such cases, learning teaches agents that they are better off continuing to choose action $a = 0$. There is no diffusion of the change in action because it is sub-optimal. We consider actions that are payoff improving and thus have the potential to diffuse.

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 χ 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 the characteristic vector θ_i consists only of an agent's location or metropolitan area and $0 < \chi < 1$, the model becomes one of local learning from neighbors, as in (Fogli and Veldkamp, 2010; Burnside et al., 2016).¹⁸

¹⁸While 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

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 χ very large, learning is dynastic.

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 ω_{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 θ_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.

3.2 Role Models

We begin with a result that connects this model both to a voluminous literature on the importance of role models and with the experiment run in the previous section. Examples of the importance of role models 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 in other contexts is reported in Appendix B.

In our model, we use the term 'role model' to denote someone who is close in characteristic space and takes an informative action. Because the role model's characteristic similarity, the follower puts more weight on their experience, making them more likely to switch actions. While not surprising, this result illustrates the basic mechanics of the model.

information sources.

Proposition 1. *Consider agent i with characteristics θ_i and a role model with characteristics θ_r such that $\|\theta_i - \theta_r\| = e$, where e is small. Consider a non-role model j such that $\|\theta_i - \theta_j\| = \gamma > e$. All agents choose $a_{jt} = 0 \forall j, t$, except one. If the role model takes the informative action $a_{rt} = 1$, there is a greater chance of influencing i 's choice (a higher $\Pr[a_{i,t+1} = 1]$) than if the non-role model is the single agent choosing $a_{jt} = 1$.*

Proof: Using the objective (4) and the fact that the payoff of action 0 is 0, then $\Pr[a_{i,t+1} = 1] = \Pr[E^r[z + \epsilon_{i,t+1} | \mathcal{I}_t] > 0]$. Since the payoff shock ϵ is mean zero and independent of all other random variables, this is simply $\Pr[E^r[z | \mathcal{I}_t] > 0]$.

Consider the case where the non-role model j takes the informative action. In this case, all $a_{j't'} = 0$, except for one action by j at date t . Using the formula for resonant updating (5), we can rewrite $\Pr[E^r[z | \mathcal{I}_t] > 0] = \Pr[1/\bar{\omega}_{it}(\alpha_t \mu_z + \omega_{ij} \beta_t (z + \epsilon_{jt})) > 0]$. Rearranging, this is the probability that j 's idiosyncratic payoff shock is sufficiently high: $\Pr[\epsilon_{jt} > 1/(\omega_{ij} \beta_t)(-\alpha_t \mu_z) - z]$.

Recall that we assumed $\mu_z < 0$. The Bayesian weights α and β are always > 0 because variances are never negative and are assumed finite. Thus, the right hand side of the inequality is decreasing in ω_{ij} . So, a higher resonance weight corresponds to a larger probability that the shock ϵ_{jt} exceeds the lower threshold. Greater resonance creates a higher probability of action switching, $\Pr[a_{i,t+1} = 1]$.

In the case of the role model choosing $a_{rt} = 1$, the probability $\Pr[a_{i,t+1} = 1]$ is the same, except with resonance weight ω_{ir} , in place of ω_{ij} . From (9), we see that the resonance weight ω_{ij} is decreasing in the distance between the characteristics of i and j , $\|\theta_i - \theta_j\|$. Since that distance is assumed to be smaller for the role model, the role model resonates more: $\omega_{ir} > \omega_{ij}$. By the logic above, the more resonant role model action produces a higher $\Pr[a_{i,t+1} = 1]$. $\square$

This finding connects the model to the experiment results where information derived from the experiences of others was more influential when the provider of the information was closer in characteristic space to the decision-maker.

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

Next, we examine how online social network platforms interact with resonance. Online social networks have affected a radical change in information access in the last two decades. This result illustrates how resonance interacts with a traditional information access friction.

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 of others with shared characteristics, without geographic constraints. At the same time, the influence of experts declined. Resonance offers an explanation for this decline.

To capture the essence of the introduction of a social network, we will think of it as an increase in the size of one’s community. Let m_t be the number of people taking the informative action at date t . With a social network, m_t , the number of people one can observe, grows.¹⁹

Consider an expert e , who has more precise signals than others. This person’s outcomes receive a higher Bayesian weight β_e , because they contain more precise information. This person is not particularly resonant with others.

Proposition 2. *Define $w_{et} := (1/\bar{\omega}_{it})\omega_{ie}\beta_e$ to be the weight that agent i places on the observation of the expert’s signal. If the introduction of social media increases m_t , the size of the set of agents that i can observe, then the weight on expert signals declines: $\partial w_{et}/\partial m_t < 0$.*

Proof: From (9), the number of signals affects the weight on any one outcome only through the scaling factor $1/\bar{\omega}_{it}$. From (8), note that $\partial \bar{\omega}_{it}/\partial m_t > 0$. Thus, the weight w_{et} on the expert outcome declines as m_t rises. $\square$

This simple result, in fact, works through two channels: the size of the data set and the increased probability of connection with a resonant peer. To see each effect separately, it is helpful to consider two extreme values of χ , which determines how much a change in characteristic distance affects resonance weights. When $\chi = 0$, resonance weights are always 1. Agents are rational Bayesians. The social media result still holds in that case because with more non-expert observations, the expert signal is still rationally weighted less.

However, when χ is very large, resonance weights ω_{ij} on all signals are near zero, unless i and j are nearly identical in their characteristics. We will call such closely located agents “role models.” When i does not have access to any role models, she learns from the people available to her. While the weights ω_{ij} are all very small, they are scaled by a large $1/\bar{\omega}_i$ value. As a result, the information i has access to, like the expert signal, have a sizeable effect on her choice. However, when social media expands the observation set, if i finds a role model, that person has a strong, almost exclusive influence on i ’s choice, reducing the influence of the expert non-community member to near zero.

This second force is particularly strong for people who are extreme in the characteristic space. Extreme agents are more likely to not have access to role models initially. Limiting the scope of their social connections makes them learn from people different from them and moderates their actions.

¹⁹One can create a richer information structure that covaries with resonance by defining an observation matrix Ψ , where the (i, j) th entry of Ψ is 1 if i observes j and j takes action $a_j = 1$. The observation matrix can take many forms. Ψ_t might have a component that is time-varying. For example, if i sees only people in her neighborhood or only facebook friends, then Ψ_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 an $n \times 1$ vector containing every agent’s action at time t , and $\mathbf{1}'_n$ is a $1 \times n$ vector of ones. Celebrities might be visible with a higher probability than others. This would be captured by a Ψ_t matrix with a random component.

3.4 Resonance as a Foundation for Labor Market Discrimination

Resonance could also 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 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 C extends our model of resonance to show how it naturally lends itself to this type of labor market effect.

3.5 Career Choice Homophily (H1)

The next two results are simple corollaries of Theorem 1. Their role is to link the theory to predictions that we will test in labor market data.

When choosing a career, there are many occupations a worker could choose, all of which have uncertain payoffs. Consider two possible occupations a young worker i might choose: $a_i = 1$ or $a_i = 2$. The worker has the same prior beliefs over payoffs of the two occupations: $z_1, z_2 \sim i.i.d.N(\mu_z, \sigma_z^2)$. Worker i observes an existing worker in each occupation.

We assume that, unbeknownst to the young worker, $z_1 = z_2 > \mu_z$. We consider the opposite case in the next result.

The young worker observes the payoffs of the two existing workers: $z_1 + \epsilon_1$ and $z_2 + \epsilon_2$. The existing workers in occupations 1 and 2 have characteristics given by vectors θ_1 and θ_2 .

