Information Resonance

Ulrike Malmendier and Laura Veldkamp*

August 22, 2022

Abstract

People process the same information differently depending on who delivers it. Information resonates with recipients when they identify with the person who communicates it or whose personal experience it reflects. Resonant information imprints itself in our memory and comes to mind when making relevant decisions. We describe and formalize the phenomenon of resonant information being more heavily weighted in the process of belief updating. The model delivers an empirical tool to identify which social characteristics determine resonance. We then work through the model implications for social transitions in response to COVID-type shocks, for the influence of role models on behavior change such as vaccine adoption, and why social media changes the influence of experts. Evidence on occupational choices within geography-ethnicity-gender groups, and their responses to labor market outcomes inside and outside these groups, suggest that resonance may be a powerful force.

Individuals and communities of people with different socioeconomic and demographic characteristics make systematically different choices in many important dimensions, including economic, financial, education, and health-care. In education, for example, we continue to see persistent gender gaps in college major choices, particularly the representation of women in STEM fields (Carrell et al., 2010). In the health sector, a large portion of the black-white male cardiovascular mortality gap has been attributed to the hesitance of Black men to receive preventative care (Alsan et al., 2019). Most recently, we have seen these demographic differences in the disproportionate impact of COVID-19 along racial lines and the variation in vaccine uptake and safety precautions (Alsan et al., 2021; Torres et al., 2021).

What explains the persistent differences along socio-demographic lines? One determinant often highlighted in economic research is access to information. A large literature on information diffusion has linked uneven access to information to precisely the types of choices outlined above, from economic and financial choices in developing countries, to schooling in underserved communities (Cai et al., 2015; Banerjee et al., 2013; Foster and Rosenzweig, 1995; Conley and Udry, 2010)

However, as information technology has made knowledge ever more accessible, we might have expected communities with access to similar information to develop more similar beliefs and make more similar choices. Yet, economic decisions and incomes of neighborhoods continue to diverge.

*We thank Nikki Azerang, Rishi Chandran, and Xiaobo Yu for excellent research assistance.

For example, as rates of internet access climb ever higher, we still see educational choices diverge across communities, house price beliefs have a strong local component, and even behavior in response to COVID depends on communities (Fogli and Guerrieri, 2019; Kindermann et al., 2021; Kuchler et al., 2020b). To some extent, differences in preferences and financial constraints play a role here. But why don't we see more convergence in beliefs if access to information is more extensive then ever before?

In this paper, we argue that the prior emphasis on "access" to information, or lack thereof, neglects an important determinant of actual information processing, which we call *information resonance*: Whether or not people account for information, and how strongly it features in their decisions, depends on how firmly the information is anchored in memory. Some information is transmitted, but not internalized. The person learning the information does not find it as relevant, and does not feel the importance of this experience in a way that will encode it in their memory for use in future decision making. For example, abstractly learned statistics and other information tends 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). This is not a question of "limited attention" or cognitive limitations, as frequently modeled in economics. Recipients did not "miss" the expert information in question, and they might even be able to reproduce and recite it, but it does not resonate with them. They simply do not identify with the person conveying the information and, as a result, put less weight on it in their own decision-making. For example, hearing from a faculty colleague that they will not vaccinate their 8-year old kid before longer-run studies are out might give us pause and resonate more with us than a statement by Anthony Fauci. Or, the investment trading behavior of a family member might inspire and influence us more than advice about longterm investment by Warren Buffet.

At the same time, it is not always ex-ante obvious whose experience will resonate with us. Standard socio-economic characteristics are a natural starting point. For example, 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), with 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 can significantly increase vaccination rates (Larsen et al., 2022).

But it also turns out that the relevant characteristics vary by informational setting and by characteristics of the observer. In the health context, for example, Alsan et al. (2021) found that messages from same-race doctors increased COVID-related information-seeking among Black people only when coming from doctors they identified with (because those doctors acknowledged unequal treatment in health care or economic hardships that were associated with tight living quarters and working in essential services). In the macro-finance realm, D'Acunto et al. (2021) show that women and Black men update their unemployment and inflation expectations in line with FOMC forecasts when the presence of a woman or Black man on the FOMC is made salient, but White women and Black men *both* respond to the presence of a White woman or a Black man, possibly because of a greater "taste for diversity" among these groups compared to White men. Most strikingly, in the household-finance realm, Stolper and Walter (2019) found that clients were more likely to follow financial advice from advisors with whom they shared multiple demographic characteristics, but that there is also heterogeneity among which demographic characteristics are given the greatest weight. For instance, women identify most with the parental and marital status of their advisor, while men identify most with age and gender.

In this paper, we aim to provide a simple framework to better understand the concept of "information resonance," and to provide a measurement tool that allows to identify the determinants of such resonance and, thus, action-relevant flow of information.

We start from exploring the idea that people with heterogeneous characteristics attribute more relevance to experiences and information that come from others they perceive to be like them, i. e., from the experiences of their community members and social circles (including their own experiences).[1] We aim to distinguish such *resonance effects* from the true payoff relevance, in the sense of a higher correlation in payoffs or other outcomes among people with similar characteristics. As a result of information-resonance effects, people with heterogeneous characteristics display heterogeneity in beliefs and choices that follow socio-economic, racial, ethnic, and political boundaries above and beyond those explicable by the covariance in characteristic-related outcomes. Such inequality is difficult to remedy with traditional information campaigns and redistributive policies. While other barriers, such as discrimination, surely exist, one important barrier that may keep the poor from escaping poverty and other adversities (such as adverse health outcomes) is the social barrier to information transmission.

Section 1 proposes a framework for thinking about information diffusion when social proximity determines how much information "resonates," i. e., how much weight individuals put on an action-outcome pair they observe. The model applies standard information diffusion mechanisms to a Bayesian learning model, where social proximity tilts the updating process. It illustrates how information spreads more quickly among people who share similar characteristics. The vector of

[1] It is no accident that community and communicate derive from the same Latin origin. Communities are and have always been groups of people who communicate or share information.

social characteristics could represent anything that might make one "identify" more or less with another person: gender, ethnicity, race, caste, class, income, location, disability, weight, height, other aspects of physical appearance, sexuality, language, political affiliation, religion, education, or even affiliation with a sports team.

We use this model to contrast the role of information resonance with that of information access. Initially, we allow people to learn from everyone else's experiences. That is, in the initial benchmark case, there is full access to information about all actions and resulting realizations of the variable(s) of interest. However, people give more weight to the experiences of people like them, including their own experiences. This could be the optimal use of all available information in the sense of full Bayesian updating. For example, one might give more weight to the experiences of "people like me" since the economic shocks are highly correlated within socially similar groups. Or it could reflect suboptimal but "neurologically realistic" Bayesian updating – the experience of someone that I identify with triggers a different emotional response and encodes itself as a more durable neural connection.[2]

Our model directly accounts for the actual correlation in outcomes across individuals, and contrast the predictive weight that would be assigned by a pur Bayesian learner with the resonance weight.

Section 2 shows that, even when all experiences are observable to everyone, the beliefs and actions of different communities, or social circles, can diverge for long periods of time. The more people discount the experience of socially distanced households, the slower the transmission and ultimate convergence across groups, and the longer "information ghettos" persist. This model set-up provides a new perspective to the debate on how to address, say, income and racial inequalities, rural poverty and the differential rates of economic progress or vaccination rates of new immigrant groups. As we illustrate in our first result, the effect of providing more and better information might pale in comparison to the effect of resonant information.

We consider a setting, in which initially no member of a community chooses a dominant action. Think of the choice to obtain more education or take beneficial health measures, such as getting vaccinated. In one scenario, multiple people in a neighboring community take the better action (that is not adopted in the home community) and provide highly visible information about the superior outcomes to the home community. In the other scenario, a single community member who is centrally located in the social characteristic space is subsidized to experiment with the new action. We show that the action of a single but central community member can be a more powerful force for change than numerous visible actions by people that the community perceives to be "not like us."

[2]This is related to the neuroscience work on synaptic plasticity: Abraham et al. (2019), McNaughton et al. (1978); Long-Term Potentiation (LTP) as underlying mechanism for learning and memory: Bliss and Lømo (1973), Malinow (1991), Otto et al. (1991), Morris et al. (1990), Bear (1996)

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

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

Section 3 turns from the theoretical modeling of information resonance to its empirical identification. Measuring resonance is useful because knowing who is persuasive can help craft effective policy. For example, having the right people convey public health measure might induce better compliance. We first illustrate why it can be difficult to identify the relevant dimensions of resonance and how confounds (sorting on irrelevant dimensions of social distance) can obfuscate the actual information-diffusion pattern. We then show how to use our model to extract the true underlying determinants of information diffusion from observed data.

Finally, in Section 4, we undertake a first step towards applying the theoretical insights to measure information diffusion in the data. We use data from the American Community Survey (ACS) from 2005 to 2020 to provide evidence on information resonance influencing occupational choices. In particular, we show that the occupational choices of 18-22 year olds are strongly influenced by the prior choices of same-ethnicity and same-gender older people living in the same

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

geographic area. To conduct this analysis, we construct a novel classification of ethnic groups based on ancestries. Using this classification, we find that, when in a given area (PUMA) and a given year, an ethnicity is overrepresented in an occupation j, then a young person from the same ethnicity is significantly more likely to choose occupation j. To help interpret this result, we further subsample by another social characteristic, gender. 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. The estimation holds after controlling for Occupation $\times$ PUMA $\times$ Year fixed effects or, alternatively, Occupation $\times$ PUMA $\times$ Ethnicity fixed effects, and – when distinguishing by gender – even Occupation $\times$ PUMA $\times$ Ethnicity $\times$ Year fixed effects. We also show that occupational layoff shocks influence the job choices of the younger generation, only if their ethnicity was overrepresented.

These results can be interpreted in three ways. First, correlated occupational choices among same ethnicity-gender-PUMA peers might reflect correlated payoffs. The inclusion of fixed effects, however, makes this first interpretation unlikely: even after controlling for the popularity of an occupation, for labor-market time trends, for ethnicity-specific occupational choices, and for geography-specific over- or under-representation of certain occupations, individuals appear to be strongly influenced by peers to whom they are socially close. Attributing those choices to group-specific payoffs, or abilities, would require that members of a certain ethnicity and gender, say young male Southeast Asians, are particularly likely to choose an occupation in Hospitality or related service occupations in, say, The Big Island, Hawaii, where the fraction of older male Southeast Asians in the same service occupations is particularly high, but not in parts of Snohomish and King county Washington, where the fraction of older male Southeast Asians is instead high in manufacturing trades, particularly machine operators, and younger male Pacific Islanders follow suit. This is the case even if the fraction of female Southeast Asians in Washington is over-proportionally high in service occupations.[4]

The latter aspects also addresses a second alternative interpretation – that the patterns in occupational choice are related to information diffusion, but of the traditional (non-resonance) type. Same-ethnicity people in the same area communicated information about job opportunities to each other, possibly also helped by a lower language barrier. The fact that same-ethnicity but other-gender representation does not have any positive effect points to peers being "role models" rather than merely providing access to otherwise unavailable information.

The third interpretation, proposed here, is that job choices among socially close community

[4]In 2006, Southeast Asian men in the Big Island were overrepresented in service occupations by 0.19 (34% of men of Southeast Asian decent worked in service occupations compared to 15% of people in the region). At the same time, men of the same ethnicity in Snohomish and King county were overrepresented in machine operations by about 0.9 (12% of Southeast Asian men worked as machine operators compared to only 3% of people in the region). While Southeast Asian women in this region were overrepresented in service occupations by 0.13 (20% compared to 7% baseline), men of the same ethnicity were actually slightly underrepresented by about 0.004.

members resonate more, as do layoffs in industries where those community members are overrepresented.

In a next step, we plan to apply the measurement implementation of our theoretical tool directly and use machine-learning techniques to identify and compare the roles of different determinants of occupational choice.

In summary, this paper adapts standard tools from diffusion and social networks to formally represent information resonance, with a different measure of distance leading to a different information weighting scheme. Despite the formal similarity to existing models, adopting the lens of information resonance, instead of information access, is an important shift in thinking. In many situations, people do not need more information, they need information that speaks to them in an emotionally meaningful way. Our model suggests that policies that rely on simple information provision will fail. Information and personal experiences need to be conveyed by people who resonate with their communities. This change in focus, from information access to information resonance, could greatly enhance the effectiveness of a broad array of economic and social policies. This simple change of focus could also change the way we think of idea diffusion and the process of economic progress.

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

Those finding 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.[5] 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).

In the development sector, literature on information diffusion, particularly "seeding," has shown that community members are more effective than governments in spreading information (Alatas et al., 2016). Not only are they able to spread the information, but they can actually identify which members of their social network will diffuse the information most efficiently (Banerjee et al., 2019). Banerjee et al. (2019) dub these individuals "gossips" and note that while their ability to circulate information is largely due to their diffusion centrality, their position in the network is not at all based on leadership status, education, or even their physical location in the village. Moreover, the authors emphasize that diffusion centrality fails to fully explain why gossips are good at spreading information and note that there may be unobserved factors used to identify "gossips," including who is most listened to/trusted.[6]

1 Model

We begin by developing a framework and a measure that captures the idea of learning based on social proximity, in as simple a way as possible. At the same time, we build in the potential for standard information frictions as well, so that we can compare and contrast our new idea with the existing ones about access to information, and some outcomes being unobservable.