Proposition 3. Homophily in occupation choice. *Young worker i is more likely to select the occupation of the existing worker with higher resonance: $Pr[a_i = 1] > Pr[a_{i'} = 1]$ if $\|\theta_i, \theta_l\| < \|\theta_i, \theta_{l'}\|$, for $l, l' \in \{1, 2\}$.*

Proof: From the proof of proposition 1, we know that $Pr[E^r[z|z + \epsilon_{jt}] > 0]$ is increasing in resonance weight ω_{ij} when i observes j 's informative payoff $z + \epsilon_{jt}$. When we add a second action, and follow the same steps, we get $Pr[E^r[z_1|z_1 + \epsilon_{jt}] - E^r[z_2|z_2 + \epsilon_{j't}]]$ is increasing in the difference of resonance weights $\omega_{ij} - \omega_{ij'}$. Since resonance weights are decreasing in the characteristic distance, this means the expected payoff difference is decreasing in the difference of $\|\theta_i - \theta_1\| - \|\theta_i - \theta_2\|$. Since the worker chooses action 1 when $E^r[z_1|z_1 + \epsilon_{jt}] > E^r[z_2|z_2 + \epsilon_{j't}]$ and action 2 otherwise, this implies the result. $\square$

Resonance provides a clear rationale for why there is ethnic bunching across occupations. People observe the payoffs of those like them. If existing workers have positive signals, placing higher weight on these positive signals leads young workers to select similar careers. One interpretation of the

positive signal assumption is that the existing worker has selected this occupation because they learned that the realized payoff z was better than their belief about the payoff of the other job(s) μ_z . Of course, this would create a sample selection problem that would complicate the model without meaningfully changing the conclusion. Another way to arrive at the conclusion that more information weight makes a choice more likely is to model resonance as reducing uncertainty.

An alternative hypothesis that explains homophily in occupations is that people simply prefer working with others like themselves. 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. We derive that auxiliary prediction next.

3.6 Prediction: Resonance Transmits Negative Shocks More (H2)

Next we consider the case where the observed workers have negative payoff shocks.

Proposition 4. *Negative shocks affect resonant peers more.* Suppose there is a negative payoff shock to both occupations such that $z_1 = z_2 < \mu_z$. If the existing worker in occupation l is more resonant with young worker i ($\omega_{il} > \omega_{il'}$), then i is less likely to choose occupation l : $Pr[a_l = 1] < Pr[a_{l'} = 1]$, for $l, l' \in \{1, 2\}$.

Proof: For occupation $k \in \{l, l'\}$, define its posterior expected payoff as

$$M_k(\omega_{ik}) = \frac{\alpha\mu_z + \omega_{ik}\beta(z_k + \epsilon_k)}{\alpha + \omega_{ik}\beta},$$

where $\bar{\omega}_k = \alpha + \omega_{ik}\beta$. Taking expectations over the payoff shock and differentiating gives

$$\frac{\partial \mathbb{E}_\epsilon[M_k(\omega_{ik})]}{\partial \omega_{ik}} = \frac{\alpha\beta(z_k - \mu_z)}{(\alpha + \omega_{ik}\beta)^2} < 0,$$

because $z_k < \mu_z$. Thus, the normalization factor does not overturn the negative effect of resonance on the expected posterior following a negative shock. Because $z_l = z_{l'}$ and $\omega_{il} > \omega_{il'}$, the difference $M_l - M_{l'}$ has a negative mean. With independent normally distributed payoff shocks,

$$\Pr(M_l < M_{l'}) > \frac{1}{2}.$$

This implies the result. $\square$

The testable prediction is that ethnic communities will respond more to negative shocks that affect industries in which their ethnic kin are heavily represented. We will see this take the form of young 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.

4 Evidence from Occupational Choice

We close by providing some evidence that aggregate patterns predicted by our model appear in labor market data. For this we focus on documenting the patterns predicted in 3.5 and 3.6, thereby linking the individual tendencies documented in our experiment, via the model, to real-world data. The goal of this section is not to empirically identify resonance (which was the task of the experiment), rather it is to examine whether aggregate patterns predicted by our model appear in the data and show how to uncover them. Though in principle we could study resonance along any characteristic, we focus on the linkage between ethnic identity and occupational choice for reasons of policy relevance and data availability.

Using the American Community Survey (ACS), we provide suggestive evidence for Prediction H1: 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. To do so, we observe young labor force entrants following their community elders into the same occupations. More specifically, we find that the choice of young people 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.

To further probe the robustness of our findings and alleviate information access concerns, we also show that the effect of resonance is stronger when matching on more demographic characteristics. In particular, we use an additional measure of social proximity, gender, and show that young members of an ethnic community enter occupations in which elders of the same ethnicity *and* gender are overrepresented. If men and women of the same ethnic communities share information about their occupations and have overlapping social networks, it makes it harder to attribute these findings solely to information access.

That said, while the evidence for Prediction H1 is consistent with information resonance we cannot fully rule out every channel behind ethnic- or gender-based occupational sorting (e.g., social stigma, identity effects, etc.). For this reason, we believe that the strongest evidence for the effect of resonance on the aggregate labor market comes from our test for Prediction H2, where we directly examine how a plausibly exogenous negative shock to members of one’s community impacts the behavior of young labor market entrants of the same group. Here, we show that young people from ethnic groups that were overrepresented in occupations most negatively affected by the 2008 financial crisis are later discouraged from entering such occupations.

This analysis makes progress in isolating resonance from other forces. But it also has clear limitations. As already mentioned for what concerns our tests of Prediction H1, 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. To end this section, we review in detail what potential channels our econometric framework can account for and what identification challenges we cannot overcome.

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.

Evidence for Prediction H1: Adult Overrepresentation We define Adult Overrepresentation (Adult 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 (Adult Rate) and the overall share of workers in that occupation (Local Rate).²⁰

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

$$Adult\ ORep_{i,j,k,t} = Adult\ Rate_{i,j,k,t} - Local\ Rate_{j,k,t}.$$

²⁰For a detailed description of how these variables were constructed, see the Appendix-Section D.


Similarly, we measure the Adult Rate 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 2 shows that young labor force entrants are following in the path of their same ethnic elders. In Panel A, we can see that as adult 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 2: Relationship Between Ethnic Overrepresentation 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 Rate* as an additional control variable.

Taking this analysis further, we run a linear regression of the youth 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 adult overrepresentation, the percent amount by which the ethnic group is overrepresented in the occupation in the PUMA that year. Our baseline specification is

$$Youth\ Rate_{i,j,k,t} = \gamma_{j,k,t} + \beta(Adult\ ORep_{i,j,k,t}) + \varepsilon_{i,j,k,t},$$

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

will increase by β percentage points.

Table D5 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 *Adult ORep* as the influence of ethnicity on occupation choice holding other local factors constant. If *Adult ORep* increases by 1.2% (the IQR of *Adult ORep*), *Youth Rate* will increase by 0.2 pp— 6% of the mean.

Evidence for Prediction H1: Information Frictions and The Role of Gender In addition to ethnicity, we consider gender as another measure of social proximity. To the extent that husbands and wives, fathers and daughters, share information about their jobs and engage in the same social and professional networks, finding that labor market entrants follow same gender elders into occupations is harder to attribute to differences in information access alone. In addition, finding that labor market entrants do not follow opposite gender elders in occupations is consistent with a high resonance weight on information obtained from individuals of the same gender.

In what follows we discuss the effect of *Adult ORep (Same Gender)*, the degree to which elders of the *same ethnic group* and *same gender* are overrepresented in an occupation, and of *Adult ORep (Wrong Gender)*, the degree to which elders of the *same ethnic group* but *other gender* are overrepresented in an occupation, on the *Youth Rate*.²¹ Table D9 and Table D10 in Appendix D show that young people follow their coethnic elders’ occupational paths within gender. Based on Column (6) of Table D9 and using the interquartile range of *Adult ORep (Same 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 0.26 pp. In contrast, Table D10 shows little evidence that young labor market entrants follow in the occupations of coethnic elders of the opposite gender. In all specifications where we control for occupation-specific effects (columns (2)-(7)), the coefficient on *Adult ORep (Wrong Gender)* is a precise zero²².

Taken together, these results are consistent with our model of resonance where individuals’ resonance weight ω_{ij} is high if their genders θ_i and θ_j are the same, and low otherwise.²³ However,

²¹To construct *Adult ORep (Wrong Gender)*, we first compute *Adult ORep (Same Gender)* as the share of elders in occupation j among those of ethnicity i and gender l , in PUMA k and in year t (i.e. *Adult Rate (Same Gender)*) minus the share of elders in occupation j among those of gender l , in PUMA k , and year t (i.e. *Local Rate (Same Gender)*). Then we permute *Adult ORep (Same Gender)* within occupation j , ethnicity i , PUMA k , and year t blocks.