The main ingredients of the model we propose are

1. Agents are uncertain about the distribution of a stochastic payoff.

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

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

2. Agents learn from observing experiences (action-payoff pairs) of other agents in the economy, including their own.

3. Agents with heterogeneous characteristics θ weight the information that comes from others with nearby characteristics more heavily.

Further, in order to compare our mechanism with traditional economic mechanism, we add variation in the access to information, as a feature that we can turn on or off.

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

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

$$U(a_{it}) = a_{it} E[z_a + \epsilon_{it} | \mathcal{I}_{it}]. \tag{1}$$

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

Information sets. If agent i takes the action $a_{it} = 1$ at time t, then each agent j potentially sees their action and payoff outcome. Define an observation matrix Ψ_t, where the (i, j)th entry of Ψ_t is 1 if i observes j and j takes action $a_{jt} = 1$.

Ψ_t might have a component that is not time-varying. For example, if i sees only people in her neighborhood or only facebook friends, then 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 a $n \times 1$ vector of every agent's time t action and $\mathbf{1}'_n$ is a $1 \times nh$ vector of 1's. Alternatively, we can consider people who are randomly visible. Or celebrities who are visible with a higher probability than others. This would be captured by a Ψ_t matrix with a random component.

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

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

Belief updating For agent i taking action $a = 1$, belief updating takes the form:

$$E[z_i|\mathcal{I}_{it}] = \alpha_{iz}\mu_z + \overline{\omega}_{iz} \sum_{j=1}^{N} \sum_{t'=1}^{t} \omega_{ij}\beta_{izj}(z_j + \epsilon_{jt'}), \tag{2}$$

where α_{iz} and β_{izj} are the Bayesian weights on priors and signals:

$$\beta_{iz} =(\mathcal{Z}_{it}\Sigma_z\mathcal{Z}_{it}' + \sigma_\epsilon^2 I)^{-1}\mathcal{Z}_{it}\Sigma_z\mathbf{1}_i \tag{3}$$

$$=\frac{Cov(z_i, z + \epsilon)}{Var(z + \epsilon)} \tag{4}$$

$$\alpha_{iz} = 1 - \sum_{j=1}^{N} \beta_{izj}. \tag{5}$$

The indicator variable $\mathbf{1}_i$ is a $N \times 1$ vector of zeros with 1 in row i, and I is an identity matrix with size K_{iz}. The equality in (4) shows that the Bayesian weight on signals is simply the standard OLS estimator derived from linear models.

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

$$\overline{\omega}_{iz} \sum_{j=1}^{N} \sum_{t'=1}^{t} \omega_{ij}\beta_{izj} = \sum_{j=1}^{N} \sum_{t'=1}^{t} \beta_{izj}. \tag{6}$$

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

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

Belief weights and characteristics. For an observer i, different observations have different relevance. When i observes j, the weight i places on that observation, 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. Then, if i observes j's actions,

$$\omega_{ij} = \underbrace{(2 - 2\Phi(\chi||\theta_i, \theta_j||))}_{} \cdot \sigma_\epsilon^{-2} \tag{7}$$

$$= \quad \rho_{ij} \quad \cdot \sigma_\epsilon^{-2}, \tag{8}$$

where $||\theta_i, \theta_j||$ is the Euclidean distance between agents' characteristics θ_i and θ_j.[7] 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 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 and then declines quickly towards zero. These terms define the resonance, or relevance, ρ_{ij} that a signal coming from j has for observer i.

We scale the resonance weight ρ with the (true) accuracy of the signal, σ_ϵ^{-2}, to reflect the idea that agents will not weight signals that have no information content at all. Thus, the overall weight ω_{ij} is the product of the resonance weight ρ_{ij} and the precision weight σ_ϵ^{-2}.

Equilibrium Everyone acts simultaneously in every period, knowing the outcomes of the previous periods. Agents maximize expected utility as specified in (1). They update their prior of the unknown payoffs z_a for each action a using Bayes' rule as spelled out in equations (2), with weights ω_{ij} as defined in (7)-(8).

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

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

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.

Interpretations or mechanisms underlying the assumptions What do we mean to capture with the "characteristic θ"? As discussed in the introduction in the context of social-learning theory, a characteristic could be gender, familiarity, race, or distance in a social network. We could also think of religion, political affiliation, or socio-economic status. The idea is that people identify with, and internalize, the experiences of others who are like them or socially connected to them. A characteristic could also be the time at which the data point was observed. This would discount older data. These types of characteristics could all feed into information resonance.

What 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 the information access matrix $\mathcal{Z}_{it}$, which would be related to the entries of θ_j.

Resonance alters memory or how beliefs are formed. The most literal interpretation of our model is that (2) represents a belief formation process. Some events get weighted more than others when we form beliefs. Another possibility is that this is the average belief formation process. Perhaps people do form beliefs according to Bayes Law, but some information is forgotten. If less resonant information is more likely to be dropped from memory, then (2) could represent the effective average beliefs, after adjusting for forgotten data and applying a standard verison of Bayes Law.

Resonance alters how beliefs are used to choose actions. Another possibility is that agents have beliefs and then they have some other criteria they use to choose actions. One could argue that this is simply a pedantic difference about what then is a belief. But the relevance of that difference is that survey data might be inconsistent with chosen actions. When asked what they believe, people may report Bayesian beliefs. But when choosing a action that has uncertain payoffs in their own lives, they may twist those beliefs in the direction of more resonant information. Other theories in which agents twist beliefs include ambguity aversion or robust control.

A rational interpretation. Finally, it is possible that what we call resonant information is simply more statistically relevant information, in some settings. 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 (2) 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, some of our examples highlight applications, such as a vaccine, in which the variability of payoffs across individuals is very unlikely to depend so significantly on social characteristics. In other words, resonance may be seeded by a grain of truth: Some choices have payoffs that are correlated with group characteristics. However, these characteristics might also be over-weighted, relative to their information content, when forming beliefs or choosing actions.

Comparing related learning mechanisms. The idea of information diffusion among communities traces back more than half a century:

> While mass media play a major role in alerting individuals to the possibility of an innovation, it seems to be personal contact that is most relevant in leading to its adoption. Thus, the diffusion of an innovation becomes a process formally akin to the spread of an infectious disease (Arrow, 1969).

In our setting, contact does not "infect" someone with a belief (as in Burnside et al. (2016)). Instead, an accumulation of evidence changes their opinion.