²²Note that when the only identifying variation is gender l , *Adult ORep (Wrong Gender)* is perfectly collinear with *Adult ORep (Same Gender)* (in fact it is its additive inverse). Thus, the coefficient in column (8) of Table D10 is the additive inverse of the coefficient in column (8) of Table D9. However, when the identifying variation is more than just gender l , the relationship between *Youth Rate*, *Adult ORep (Same Gender)*, and *Adult ORep (Wrong Gender)* is not mechanical and economically meaningful.

²³Table D9 shows that if we put both *Adult ORep (Same Gender)* and *Adult ORep (Wrong Gender)* on the

we still cannot fully rule out potential channels other than resonance that could drive the gender results. For instance, incentives to work long, and particular, hours might push women out of some occupations and into others (Goldin, 2014) or men entering women dominated occupations might face social stigma from peers (Williams, 1992).

Evidence for Prediction H2: Economic Shock To do away with these concerns that arise from observational regressions, we now provide evidence for Prediction H2 exploiting a plausibly exogenous shock to ethnic representation in specific occupations. 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.²⁴ We include an indicator variable, *POST*, indicating that the observation is from the year 2010 or later (i.e., post-recession).

Table D14 shows a positive coefficient on the interaction of the high layoffs dummy and pre-recession adult 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* $\times$ *POST* $\times$ *Adult 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 0.2 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* $\times$ *Adult ORep* is insignificant)—meaning that young people are continuing to

RHS, the coefficient on *Adult ORep (Same Gender)* is of similar magnitude, while *Adult ORep (Wrong Gender)* decreases in magnitude and becomes statistically significant in some specifications. In Table D13 we use a simple FWL decomposition to show that the decrease in magnitude in the *Adult ORep (Wrong Gender)* is due to the fact that the *Adult ORep (Same Gender)* coefficient is relatively large and that *Adult ORep (Same Gender)* is strongly predictive of *Adult ORep (Wrong Gender)* (more details can be found in Appendix D). This empirical fact is not directly inconsistent with resonance, since in our model we do not make statements about how people use non-resonant information *holding the level of resonant information constant* to update their beliefs and vice versa.

²⁴For this analysis, we fix the adult 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.

follow their same-ethnic elders' career paths in the absence of a shock.

Identification: What We Can Rule Out The preceding empirical results serve as motivating evidence that people place greater weight on information learned from people they are socially close to. While, as we discuss below, we cannot rule out all alternative channels that could generate homophily in occupational choice, the fixed effects in these specifications help control for some.

In the ethnicity and gender regressions, our preferred specifications²⁵ include either occupation-by-PUMA-by-year or occupation-by-ethnicity-by-year effects. The former set of fixed effects compare ethnic groups within the same occupation, local labor market, and year. They therefore absorb all shocks and conditions that are common to an occupation in a PUMA-year, such as local changes in labor demand, occupation-specific hiring opportunities, local wages, or occupation size. Occupation-by-ethnicity-by-year fixed effects instead compare PUMAs within the same ethnic group, occupation, and year. They absorb national occupation-specific trends for each ethnic group, such as changes in an ethnic group's aggregate attachment to an occupation, national ethnic networks in that occupation, immigration-composition changes affecting that occupation, or ethnicity-specific barriers and opportunities that vary over time at the occupation level.

Thus, the results are not driven mechanically by either local occupation-year shocks common to all ethnic groups or by national ethnicity-specific occupational trends.

Identification: Challenges However, identification of information resonance poses many questions that this empirical exercise does not address. In addition to alternative channels driving gender-based occupational sorting mentioned above, 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 (Akerlof and Kranton, 2000), 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.²⁶

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 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.

²⁵Column (6) and Column (7) of Table D5, Table D9, Table D10, and Table D11.

²⁶The 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.

5 Conclusion

This paper recasts the idea that people listen to people like them, as a systematic information-weighting bias that we call information resonance. First, we measure information resonance using a tightly controlled laboratory experiment that rules out alternative mechanisms like memory, relevance and trust, and nonetheless uncovers a large premium on advice that resonates through shared beliefs, preferences or other characteristics. Next, we build a model of the information weighting found in the experiment and study its aggregate consequences. 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 can eclipse the influence of many distant experts and generates homogeneous leadership structures despite neutral promotion rules. Third, we examine these aggregate predictions using U.S. labor data and find patterns consistent with the same kind of bias at scale: young workers sort toward occupations dominated by coethnic incumbents and react more when incumbent workers like them suffer negative 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 information and advice 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 misinformation, conveyed in chat rooms populated by those with similar interests, can often be so effective.

Finally, information diffusion happens not only in geographic space, but in social characteristic space. A change in policy 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

- G. A. Akerlof and R. E. Kranton. Economics and identity. *The Quarterly Journal of Economics*, 115(3):715–753, 2000. ISSN 0033-5533.
- 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. Stroebe. 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.

- S. Currarini, M. O. Jackson, and P. Pin. An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4):1003–1045, 2009. ISSN 1468-0262. doi: 10.3982/ECTA7528.
- P. E. Downes, E. R. Crawford, S. E. Seibert, A. C. Stoverink, and E. M. Campbell. Referents or role models? the self-efficacy and job performance effects of perceiving higher performing peers. *Journal of Applied Psychology*, 106(3):422–438, 2021. doi: 10.1037/apl0000519.
- F. D’Acunto, A. Fuster, and M. Weber. Diverse policy committees can reach underrepresented groups. Technical report, National Bureau of Economic Research, 2021.
- M. Elliott, B. Golub, and M. O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–3153, Oct. 2014. ISSN 0002-8282. doi: 10.1257/aer.104.10.3115.
- 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.
- J. Gladstone and A. Cimpian. Which role models are effective for which students? a systematic review and four recommendations for maximizing the effectiveness of role models in STEM. *International Journal of STEM Education*, 8(1):59, 2021. doi: 10.1186/s40594-021-00315-x.
- C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4): 1091–1119, Apr. 2014. ISSN 0002-8282. doi: 10.1257/aer.104.4.1091.
- P. Goldsmith-Pinkham and G. W. Imbens. Social networks and the identification of peer effects. *Journal of Business and Economic Statistics*, 31(3):253–264, July 2013. ISSN 0735-0015. doi: 10.1080/07350015.2013.801251.
- B. Golub and M. O. Jackson. How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338, Aug. 2012. ISSN 0033-5533. doi: 10.1093/qje/qjs021.
- B. S. Graham. An econometric model of network formation with degree heterogeneity. *Econometrica*, 85(4):1033–1063, 2017. ISSN 1468-0262. doi: 10.3982/ECTA12679.
- 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.
- M. O. Jackson and D. López-Pintado. Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science*, 1(1):49–67, Apr. 2013. ISSN 2050-1242, 2050-1250. doi: 10.1017/nws.2012.7.
- 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.
- E. Lawner, D. M. Quinn, G. Camacho, B. T. Johnson, and B. Pan-Weisz. Ingroup role models and underrepresented students’ performance and interest in STEM: A meta-analysis of lab and field studies. *Social Psychology of Education*, 22(4):1169–1195, 2019. doi: 10.1007/s11218-019-09518-1.

- 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.
- C. Verniers. The double-edged sword of role models: A systematic review. *Review of Educational Research*, 2024. doi: 10.3102/0091732X241261310. Advance online publication.
- C. L. Williams. The glass escalator: Hidden advantages for men in the “female” professions. *Social Problems*, 39(3):253–267, Aug. 1992. ISSN 0037-7791. doi: 10.2307/3096961.

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 A1 and A2 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 A3 shows a sample Stage 2 choice question and Figure A4 shows a sample Stage 1 choice question. Finally, Figure A5 shows a sample characteristics question (these questions are not randomized and are the same in Stage 1 and Stage 2).

Table A1 lists the characteristics of experiment participants that we elicited (also called recommender characteristics or congruence 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 A2 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 National Longitudinal Study of Adolescent to Adult Health (Add Health)²⁷. 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 A1) we wrote that we would randomly select one of the choice questions participants answered to (for a randomly 10% of participants). If

²⁷Add 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 from Yahoo Finance²⁸ for pairs of ETFs covering the same sector—for instance, two ETFs tracking healthcare companies or gold—from January 1, 2020, through November 1, 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.
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 ensured 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 A1) 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 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’:

²⁸<https://finance.yahoo.com/>

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 A1: Categorized Recommender/Respondent Characteristics

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

Note: List of the 25 recommender/respondent characteristics pairs used in the experiment.

Table A2: Categorized Choice Variables

Choice Variables
<i>Books (Fiction)</i>
Suldrun’s Garden (The Riddle-Master of Hed).
The Futurological Congress (The Shockwave Rider).
The Game of Kings (The King’s General).
The Leavenworth Case (The Mystery of a Hansom Cab).
The Monk (Melmoth the Wanderer).
<i>Books (Non-Fiction)</i>
Deep Work (The Power of Habit).
Emotional Agility (Insight).
Learn Python the Hard Way (Python Crash Course).
The Joy of Less (The Power of Purpose).
The Simple Path to Wealth (Your Money or Your Life).
<i>Charities</i>
Better Angels (The Bridge Alliance).
ChildFund International (Plan International).
Disability Rights Fund (Disability Rights International).
Jump\$tart Coalition for Financial Literacy (National Endowment for Financial Education).
Water.org (WaterAid).
<i>ETFs</i>
Emerging Markets Fund A (B).
Gold Fund A (B).
Healthcare Fund A (B).
Japan Exposure Fund A (B).
Real Estate Fund A (B).
<i>True or False</i>
Among people who like to drink alcohol or smoke marijuana, share of registered voters is higher (lower) than 60%.
Among people who consider themselves vegetarians share who filed a tax return in the past 12 months is higher (lower) than 70%.
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%.
Among people who have ever bought a lottery ticket, share who have eaten breakfast in the past 7 days is higher (lower) than 50%.
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 the 25s 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 A1: Stage 2 instructions in the order displayed to the participants.

Survey Instructions: Making Choices

We will present you with three kinds of questions:

1. We will present you with a series of pairs of charities and books. We are interested in knowing which charities and 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 choices, 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 holds 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**.

Your answers to these questions may be used as recommendations or advice for participants in a future survey.

Figure A2: Stage 1 instructions.

Choice 2

Please choose between these two **non-fiction books** about **emotional intelligence and self-awareness**:

Participant WPVC82836 recommended

Emotional Agility: Get Unstuck, Embrace Change, and Thrive in Work and Life

Trivia 🎉: participant WPVC82836 **likes Pepsi**.

Participant ZUDZ50727 recommended

Insight: The Surprising Truth About How Others See Us, How We See Ourselves, and Why the Answers Matter

Trivia 🎉: participant ZUDZ50727 **likes Coca-Cola**.

Click here to display more information about 'Emotional Agility: Get Unstuck, Embrace Change, and Thrive in Work and Life'.

Click here to display more information about 'Insight: The Surprising Truth About How Others See Us, How We See Ourselves, and Why the Answers Matter'.

Figure A3: Example of a Stage 2 choice question.

Choice 25

Please choose between these two non-fiction books about emotional intelligence and self-awareness:

Insight: The Surprising Truth About How Others See Us, How We See Ourselves, and Why the Answers Matter

Insight by Tasha Eurich explores the transformative power of self-awareness in achieving personal and professional success. Eurich, an organizational psychologist, combines compelling research with relatable stories to reveal the surprising truth about how we perceive ourselves and how others perceive us. The book introduces two types of self-awareness: internal (understanding our values, emotions, and impact) and external (grasping how others view us). Eurich emphasizes the importance of balancing both types to make better decisions, improve relationships, and lead more effectively. Through practical exercises, Insight provides tools to enhance self-awareness, such as identifying blind spots, seeking honest feedback, and challenging ingrained beliefs. Eurich highlights the barriers that prevent people from achieving self-awareness, including cognitive biases and societal pressures, while offering strategies to overcome them. The book delves into the benefits of aligning actions with personal values, fostering authenticity and growth. By enhancing self-awareness, Eurich argues, individuals can unlock their potential and create meaningful connections. Insight is a thought-provoking guide for anyone seeking greater clarity about themselves and their interactions with others, making it an invaluable resource for personal development and leadership growth.

Emotional Agility: Get Unstuck, Embrace Change, and Thrive in Work and Life

In Emotional Agility, psychologist Susan David presents a revolutionary framework for navigating life's challenges with adaptability and resilience. Drawing on decades of research, David introduces emotional agility as the ability to recognize, understand, and manage emotions in ways that foster growth and align with personal values. The book provides a four-step process for cultivating this skill: showing up (accepting emotions without judgment), stepping out (distancing oneself from harmful patterns), walking your why (acting in alignment with values), and moving on (making small, intentional adjustments). David uses engaging anecdotes and practical exercises to illustrate how emotional agility can lead to stronger relationships, improved well-being, and professional success. She emphasizes the importance of embracing, rather than suppressing, difficult emotions to foster self-compassion and authentic living. The book also addresses the psychological traps that hinder progress, offering strategies to overcome them and build resilience. By learning to respond flexibly to life's uncertainties, readers can thrive in a constantly changing world. Emotional Agility is an essential guide for anyone looking to enhance their emotional intelligence and lead a more intentional, fulfilling life.

Figure A4: Example of a Stage 1 choice question.

Question 6

Please select the option that best describes you:

A person who **likes Coca-Cola**.

A person who **likes Pepsi**.

Does not apply.

Figure A5: Example of a Stage 1 or Stage 2 characteristics question.

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

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

Table A3 reports how many Stage 2 participants are best described by the various respondent characteristics. As we can see, the characteristics are mostly balanced within pairs. Among those characteristics pairs that are more skewed towards a specific type, we observe, for instance, disproportionately more participants who like coffee (247) than who do not (49), or who consider themselves optimists (242) as compared to pessimists (49).

However, as we show in Tables A4 and A5 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 break down 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 terms of characteristics in Figure A6.

Table A3: Stage 2 Respondent Type Counts

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

Note: Frequency of characteristics (congruence 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 to the congruence 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 A4: Resonance Rates by Recommender Characteristics Pair

Characteristic	Congruent Choice	
	Unweighted (1)	Weighted (2)
<i>Beliefs</i>		
AI will improve (worsen) society	0.221 (0.0276) [0.0000]	0.182 (0.0279) [0.0000]
Democrat (Republican)	0.236 (0.0287) [0.0000]	0.240 (0.0279) [0.0000]
Historical events are mostly random (follow a pre-determined path)	0.190 (0.0278) [0.0000]	0.189 (0.0279) [0.0000]
The assassin of JFK acted alone (with others)	0.214 (0.0283) [0.0000]	0.198 (0.0279) [0.0000]
We should increase (decrease) reliance on nuclear power	0.215 (0.0279) [0.0000]	0.215 (0.0279) [0.0000]
<i>Biographical</i>		
Parents divorced or separated (did not divorce or separate) during childhood	0.106 (0.0275) [0.0001]	0.086 (0.0279) [0.0019]
Played (did not play) a musical instrument for at least 5 years growing up	0.061 (0.0275) [0.0272]	0.072 (0.0279) [0.0093]
Practiced (did not practice) a sport for at least 5 years growing up	0.136 (0.0275) [0.0000]	0.114 (0.0279) [0.0000]
Saw (didn't see) most winters growing up	0.145 (0.0275) [0.0000]	0.112 (0.0279) [0.0001]
Was (was not) allowed to watch TV often as a child	0.109 (0.0276) [0.0001]	0.172 (0.0279) [0.0000]
<i>Demographic</i>		
From the North (South) of the US	0.131 (0.0282) [0.0000]	0.130 (0.0279) [0.0000]
Man (Woman)	0.107 (0.0274) [0.0001]	0.107 (0.0279) [0.0001]
Married (not married)	0.146 (0.0275) [0.0000]	0.136 (0.0279) [0.0000]
Religious (not religious)	0.172 (0.0277) [0.0000]	0.192 (0.0279) [0.0000]
White (not white)	0.130 (0.0275) [0.0000]	0.103 (0.0279) [0.0002]
<i>Personality</i>		
Analytical (creative) thinker	-0.037 (0.0276) [0.1755]	-0.035 (0.0279) [0.2094]
Cooperative (competitive)	0.200 (0.0277) [0.0000]	0.194 (0.0279) [0.0000]
Extrovert (introvert)	0.159 (0.0275) [0.0000]	0.159 (0.0279) [0.0000]
Morning person (night owl)	0.171 (0.0277) [0.0000]	0.167 (0.0279) [0.0000]
Optimist (pessimist)	-0.040 (0.0278) [0.1546]	-0.028 (0.0279) [0.3180]
<i>Preferences</i>		
Cat (dog) person	0.219 (0.0282) [0.0000]	0.210 (0.0279) [0.0000]
Has (has not) fired a gun	0.123 (0.0275) [0.0000]	0.113 (0.0279) [0.0001]
Likes (does not like) coffee	0.196 (0.0275) [0.0000]	0.146 (0.0279) [0.0000]
Likes Coca-Cola (Pepsi)	0.233 (0.0288) [0.0000]	0.231 (0.0279) [0.0000]
Prefers a beach (mountain) vacation	0.145 (0.0277) [0.0000]	0.145 (0.0279) [0.0000]
R^2	0.106	0.099
Adj. R^2	0.103	0.096
Sample Size	7275	7275