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. Work on experiential learning has highlighted the importance of one's own experiences in shaping beliefs and economic decisions (e.g., Malmendier and Nagel (2011, 2016)) or the historical experience of one's family (Boerma and Karabarbounis, 2017).[8]

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. There is no relationship between social distance and information weights. Theories of belief scarring show that when we learn from our collective experiences, transitory but extreme outcomes can have long-lived effects on aggregate outcomes (Kozlowski et al., 2020).

In between these polar cases, cities or metropolitan areas have beliefs that evolve differently because of their different experiences (Chodorow-Reich et al., 2021). If we think of the characteristic vector $0 < \chi < 1$ including one's location or metro area, the model becomes one of local learning from neighbors, as in (Fogli and Veldkamp, 2010; Burnside et al., 2016). This connects our framework to an extensive literature on spatial diffusion of technology (Comin et al., 2013), and the evidence on ideas spreading from one farm to an adjacent farm in developing countries (Conley and Udry, 2010; Munshi, 2004). Instead of a technology spreading, we consider information about the returns to college, to entrepreneurship, or the pathway to a new career spreading among

[8]Learning from own experiences is also a feature of the literature on active experimentation (Bergemann and Välimäki, 2000). The difference is that one learns not only from pulling the bandit's lever one's self. In our setting, people also learn from seeing the outcomes of others nearby in characteristic space. They can also free ride on others to produce relevant information for them to learn from.

community members, and substitute geographic distance with broader notions of social distance. If we allow for the observation Z_{ij} to be a function of the physical distance between i and j. In this view of the world, all information that is observed is weighted by its precision. Geography simply limits what is observed. [9]

Dynastic learning 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. People learn only from the past experiences of their own family members.

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.

Outside of economics, the notion of information resonance, and its determinants, builds on a long literature on social learning, as discussed above. We will discuss more of the related theory literature after we have introduced our model.

2 Illustrative Numerical Examples

We start by contrasting the essential difference between "information resonance" and the traditional emphasis on "information frictions," and then explore how the two forces interact. Information frictions prevent people from obtaining, or observing, information. Information resonance captures how people weight the information they see. Some information affects us much more: the horrible act committed on a child that looks just like one's own, the tragedy experienced by a close friend, or the victory achieved by someone in a similar position, that you can envision being you the next time around. These events imprint themselves in our psyche and come to mind when making relevant decisions. When information access changes, people can access more resonant information. This causes them to downweight less resonant, but more reliable expert opinions.

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

2.1 Parameter Choices and Procedure

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

The parameters we use are as follows: Action payoffs are 0 for action $a = 0$ and $z_i = 4\ \forall i$, meaning that all agents have the same persistent component of payoffs to action $a = 1$. The individual transitory shocks to payoffs are very noisy, with $\epsilon_{ait} \sim N(0, 100)$ for both actions a, independently drawn for actions, agents, and time. Thus, in the absence of taxes or subsidies the total payoff to actions, net of action costs are: Action 0 has payoff $\sim N(0, 100)$. Action 1 has payoff $z_1 + \epsilon_{it} \sim N(4, 100)$. That is, action 1 is the better (higher expected payoff) action. If agents knew the true value of z_1, they would maximize their utility, as specified in (1), by choosing action 1.

Prior beliefs for action 1 are $z_1 \sim N(-3, 10)$. So agents believe that action 0 is superior, when in truth, it is not. The action costs and subsidies are known to all agents. Agents update their beliefs according to (2).

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

The simulations below differ only in their initial conditions, including the distribution of characteristics (uniform or normal), the presence of taxes (or subsidies). What they all share in common is the following procedure:

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

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

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

4. Draw random payoff shocks (ϵ_{it}) for each chosen action.

5. Allow agents to observe the total payoff $z_j + \epsilon_{jt}$ for each $j \epsilon \tilde{\mathcal{Z}}_{it}$. Add this new data to each agent i's information set and update each agent's beliefs, according to (2).

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

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

2.2 Representing Diffusion in Characteristic Space

Once we do a change of distance measure, this model becomes simple to solve and to understand. We start by describing what this new measure (a topology) is and what model outcomes looks like in geographic space and in this new characteristic space. The technical insight is that after re-casting the outcomes in this new topological space, resonance functions just like an information diffusion friction. Armed with this insight, our predictions look pedestrian. However, the practical challenge will be knowing what the right distance measure is, or equivalently, what dimensions on which to plot the data. Sections 3 and 4 wrestle with this measurement challenge.

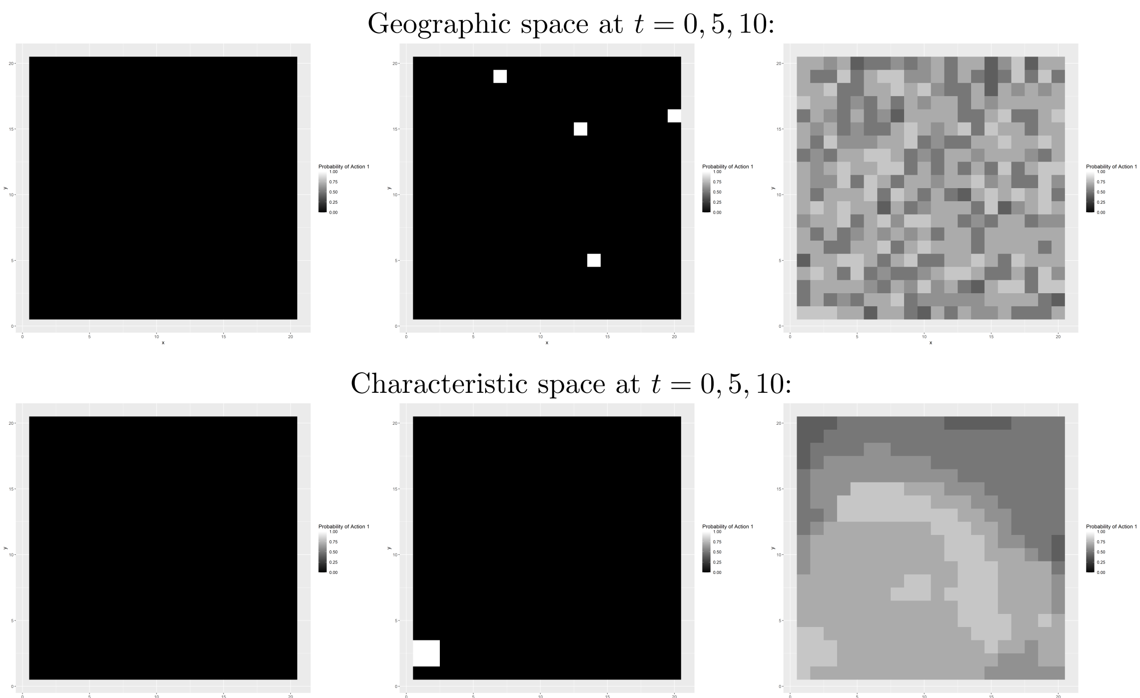

Figure 1: A diffusion process represented in geographic space (top) and characteristic space (bottom).

Figure 1 plots a diffusion process that results from our model. For now, all that matters is that the top row of panels and the bottom three panels represent the same result in two different ways: in geographic space versus characteristic space.[10] In the top row, all 400 agents are sorted along the x- and y-axes that represent their geographic location. In the bottom row, instead, they are sorted a long the v- and w-axes, which represent their social characteristics.

[10] Appendix B.2 describes the simulation and parameters that give rise to these results.

We illustrate the choice of action 0 in black, and the choice of action 1 in white. If the fraction of times an agent choose action 1 (across simulations) lies between 0 and 1, we choose shades of grey accordingly. For example, the leftmost panel in both rows is entirely black, illustrating that all agents choose action 0 in all simulations of period 0.

The top row shows how learning about the higher-payoff action 1 spreads across geographic space. We see that, by period 5, four agents have started choosing action 1, and by period 10 all agents have chosen action 1 at least in some period. While it is always possibile to have a simulation with many negative payoff shocks such that no agents switch to action 1, the light grey colors of squares tells us that this happens very infrequently. Importantly, notice that the four initial switchers are not located nearby. And the jumble of colors in the rightmost panel indicates that adjacent people may transition from action 0 to action 1 at very different rates. An economist looking at this data would not see any geographic patterns and might well conclude that no diffusion is at work. Instead, this looks like a sequence of independent shocks.

If, instead of plotting people according to their geographic location, we now organize them according to their characteristics, the pattern becomes clear. One node, with characteristics in the bottom left of the plot initially switches from action 0 to action 1 at random (given her random payoff realization $z_0 + \epsilon_{0i,t=1}$ in period 1). Three nodes with similar characteristics quickly follow suit. Then, as others with nearby characteristics (i. e., those in close social proximity) observe the actions and payoffs of others they identify with, that information resonates and causes then to change actions as well. The bottom row of panels looks like a classic diffusion model. But to see this classic diffusion, one needs to understand that information only gets incorporated if it resonates. Understanding the nature of resonance leads one to use a different notion of distance, a new topology, based on distance in characteristic space.

The following subsection explores what prompts the diffusion of a new action, represented in characteristic space.

2.3 Experts vs. Role Models

We now illustrate the basic insight emerging from our learning mechanism: the differing influences of "access to information," which is at the core of many traditional learning models, versus "social proximity" (information resonance) to the underlying realizations.

We denote someone whom many people feel socially close to a 'role model.' Their actions and experiences resonate with many people. That is, others put significant weight on and thus learn from the information provided by role models. We contrast the role model with a person about whose action and experiences many people are informed about, which we dub a 'celebrity.' The contrast between the influence of role models versus celebrities mirrors the difference between the effect of "social proximity" (information resonance) versus "access to information" that we would like to illustrate. For example, we saw during the COVID pandemic that famous people recommending vaccines did not necessarily work. The fact that many other people in the country got vaccines with no adverse effects also seemed to have little influence. However, having a local leader, like a priest or community doctor, recommend the same action was more effective. Examples of such local influence 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.[11] 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.[12]

[11]NPR reported about the efforts of to "keep an open line of communication with recalcitrant rabbis ... knowing they would have the most influence to convince their followers to get vaccinated," and the ultimate success of this approach: "The rabbis agreed to be vaccinated — and their adherents followed suit" (April 22, 2021). The Australian ABC Network describes how "the knowledge and trust Aboriginal health organizations had with local communities was another reason for Victoria's success" in convincing the Aboriginal population, in particular by "re-framing government messaging for local communities" (August 19, 2021). As for the U.S., the PR Newswire explains, "With the disproportionate impact of the pandemic in rural areas across the country, it's critically important that Americans in rural communities hear these powerful stories from each other, their neighbors" (October 19, 2021). The Villager report that "Union leaders receive COVID-19 vaccine in Harlem, hoping to inspire others to do the same" (February 27, 2021). And the New York Times attributes the narrowed vaccination gap between Black and other racial groups to "decisions made in many states to send familiar faces to knock on doors and dispel myths about the vaccines' effectiveness," such as the case of a store owner and county commissioner in Panola, a small rural town near the Mississippi border, who "led the effort to vaccinate nearly all of her majority Black community," and pro-vaccination campaigns including descendants of participants in the Tuskegee study (October 13, 2021). The campaign successfully eased vaccine hesitancy derived from medical mistrust by positioning community members who had directly suffered from the betrayals of the healthcare industry as trusted informants. The Ad Council noted that "[the descendants] are so perfect to be the storytellers because if they can move past it and see the lessons and drive towards getting people vaccinated, it feels like everyone should be able to do the same."