Note: OLS estimates of congruent choice (centered at 0) on congruence characteristics. Heteroskedasticity-robust standard errors clustered at the participant level are reported in parentheses. P-values are reported in square brackets. Column (2) adjusts observations so that each congruence type contributes equally. The weights are constructed as follows: (i) we select respondent's choices with the same recommender's congruence characteristics, (ii) we count how many respondents are of either congruence 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 A5: Resonance Rates by Choice Pair

Characteristic	Congruent Choice	
	Unweighted (1)	Weighted (2)
<i>Books (Fiction)</i>		
Melmoth the Wanderer (The Monk)	0.104 (0.0279) [0.0002]	0.112 (0.0284) [0.0001]
Suldrun's Garden (The Riddle-Master of Hed)	0.163 (0.0279) [0.0000]	0.171 (0.0281) [0.0000]
The Futurological Congress (The Shockwave Rider)	0.144 (0.0278) [0.0000]	0.128 (0.0281) [0.0000]
The Game of Kings (The King's General)	0.198 (0.0280) [0.0000]	0.187 (0.0280) [0.0000]
The Leavenworth Case (The Mystery of a Hansom Cab)	0.166 (0.0281) [0.0000]	0.142 (0.0280) [0.0000]
<i>Books (Non-Fiction)</i>		
Deep Work (Power of Habit)	0.114 (0.0281) [0.0001]	0.116 (0.0282) [0.0000]
Emotional Agility (Insight)	0.159 (0.0279) [0.0000]	0.156 (0.0279) [0.0000]
Essentialism (Joy of Less)	0.175 (0.0280) [0.0000]	0.194 (0.0280) [0.0000]
Learn Python the Hard Way (Python Crash Course)	0.131 (0.0279) [0.0000]	0.136 (0.0278) [0.0000]
The Simple Path to Wealth (Your Money or Your Life)	0.131 (0.0279) [0.0000]	0.121 (0.0273) [0.0000]
<i>Charities</i>		
Better Angels (Bridge Alliance)	0.129 (0.0280) [0.0000]	0.129 (0.0281) [0.0000]
ChildFund (Plan) International	0.140 (0.0281) [0.0000]	0.143 (0.0280) [0.0000]
Disability Rights Fund (International)	0.112 (0.0283) [0.0001]	0.116 (0.0284) [0.0000]
Jump\$tart Coalition (National Endowment) for Financial Literacy	0.178 (0.0281) [0.0000]	0.185 (0.0278) [0.0000]
Water.org (WaterAid)	0.224 (0.0279) [0.0000]	0.212 (0.0281) [0.0000]
<i>ETFs</i>		
Emerging Markets Fund A (B)	0.190 (0.0282) [0.0000]	0.211 (0.0282) [0.0000]
Gold Fund A (B)	0.137 (0.0278) [0.0000]	0.134 (0.0278) [0.0000]
Healthcare Fund A (B)	0.176 (0.0279) [0.0000]	0.144 (0.0287) [0.0000]
Japan Exposure Fund A (B)	0.147 (0.0278) [0.0000]	0.114 (0.0280) [0.0001]
Real Estate Fund A (B)	0.149 (0.0280) [0.0000]	0.122 (0.0276) [0.0000]
<i>True or False</i>		
Share who filed a tax return is higher (lower) than 70%	0.153 (0.0280) [0.0000]	0.134 (0.0284) [0.0000]
Share who did housework is higher (lower) than 80%	0.127 (0.0282) [0.0000]	0.142 (0.0283) [0.0000]
Share who ate breakfast is higher (lower) than 50%	0.096 (0.0282) [0.0007]	0.084 (0.0287) [0.0033]
Share who wake up before 5AM is higher (lower) than 20%	0.123 (0.0281) [0.0000]	0.108 (0.0281) [0.0001]
Share of registered voters is higher (lower) than 60%	0.099 (0.0282) [0.0004]	0.106 (0.0280) [0.0001]
R^2	0.090	0.085
Adj. R^2	0.087	0.082
Sample Size	7275	7275

Note: OLS estimates of congruent choice (centered at 0) on congruence characteristics. Heteroskedasticity-robust standard errors clustered at the participant level are reported in parentheses. P-values are reported in square brackets. Column (2) adjusts observations so that each congruence type contributes equally. The weights are constructed as follows: (i) we select respondent's choices with the same recommender's congruence characteristics, (ii) we count how many respondents are of either congruence 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 A6: OLS estimates of congruent choice (centered at 0) on congruence characteristics using heteroskedasticity-robust standard errors clustered at the participant level 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 congruence characteristics, (ii) we count how many respondents are of either congruence 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 Choice Options and of Recommender Characteristics

In this subsection we shed light on whether specific choice options or recommender characteristics carry additional informational content. That is, we first investigate whether participants are indifferent between choice options and, second, informational content of the recommenders predicts resonance rates.

We start by investigating whether, for each choice pair, participants are indifferent between the given choice options. In other words, we look at whether participants' choices are balanced. Figures A7 and A8 show balance rates at the choice pair level using Stage 1 and Stage 2 choices respectively. For each choice pair, we define a balance rate to be the proportion of the most frequent choice option chosen by participants normalized around 0.5. On average, Stage 1 participants were indifferent for about half of the choice pairs, while they disproportionately preferred a specific choice option for the other half of the choice pairs. We see a similar picture in Stage 2, but the rates are overall smaller and the number of choice pairs for which participants are indifferent is larger. The difference between Stage 1 and Stage 2 balance rates could be due to the fact that in Stage 1 choices we always show the additional information, while in Stage 2 we allow participants to click to display additional information about the choice options. The fact that choices were not recommended in Stage 1 could also play a role. Either way, Figures A7 and A8 show that while for each choice pair choice options were chosen to minimize orthogonal characteristics of choices that might interfere with the measurement of resonance (that is they were deliberately chosen to be similar), choice options are still significantly different to participants for a large number of choice pairs.

Now we investigate whether higher choice imbalance rates are associated with lower resonance rates. In doing so, we also address whether resonance rates are partly driven by the informational content of the recommender characteristics. Figure A11 plots Stage 2 resonance rates against Stage 1 choice balance rates at the choice pair, Stage 1 recommender characteristic pair, and Stage 2 respondent type level. We use Stage 1 data for the balance rates because we want an unconditional (i.e. not recommended) measure of recommenders choices.

More specifically, we calculate the balance rates as follows: (1) we use Stage 1 data / recommender data; (2) we condition on choice pair (e.g., Real Estate Fund A v. Real Estate Fund B) and recommender type (e.g., Democrat); (3) we observe which choice option is most frequent (for instance Democrats might choose Real Estate Fund B disproportionately); (4) we calculate the balance rate.

We calculate resonance rates as follows: (1) we use Stage 2 data / respondent data; we condition on (a) choice pair (i.e. Real Estate Fund A v. Real Estate Fund B), (b) Stage 1 recommender characteristic pair (e.g., Democrat v. Republican), and (c) Stage 2 respondent type (e.g., Democrat); (2) we calculate the resonance rate.

At this level of conditioning, we can see whether the information that Stage 1 recommenders hold about specific choice options affects Stage 2 participants' choices. In addition, further breaking down choices by Stage 2 respondent types allows us to control for the fact that given types might use information to make choices differently.