[12]`www.ccbh.net/wp-content/uploads/2021/09/Field-Guide_CDC_Final_2021.pdf`

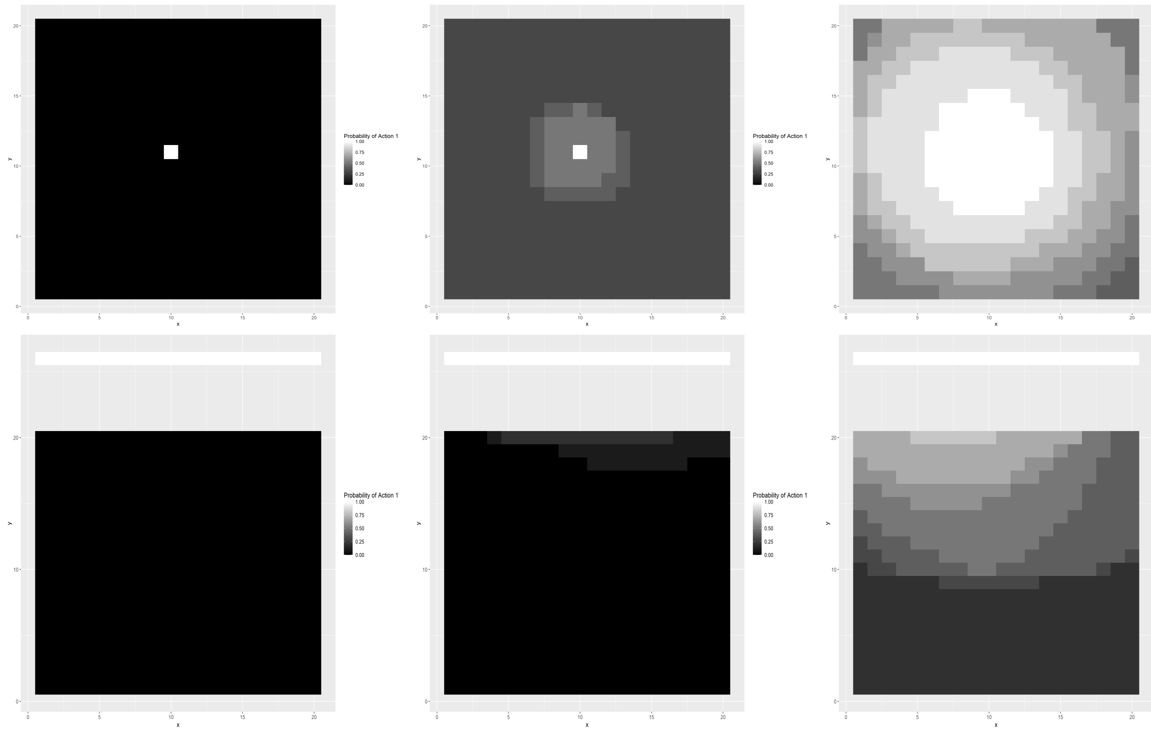

Figure 2: Transitions inspired by local leaders (top) vs. transitions inspired by a neighboring community (bottom). Black dots represent choice of action 0. White dots represent action 1. Darker shades represent a larger fraction of simulations that result in action 0 being chosen. Left plots are initial conditions, middle are $t = 10$, right figures represent actions at $t = 20$. All are plotted in characteristic space.

Simulation details. To contrast the effectiveness of local leaders, we start from a simulation. As in our baseline setting, we assume that action 1 yields a higher average payoff than action 0, $z_{0t} - c_0 < z_{1t} - c_1$, but that a community of agents initially starts with a lower expected value of action 1, $E[z_0 - c_0|\mathcal{I}_{i0}] >> E[z_1 - c_1|\mathcal{I}_{i0}]$. (For the exact parameter choices see Appendix B.1.) This situation could arise because historically, in communities where action 1 generated a chance negative outcome early on, members chose not to take that action. In the COVID example, communities scarred by past negative experiences might not trust the government and medical institutions and therefore refuse a medical interventation (vaccination) advertised by those institutions (cf. famously the Tuskegee example). In such a community, all agents are rationally stuck with low beliefs about this more desirable action.[13]

Building on these initial assumptions, we consider two scenarios that aim to contrast the role of a community leader with highly visible (but not social close) information provision.

In the first set of simulations, we consider the effect of a single centrally-located local leader who takes the higher-payoff action. We use the standard simulation procedure defined above, except that, as an initial condition, we give one local leader a subsidy to change actions. The simulation

[13]Similar results arise if agents were uncertain about the payoffs from action 1 and risk averse.

results are shown in the top row of Figure 2, separately for $t = 1$, $t = 10$, and $t = 20$. We see the local leader's choice of action 1 as the white dot in the first plot, located centrally in the characteristic space.

In the second set of simulations, shown in the second row of plots, we remove the subsidized local leader but add highly visible, though socially more distant information about action 1. Specifically, we add a group of agents on top of the 20x20 grid, with an additional set of five empty rows in between. These agents are socially more distant. However, given that there are twenty of these neighboring community nodes, they provide a large number of informative signals. In this simulation, everyone is equally visible ($\psi_{ij} = 0.2,\ \forall i, j$). Like the local leader in the first simulation, they choose action 1, because they get a subsidy to do so. Everyone else in the community, including the local leader from the previous example, starts with the standard beliefs and no subsidies. The simulation procedure is the standard one. As a result, all members of the core community, choose action 0, seen as the uniform black box at $t = 0$ on the bottom left panel.

Comparing scenario 1 and scenario 2, notice that the neighboring agents in the second set of simulations offer substantially more access to information about action 1 than the local leader does: They produce 20 potential signals about the superior payoffs of action 1, in every period. However, as the middle panels (from $t = 10$) and the rightmost panels (from $t = 20$) reveal, this information does not disseminate as fast as coming from the neighboring agents as that coming from the local leader. The reason is, of course, that information coming from sources with distant characteristics do not resonate with the community. The community members put less weight on such information when deciding about their own action, and thus respond fairly slowly.

In contrast, the single source of resonant information switches the most nodes from black (action 0 for sure) to light shades (higher probability of action 1). Because the leader's characteristic similarity leads others to put more weight on their experience, getting the one local leader to switch actions brings about faster learning and a faster change in actions from the community. This illustrates the power of resonant information and contrasts it with access to abundant information that does not resonate with a given community.

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

At date $t = 1$, either

1. A local leader, located in characteristic space at $(\lfloor n/2 \rceil, \lfloor n/2 \rceil)$ takes action 1, where $\lfloor \cdot \rceil$

represents rounding to the nearest integer; or

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

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

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

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

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

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

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

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

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

Crisis as a Time of Re-invention In a significant crisis, (think Financial Crisis or Covid-19 Pandemic) such an event can shock agents out of their 'comfort zone' and induce them to experiment with changing their action and ultimately to change actions permanently. Such change can happen in response to the shock for at least two reasons: The crisis might provide access to new information; or the crisis makes information resonate more with agents as it affects other agents who are socially close. While both channels can be at work (and interact), numerical and theoretical results in the Appendix illustrate how the characteristic of the transition and the visible patterns of behavior that result from "more access" versus "more resonance" are distinct. The implications of this finding for the COVID example are that members of communities who tend to pursue low-skill work and tend to not complete highschool and attend college (e. g., who see low-skill work as a safe option and higher education as risky), might fully realize that low-skill income is riskier than they thought and might be motivated the to acquire more education.

2.4 Social Networks and the Decline of Expertise

The concept of information resonance also speaks to the role of social networks, and sheds light on a different aspect than emphasized in prior literature. Take the example of Facebook. Pre-Facebook, access to information was constrained by geography. It was hard to learn news about the community you grew up in once you moved away, and it was hard to find information from people "similar to you" if you were not part of the local majority group. Post-Facebook, however, you can form communities with people from all over. You weight that community information more, and average information weights increase. At the same time, the influence of (socially more distant) experts and celebrities declines.

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

Simulation details. The agents are distributed uniformly in geographic space, with one person assigned to each node in a grid. However, while characteristics were assigned uniformly in the pervious simulations, they are assigned normally here. The reason for this change is that we want to illustrate the differential impact of social networks on members of the majority community and those in the majority ("in the tails"), such as ethnic or racial minorities or political extremists. Specifically, we draw two characteristics v and w from the distribution $N(10, 5)$. When we plot agents in characteristics space rather than geographic space, we assign each agent is assigned to the node with the closest integer value. For example, if an agent draws $(v, w) = (2.7, 9.1)$, they are assigned to the $(3, 10)$ on the characteristic grid. In many cases, multiple agents are assigned to the same box. The number of agents assigned to a box is not visible on the plot, but it is systematically higher at the center of the grid since many agents draw characteristics closer to the mean. Around the edges of plots in characteristic space, instead, there will be many white (empty) boxes indicating cases where no agent drew characteristics in this space.