As we can see from Figure A11, the slope of Stage 2 resonance rates on Stage 1 imbalance rates is negative but very small. This indicates that choice imbalance slightly attenuates resonance rates. This is in line with the predictions of our model with regards to the relationship between information access and resonance. That is, the more information a specific choice option holds, the

less participants will rely on resonance. The negative slope also speaks against a story of resonance being driven by the information content of recommenders.



Figure A7: 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 A8: 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.

If Stage 2 participants disproportionately followed advice from recommenders when those recommenders were more knowledgeable about a given choice option, we would expect a positive relationship between resonance rates and choice balance rates. We look for that relationship next. Figures A9, A10, and A11 below are scatter plots where we have resonance rates on the y-axis (the rate at which respondents choose the congruent option centered at 0.5) and balance rates (the probability that the most frequent option is chosen centered at 0.5) on the x-axis. What differentiates the figures below is the set of choices we use to calculate the resonance and balance rates (in other words, the conditioning level). These different levels of conditioning allow us to answer slightly different questions.

In Figure A9, we condition at the choice level only, for example Real Estate Fund A v. Real Estate Fund B. At this level of conditioning we are focusing on the unconditional relationship between how desirable a specific choice option is and how high resonance rates are. In other words, we are relating indifference about choice options and resonance rates.

In Figure A10, we condition at the choice level, for example Real Estate Fund A v. Real Estate Fund B, and at the Stage 1 and Stage 2 participant type level, for example Democrat. First, looking at the balance rates, we are capturing whether there might be heterogeneity in the unconditional balance rates in Figure A9 by recommender type; for instance, whether Democrats as compared to Republicans are better at picking which real estate fund has a higher return. Second, by relating these balance rates to resonance rates, we can see whether there is a relationship between the information participant types hold and how they make congruent choices. This graph is therefore less about the information content of Stage 1 recommenders and more about the inherent knowledge that Stage 1 and Stage 2 Democrats might have.

In Figure A11, we calculate the balance rates as follows. First, we use Stage 1 data, that is recommender data. Second, we condition on the choice pair, for example Real Estate Fund A v. Real Estate Fund B, and recommender type, for example Democrat. Third, we observe which choice option is most frequent, for instance Democrats might choose Real Estate Fund B disproportionately. Fourth, we calculate the balance rate. We calculate resonance rates as follows. First, we use Stage 2 data, that is respondent data. We then condition on (a) the choice pair, that is Real Estate Fund A v. Real Estate Fund B, (b) the Stage 1 recommender characteristic pair, for example Democrat v. Republican, and (c) the Stage 2 respondent type, for example Democrat. We then calculate the resonance rate. At this level of conditioning, we can see whether the information that Stage 1 recommenders hold about specific choice options affects whether Stage 2 participants disproportionately follow their recommendations.

As we can see from Figure A11, the slope of Stage 2 resonance rates on Stage 1 imbalance rates is negative but very small. This indicates that choice imbalance slightly attenuates resonance rates. This is in line with the predictions of our model with regards to the relationship between information access and resonance. That is, the more information a specific choice option holds, the less participants will rely on resonance. The negative slope also speaks against a story of resonance being driven by the information content of recommenders. In fact, if Stage 2 participants disproportionately followed advice from recommenders when those recommenders were more knowledgeable about a given choice option, we would expect a positive relationship between resonance rates and choice balance rates.




Figure A9: Scatter plot of Stage 2 resonance rates on Stage 1 choice imbalance rates at the choice question (e.g., Real Estate Fund A v. Real Estate Fund B) level. Stage 1 imbalance rates are constructed as follows: condition on a choice question, observe which choice option is most frequent, calculate the probability that the most frequent item is chosen, center at 0.5. Stage 2 resonance rates are constructed as follows: condition on a choice question calculate resonance rates.




Figure A10: Scatter plot of Stage 2 resonance rates on Stage 1 choice imbalance rates at the choice question (e.g., Real Estate Fund A v. Real Estate Fund B) and recommender/respondent type (e.g., 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, center at 0.5. Stage 2 resonance rates are constructed as follows: condition on a choice question, and Stage 2 respondent type, calculate resonance rates.




Figure A11: 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 characteristic pair, and Stage 2 respondent type, calculate resonance rates.

A.4 Resonance and Demand for Information

We now investigate the relationship between resonance rates and information seeking.

As already mentioned in Section 2 and as can be seen from Figure A3, we designed Stage 2 choices such that participants could click on a link to display additional information about each choice option in a given choice pair (except for the True or False choices).

In Figures A12 and A13, we focus on book, charity, and ETF choices, for which participants could seek additional information. We estimate resonance rates at the recommender-characteristic-pair level separately according to whether participants clicked an ‘additional information’ link. Specifically, we interact the recommender-characteristic-pair indicators with a choice-level indicator equal to one if the participant clicked at least one of the two links.

First, we focus on the point estimates in Figure A12. As we can see, resonance rates calculated on choices whereby participants checked the additional descriptions are generally smaller as compared to those calculated on the remaining choices (this is the case for 21 recommender characteristic pairs out of 25). However, separate F-Tests of coefficient equality reveal that only in a minority of cases we can reject that the point estimates on the choices whereby participants checked additional information are different to those whereby participants did not check the additional information (this is the case for 6 recommender characteristic pairs out of 25). Furthermore, when we can reject that the two estimates are not the same, the resonance rate on the choices whereby participants sought more information are always smaller in magnitude.

Second, we add confidence intervals to the point estimates in Figure A13. Here we are interested to see whether we can still reject that resonance rates are equal to zero conditional on information seeking. As we can see information demand does attenuate the majority of the resonance rates, however we see that for 9 out of the 25 recommender characteristic pairs, we can still reject that resonance rates are equal to zero. Ultimately, some of these results are also driven by the smaller number of choices in the recommender characteristic pairs and information seeking cells (as can be seen by the larger confidence intervals).

In summary, we interpret the results above as evidence that at least to some extent subjects use congruence as a substitute for other types of information. The lower resonance rates are ultimately driven by selection, since we allow participants to click on the links to obtain additional information. That is, participants who need more information to make a decision *a priori* are also the participants who will be less likely to strongly rely on resonance.




Figure A12: OLS estimates of congruent choice (centered at 0) on congruence characteristics *interacted with* a dummy equal to 1 for those choices where participants clicked on at least one of the two ‘more information’ links at the bottom of the choice items. Conditioning on books, charities, and ETFs choice questions (since the possibility to check additional information was not present for the True or False choice questions). Blue diamonds show coefficient estimates for choices where participants sought information. Green diamonds show resonance rates for those choices where participants did not seek information. The red bar indicates whether we can reject separate F-Tests of coefficient equality at the 95% confidence level (we use heteroskedasticity-robust standard errors clustered at the participant level for the F-Tests).




Figure A13: OLS estimates of congruent choice (centered at 0) on congruence characteristics *interacted with* a dummy equal to 1 for those choices where participants clicked on at least one of the two ‘more information’ links at the bottom of the choice items. Conditioning on books, charities, and ETFs choice questions (since the possibility to check additional information was not present for the True or False choice questions). Blue diamonds show coefficient estimates for choices where participants sought information. Green diamonds show resonance rates for those choices where participants did not seek information. We plot 95% asymptotic confidence intervals.

B Evidence on Role Models

The academic literature on role models spans psychology, sociology, education, labor and health (e.g. (Gladstone and Cimpian, 2021), (Lawner et al., 2019), (Verniers, 2024), (Downes et al., 2021)).

In public health campaigns, role models—people whom others feel close to or with whom they identify—play a key role. For example, during the COVID-19 pandemic, recommendations from local leaders, such as priests or community doctors, were often more effective than recommendations from celebrities. NPR described efforts 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). In 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 reports 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

C 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 $z \in \mathbb{R}$. 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:

$$s_i = z + \xi_i, \quad \xi_i \sim N(0, \sigma_i^2), \quad (10)$$

where σ_i^2 is the variance of worker i 's signal noise. A lower σ_i^2 , and hence a higher precision σ_i^{-2} , means that worker i transmits more precise and more useful information to the boss. A richer theory might have a boss learn about the worker's knowledge. But we assume σ_i^2 is common knowledge for simplicity. Employees truthfully communicate their signals to the boss.

Let the boss be type b . Each employee belongs to one of two groups, indexed by $\tau_i \in \{a, b\}$. Signal precisions σ_i^{-2} are independently drawn from the same continuous distribution for both worker types, and their realized values are common knowledge.

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

$$\mathbb{E}[(x - z)^2 \mid I_B], \quad (11)$$

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

$$\mathbb{E}[z \mid I_B] = \frac{\sum_i w_i s_i}{\sum_i w_i}, \quad (12)$$

where $w_i = \rho_{ib} \sigma_i^{-2}$ represents the resonance-adjusted weight assigned to worker i 's signal. The weight depends on the signal's precision, σ_i^{-2} , and on worker i 's social distance from the boss through ρ_{ib} . We normalize $\rho_{ib} = 1$ when $\tau_i = b$. When $\tau_i = a$, social distance is positive and $\rho_{ib} = \rho_a \in (0, 1)$.

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:

$$i^* = \arg \max_i \left| \frac{\partial x}{\partial s_i} \right|. \quad (13)$$

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.

Result: Resonance-Based Labor Market Discrimination Resonance implies that the boss assigns greater weight to messages from workers who share the boss's type. Holding signal precision fixed, a type b worker's signal therefore has a greater marginal effect on the boss's action than a type a worker's signal. Consequently, a promotion rule based solely on impact can favor workers who share the boss's type, even though the promotion criterion is formally type-neutral and the boss has no direct preference over the type of worker promoted.

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 distributions of signal precision.

Proposition 5. *Workers who share the leader's type are promoted more frequently.*

Holding fixed the type composition and weights of the other workers,

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

Proof: Let $q_i = \sigma_i^{-2}$ denote worker i 's signal precision. The resonance-adjusted weight assigned to worker i is

$$w_i = \rho_{ib}q_i,$$

where $\rho_{ib} = 1$ if $\tau_i = b$ and $\rho_{ib} = \rho_a \in (0, 1)$ if $\tau_i = a$.

Given the workers' signals, the boss chooses

$$x^* = \mathbb{E}[z \mid I_B] = \frac{\sum_j w_j s_j}{\sum_j w_j}.$$

The marginal impact of worker i 's signal on the boss's action is therefore

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

Because the denominator is common across workers, promoting the worker with the greatest marginal impact is equivalent to promoting the worker with the greatest resonance-adjusted weight.

For any given precision $q_i > 0$, a type b worker receives weight

$$w_i(b) = q_i,$$

whereas a type a worker receives weight

$$w_i(a) = \rho_a q_i < q_i.$$

Because signal precisions are drawn from the same distribution for both types, the distribution of w_i conditional on $\tau_i = b$ first-order stochastically dominates its distribution conditional on $\tau_i = a$.

Let

$$W_{-i} = \max_{j \neq i} w_j$$

denote the greatest weight among worker i 's competitors. Holding the competition environment fixed, stochastic dominance implies

$$P(w_i > W_{-i} \mid \tau_i = b) > P(w_i > W_{-i} \mid \tau_i = a).$$

Because the worker with the greatest weight has the greatest marginal impact and is promoted, it follows that

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

□

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 labor market disparities. Unlike traditional models of discrimination, which rely on taste-based preferences or statistical differences across groups, this framework shows how resonance-based information weighting can generate unequal promotion outcomes under a formally type-neutral promotion rule.

D Labor Market Analysis: Data, Variable Construction and Additional Statistics

D.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,²⁹ 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

²⁹Uninterpretable 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 conclusions 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.

D.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 *Adult Rate*. 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:

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

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

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

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

$$Adult\ ORep_{i,j,k,t} = Adult\ Rate_{i,j,k,t} - Local\ Rate_{j,k,t}$$

Next, we construct the variable *Youth Rate* similarly to *Adult Rate* 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 *Adult Rate* and *Adult ORep* to be gender-specific. Here, *Adult Rate (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 .

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

Adult 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.

$$Adult\ ORep\ (Same\ Gender)_{i,l,j,k,t} = Adult\ Rate\ (Same\ Gender)_{i,l,j,k,t} - Local\ Rate_{j,l,k,t}$$

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

$$Adult\ ORep\ (Wrong\ Gender)_{i,-l,j,k,t} = Adult\ Rate\ (Wrong\ Gender)_{i,-l,j,k,t} - Local\ Rate_{j,-l,k,t}$$

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 receive 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.

$$Job\ Loss_{j,k} = Occupation\ Size_{j,k,2006} - Occupation\ Size_{j,k,2010}$$

$$Percent\ Layoffs_{j,k} = \frac{Job\ Loss_{j,k}}{Occupation\ Size_{j,k,2006}}$$

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

Next, we fix the *Adult ORep* variable to *Adult ORep₂₀₀₅*, which is the average of the *Adult 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 (*Adult ORep₂₀₀₅* was positive) and ethnic groups that were underrepresented (*Adult ORep₂₀₀₅* 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.

D.3 Summary Statistics

To construct the baseline estimates of the *Adult Rate* and *Local Rate* 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.³⁰ Table D1 shows the average number of observations of each ethnicity in each occupation/PUMA/year group. Table D2 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 D3 and D4 provide summary statistics of the constructed variables. Table D3 contains values for the variables dependent on ethnicity only, while Table D4 provides the gender-specific values.

Table D1: Number of Observations of Each Ethnic Group ($Occ\ Size_{i,j,k,t}$ (unweighted))

Ethnicity	Average	Standard	25th	Median	75th	Maximum
		Deviation	Percentile		Percentile	
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

³⁰At 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 D2: Number of People in Each Ethnic Group ($Occ\ Size_{i,j,k,t}$ (weighted))

Ethnicity	Average	Standard	25th	Median	75th	Maximum
	Deviation	Percentile			Percentile	
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 D3: 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
$Adult\ Rate_{i,j,k,t}$	1,203,328	0.031	0.046	0.003	0.014	0.039	0.566
$Adult\ ORep_{i,j,k,t}$	1,203,328	-0.000	0.019	-0.007	-0.001	0.005	0.381
$Adult\ ORep_{i,j,k,t-1}$	1,122,624	-0.000	0.019	-0.007	-0.001	0.005	0.417
$Local\ Rate_{j,k,t}$	1,203,328	0.031	0.042	0.006	0.015	0.039	0.400
$Adult\ ORep_{2005}$	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 D4: Summary of Constructed Variables (Ethnicity, Gender)

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

D.4 Baseline Results

Table D5: Occupation Choice: Ethnicity

	Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adult ORep	0.571*** (0.0503)	0.176*** (0.0137)	0.175*** (0.0137)	0.234*** (0.0133)	0.170*** (0.00874)	0.180*** (0.0106)	0.238*** (0.0133)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes
Observations	1203328	1203328	1203328	1203328	1202976	1050624	1203296

Clustered standard errors in parentheses.

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

Table D6: Occupation Choice: Ethnicity, 1-year Lag

	Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adult ORep (Lag)	0.533*** (0.0513)	0.152*** (0.0142)	0.151*** (0.0141)	0.207*** (0.0139)	0.148*** (0.00905)	0.156*** (0.0110)	0.212*** (0.0138)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes
Observations	1122624	1122624	1122624	1122624	1122208	978592	1122624

Clustered standard errors in parentheses.

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

Table D7: Occupation Choice: Youth Rate on Adult Rate

	Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adult Rate	0.363*** (0.0255)	0.166*** (0.0123)	0.166*** (0.0123)	0.277*** (0.0119)	0.171*** (0.00872)	0.180*** (0.0106)	0.279*** (0.0119)
Local Rate	0.880*** (0.0288)	0.307*** (0.0184)	0.280*** (0.0187)	0.216*** (0.0156)	0.0969*** (0.0209)	0 (.)	0.182*** (0.0158)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes
Observations	1203328	1203328	1203328	1203328	1202976	1050624	1203296

Clustered standard errors in parentheses.

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

Table D8: Occupation Choice: Youth ORep on Adult ORep

	Youth ORep						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adult ORep	0.148*** (0.00751)	0.152*** (0.00752)	0.152*** (0.00753)	0.133*** (0.00790)	0.160*** (0.00808)	0.180*** (0.0106)	0.132*** (0.00798)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes
Observations	1203328	1203328	1203328	1203328	1202976	1050624	1203296

Clustered standard errors in parentheses.