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

(a) in geographic space:

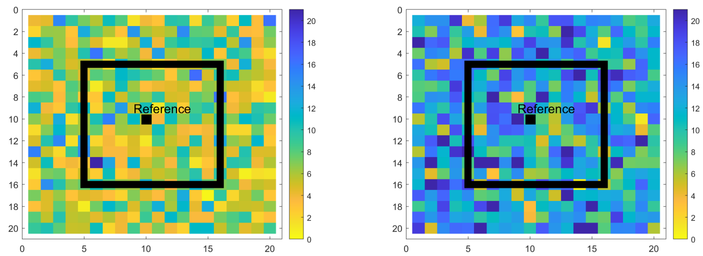

(b) in characteristic space, with geographic information constraints:

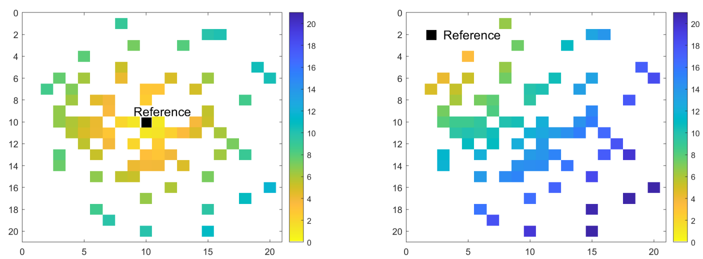

(c) in characteristic space, with no information constraint:

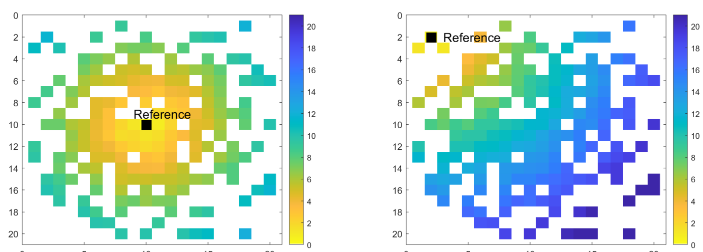

Figure 3: Social proximity in geographic and social space. Extreme agents get far more information flow when geographic constraint is relaxed. Social proximity to a reference person is represented in warmer colors (yellow tones); social distance in cooler colors (blue tones). The black frame in Panel (a) represents the geographic restriction (with $\Psi = 1$ inside the frame, and $\Psi = 0$ outside the frame) prior to social networks. The middle plots in Panel (b) include only geographically nearby agents (those inside the black frame in Panel (a), i. e., with $\Psi = 1$), plotted in characteristics space. The bottom two plots in Panel (c) include both geographically nearby and distant agents.

Figure 3 illustrates the differential effect of social networks on members of the majority and the minority groups. In Panel (a) (top row), we plot community members according to their physical (geographic) location. One agent lives at each node. Differently from prior figures, we use colors to illustrate social proximity relative to a reference person, shown as a black box at the center of the grid. (Note that the inital location geographic has no particular significance. We are just examining the 20-by-20 neighborhood around a given agent's location.) Warmer (yellow) tones indicate social proximity, and cooler (blue) tones indicate social distance, and each agent is given a color, according to their social proximity to the reference point. The variegated color pattern on the geographic topology represents the fact that social type and geographic location are not correlated in this example.

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

Focusing at first on the agents inside the black frame ("pre-Facebook"), we re-plot all such agents in their respective colors in characteristics space in Panel (b) of Figure 3. These panels show the social proximity of observable community members, i. e., those inside the black frame, both in color and spatially. They are then weighted according to how resonant their information is. The empty (white) boxes are characteristics pairs not represented by any agent to whom the reference person is connected. Note that, in the left graph of Panel (b), the reference person remains centrally located also in characteristics space. This reflects that we picked a socially central person, akin to the local leader in Section 2.3 before. In the right graph, instead, the reference person moves to the upper left corner in characteristics space. This reflects that we picked a socially more isolated, or outlier set of characteristics. Here the reference person has extreme characteristics and is thus located at the fringe of the social space.

Now let's turn to the effect of widened communities, due to social networks. In Panel (c), we allow the reference node to connect to all agents from Panel (a), whether inside our outside the black frame. The change can be interpreted as representing the ability of social media to connect across large geographic distances. Correspondingly, the bottom two panels show the social location of all nodes from Panel (a), just in different locations, representing their social distance from the reference node. Note that some locations on this plot have multiple nodes, representing multiple agents in that characteristic space location. Other nodes remain white since no agent has the corresponding characteristics. Naturally, we now see fewer white boxes. Relaxing the geographic constraint on information access allows a denser set of agents to be observed. This relaxation of the geographic constraint is meant to represent the way in which social media allows for communication and community formation, without regard to geographic location.

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

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

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

Figure 3 illustrates the logic of this result. This scarce resonant information is the large number of white boxes and few warm-colored boxes around the reference agent in the center-right pane of Figure 3. Giving extremist agents like the reference agent on the right access to information from geographically distant but socially proximate agents greatly increases their resonant-information set (bottom-right panel). For the same reason, it also greatly decreases the weight they place on information from authorities.

The difference between the left and right columns of panels in Figure 3 is that the left column agent is socially central. This is someone in the majority. Since social characteristics are normally distributed, someone near the average has a high density of socially proximate neighbors. In contrast, on the right of Figure 3, the reference agent has outlier social characteristics. Because they are located in the tail of the normal distribution, this agent has fewer socially proximate neighbors to identify with and learn from. When social media relaxes the geographic constraint on information access and allows this agent to observe all nodes, some resonant information sources (yellow boxes) appear. This person with extreme characteristics will now place more weight on those resonant sources of information and less weight on all other information, including reliable and informed authorities.

3 Using Theory for Resonance Measurement

Measuring resonance is useful because knowing who is persuasive can help craft effective policy. For example, having the right people convey public health measure might induce better compliance. While this idea is obvious, identifying the relevant dimensions of resonance may not be. Suppose one observes a collection of actions of agents with heterogeneous characteristics. If an economist studies these actions, but sorts them on some irrelevant dimensions, the result may suggest to the economist that actions are independent. There is no social diffusion of behavior. Yet, sorted on the correct dimensions, the same choices exhibit a clear diffusion process where social learning is obvious. Figure 4 illustrates this possibility. In each of the three figures, a location represents a person in a characteristic space. The dimensions could represent geography, skin color, partisanship, etc. The color associated with that location represents the fraction of times that a person with these characteristics chose action 1. White squares represent a consistent choice of action 1. Darker squares indicate a higher prevalence of action 0, for an agent with these characteristics. In the left panel, the colors have no discernable pattern. Actions appear random, with no clear pattern. On the right, the same actions are sorted on the relevant dimensions of resonance. A clear pattern of spatial correlation appears. This is not necessarily spatial, in the sense of geographic correlation. But this similarity of nearby colors represents the idea that people with similar characteristics, along these dimensions, tend to make similar action choices. In between, the middle panel illustrates the data sorted on dimensions that are nearly optimal, but not quite. The diffusion patterns emerges, but is not in clear focus. Our next set of results shows how to use our model to extract this clear focus from observed data.