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

D.5 Gender-Specific Results

Figure D1 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 D9, Table D10, and Table D11 show the effects of ethnic-gender overrepresentation on youth rates in regression form.

Figure D1: Relationship Between Ethnic Overrepresentation, Youth Entry, and Gender




Table D9: Occupation Choice: Ethnicity, Gender (Same)

	Youth Rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adult ORep (Same Gender)	0.455*** (0.0350)	0.202*** (0.0118)	0.201*** (0.0118)	0.250*** (0.0127)	0.202*** (0.0102)	0.216*** (0.0121)	0.253*** (0.0128)	0.333*** (0.0262)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes	.
Occupation $\times$ PUMA $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	.	Yes
Observations	2108064	2108064	2108064	2108064	2107904	2056448	2108032	1809472

Clustered standard errors in parentheses.

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

Table D10: Occupation Choice: Ethnicity, Gender (Wrong)

	Youth Rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adult ORep (Wrong Gender)	0.286*** (0.0307)	0.00421 (0.0121)	0.00370 (0.0121)	0.00326 (0.0131)	-0.00781 (0.0102)	-0.0146 (0.0122)	0.00433 (0.0130)	-0.333*** (0.0262)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes	.
Occupation $\times$ PUMA $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	.	Yes
Observations	2108064 2108064 2108064 2108064 2107904 2056448 2108032 1809472							

Clustered standard errors in parentheses.

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

Table D11: Occupation Choice: Ethnicity, Gender (Both)

	Youth Rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adult ORep (Same Gender)	0.406*** (0.0296)	0.221*** (0.0131)	0.221*** (0.0131)	0.253*** (0.0132)	0.218*** (0.0121)	0.230*** (0.0139)	0.255*** (0.0133)	0.333*** (0.0262)
Adult ORep (Wrong Gender)	0.157*** (0.0227)	-0.0636*** (0.0129)	-0.0639*** (0.0129)	-0.0225* (0.0127)	-0.0634*** (0.0117)	-0.0642*** (0.0136)	-0.0203 (0.0126)	0 (.)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes	.
Occupation $\times$ PUMA $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	.	Yes
Observations	2108064	2108064	2108064	2108064	2107904	2056448	2108032	1809472

Clustered standard errors in parentheses.

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

Table D12: Auxiliary Regression: Wrong Gender on Same Gender

	Adult ORep (Wrong Gender)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adult ORep (Same Gender)	0.314*** (0.00691)	0.302*** (0.00663)	0.302*** (0.00663)	0.101*** (0.00761)	0.251*** (0.00647)	0.213*** (0.00816)	0.0959*** (0.00769)	-1.000*** (9.32e-16)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	.	.	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	.	.
Occupation $\times$ PUMA FE	.	.	.	.	Yes	.	.	.
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.	.
Occupation $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	Yes	.
Occupation $\times$ PUMA $\times$ Ethnicity $\times$ Year FE	.	.	.	.	.	.	.	Yes
Observations	2108064	2108064	2108064	2108064	2107904	2056448	2108032	1809472

Clustered standard errors in parentheses.

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

Table D13 shows a FWL decomposition of the regression coefficient on *Adult ORep (Wrong Gender)*. Consider the following short regression (SR), long regression (LR), and auxiliary regression (AR):

$$\text{LR: } Youth\ Rate_{i,l,j,k,t} = \alpha_0 + \alpha_1 Adult\ ORep\ (Same)_{i,l,j,k,t} + \alpha_2 Adult\ ORep\ (Wrong)_{i,l,j,k,t} + \varepsilon_{i,l,j,k,t},$$

$$\text{SR: } Youth\ Rate_{i,l,j,k,t} = \delta_0 + \delta_1 Adult\ ORep\ (Wrong)_{i,l,j,k,t} + v_{i,l,j,k,t},$$

$$\text{AR: } Adult\ ORep\ (Wrong)_{i,l,j,k,t} = \gamma_0 + \gamma_1 Adult\ ORep\ (Same)_{i,l,j,k,t} + \eta_{i,l,j,k,t},$$

where the coefficients in the LR can be read in Table D11, the coefficients in the SR can be read in Table D10, and the coefficients in the AR can be read in Table D12.

This set-up allows us to interpret the coefficient on *Adult ORep (Wrong Gender)* in Table D11 formally as

$$\delta_1 = \alpha_2 + \gamma_1 \alpha_1 \quad \Longleftrightarrow \quad \alpha_2 = \delta_1 - \gamma_1 \alpha_1. \quad (14)$$

This decomposition is shown in Table D13, where we can clearly see that the reason why the coefficient on *Adult ORep (Wrong Gender)* in the long regression (Table D11) is smaller than that in the short regression (Table D10) is because the coefficients on *Adult ORep (Same Gender)* in the long and auxiliary regressions dominate it.

Table D13: Decomposing the Adult ORep (Wrong) Coefficient in the Long Regression

Spec	(Wrong SR: δ_1) – (Same AR: γ_1) $\times$ (Same LR: α_1) = (Wrong LR: α_2)							N
(1)	0.2863	–	0.3138	$\times$	0.4061	=	0.1569	2,108,064
(2)	0.0042	–	0.3024	$\times$	0.2209	=	-0.0636	2,108,064
(3)	0.0037	–	0.3020	$\times$	0.2207	=	-0.0639	2,108,064
(4)	0.0033	–	0.1012	$\times$	0.2528	=	-0.0225	2,108,064
(5)	-0.0078	–	0.2514	$\times$	0.2182	=	-0.0634	2,107,904
(6)	-0.0146	–	0.2130	$\times$	0.2296	=	-0.0642	2,056,448
(7)	0.0043	–	0.0959	$\times$	0.2550	=	-0.0203	2,108,032
(8)	-0.3326	–	-1.0000	$\times$	0.3326	=	0.0000	1,809,472

D.6 Economic Shock Results

Table D14: Impact of a Shock on Occupation Choice

	Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adult ORep (2005)	0.623*** (0.0866)	0.139*** (0.0335)	0.153*** (0.0338)	0.323*** (0.0347)	0.147*** (0.0291)	0.170*** (0.0318)	0 (.)
POST=1	0.0380*** (0.000697)	0.000830** (0.000346)	0 (.)	0.000788** (0.000346)	0.000774** (0.000359)	0 (.)	0.000798** (0.000365)
POST=1 $\times$ Adult ORep (2005)	-0.01000 (0.0471)	-0.0182 (0.0391)	-0.0468 (0.0396)	0.00929 (0.0390)	0.00523 (0.0403)	-0.0456 (0.0416)	0.00557 (0.0412)
High Layoffs=1	0.0193*** (0.000535)	0.00163*** (0.000434)	-0.00000514 (0.000425)	0.00143*** (0.000416)	0 (.)	0 (.)	0 (.)
High Layoffs=1 $\times$ Adult ORep (2005)	-0.0450 (0.122)	0.249*** (0.0664)	0.225*** (0.0661)	0.153** (0.0637)	0.262*** (0.0559)	0.229*** (0.0613)	0 (.)
POST=1 $\times$ High Layoffs=1	-0.0407*** (0.000824)	-0.00327*** (0.000539)	0.000127 (0.000503)	-0.00310*** (0.000537)	-0.00307*** (0.000557)	0 (.)	-0.00317*** (0.000568)
POST=1 $\times$ High Layoffs=1 $\times$ Adult ORep (2005)	-0.169** (0.0831)	-0.241*** (0.0763)	-0.184** (0.0743)	-0.252*** (0.0752)	-0.259*** (0.0759)	-0.204** (0.0793)	-0.259*** (0.0759)
Occupation FE	.	Yes	Yes	Yes	Yes	Yes	Yes
Occupation $\times$ Year FE	.	.	Yes	.	.	Yes	.
Occupation $\times$ Ethnicity FE	.	.	.	Yes	.	.	Yes
Occupation $\times$ PUMA FE	.	.	.	.	Yes	Yes	Yes
PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
PUMA $\times$ Ethnicity FE	.	.	.	.	.	.	Yes
Occupation $\times$ PUMA $\times$ Year FE	.	.	.	.	.	Yes	.
Occupation $\times$ PUMA $\times$ Ethnicity FE	.	.	.	.	.	.	Yes
Observations	275427	275427	275427	275427	274986	242493	273211

Clustered standard errors in parentheses.

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