Diffusion at $t = 10$:

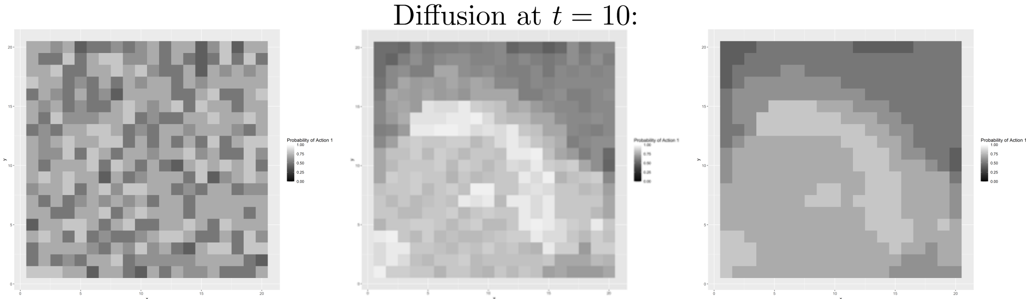

Figure 4: Diffusion pattern at $t = 10$ when sorting on completely wrong, partially wrong and correct characteristics.

To extract resonance from observed data, we need to use the structure of our model. To make the key connection between model and data, it is helpful to compare the outcomes predicted by our model to outcomes predicted by an alternative model, with no resonance effect. In other words, the alternative model has $\omega(i, j) = 1$, $\forall\ i, j$. This is what we mean by a model with "Bayesian beliefs."

Lemma 1. *There exists a covariance of payoffs Σ^R that rationalizes the actions of agents, with Bayesian beliefs.*

This result is both discouraging and helpful. The discouraging part is that, without additional knowledge, it is impossible to look at a set of actions and distinguish information relevance and information resonance. Any outcome pattern generated by an economy filled with agents who respond to information resonance could also be generated by an economy with agents who are pure Bayesian, who do not respond to resonance.

However, the helpful part of this result is that we can estimate this Σ^R. If an econometrician can also measure the true covariance of payoffs Σ_z directly, then comparing the measured covariance with this Σ^R can reveal resonance. To estimate relationships between agents with various characteristics, one needs to group agents into set with similar social characteristics that are potentially resonance-relevant. At this point, we switch from i denoting a specific individual to i' denoting a group of individuals with similar characteristics.

How one approaches measurement depends on what kind of data is available. We consider two common alternatives: survey or forecast data that reveal beliefs, or outcome data that reveal agents' actions. We begin with survey data. Suppose that an econometrician can get a set of agents to report their expected payoff. We can express this expected payoff as an expectation, conditional on everything that agent knows, $E[z_i|\mathcal{I}_{it}]$.

The econometrician should then regress each forecast of an agent with characteristics similar to i on neighboring outcomes $(z + \epsilon)$:[14]

$$E[z_{i'}|\mathcal{I}_{it}] = \alpha + \beta_i^R (z + \epsilon) + \eta_{i't}.$$

The coefficient vector β^R captures the sensitivity of one's expectations to the actions of others, from different groups.

Similarly, for binary actions, an econometrician could estimate coefficients of a logit model, using the payoffs of others with various characteristics to predict the probability that an agent of type i' will take action 1:

$$Pr[a_{i'} = 1|\mathcal{I}_{it}] = \Phi\left(\alpha + \beta_i^R (z + \epsilon) + \eta_{i't}\right).$$

If one has measured the true payoff covariance Σ_z, then simply map the sensitivity estimate β_i^R into a (scaled) resonance weight as follows:

$$w_i \bar{w}_i = \beta_i^R./\beta_i = (I + \sigma_\epsilon^2 \Sigma_z^{-1})\beta_i^R,$$

where ./ denotes an element-by-element division of two vectors. In other words, resonance is the ratio between what the estimated sensitivity is to others' outcomes and what that sensitivity should be, based on knowledge of the covariance of those outcomes.

The key here, of course, is that one needs to know the true payoff covariance. In some cases, that will not be possible to know. But in other cases, it will be feasible, or even trivial. Consider the case of vaccines. The efficacy vaccines for different age, race and socio-economic groups has been carefully studied. It is also quite obvious that one's political orientation should not affect vaccine outcomes directly. Only, perhaps indirectly, through potentially correlated variables such as health, age, gender or weight. Another case where covariances are know is where they are all identical. For example, in the experiment that asked participants to forecast national inflation, after getting information from federal reserve governors of different genders or races, the payoff of a correct inflation forecast does not depend on any personal characteristic, by construction. In

[14] Of course, this is the procedure that is consistent with our static model. With dynamic panel data, an econometrician would want to regress conditional expectations on current and past realizations of others. Depending on the type of outcome and time horizon, some time discounting might be advisable.

that case, every entry of Σ_z should be identical. Any variation in the sensitivity of forecasts to the characteristics of the sender is due to resonance.

4 Empirical Evidence

4.1 Measurement, Identification and Reflection

To be adopted, the idea of information resonance will need to be measured, calibrated, and evaluated with empirical evidence. In doing this, many questions about identification arise that the model can help resolve. One challenge of quantifying local learning 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, as in models with social norms, or due to their beliefs.

Another potential issue is Manski (1993)'s reflection problem: It is difficult to distinguish how an individual affects a group from the group's effect on the individual. The reflection problem is predominantly an issue with linear-in-means models; however, Brock and Durlauf (2001) highlights how in a backward-looking, dynamic linear-in-means model where current behaviors are driven by past beliefs the reflection problem does not apply. Given that our model relies centrally on past experiences, it is unlikely that the reflection problem would be exhibited.

We build ways to address these measurement problems into the model. After estimating the model, one can use it to perform counter-factual experiments, where one person changes actions or some signal noise is high or low. Demonstrating that localized learning is the dominant effect occurring econometrically allows for it to be incorporated into policy. A person does not change their action in isolation from their community. Our model implies that there is social multiplier since individuals will take up optimal actions demonstrated by their peers and, uniquely, highlights the importance of targeting community leaders across communities to maximize the social multiplier with minimal initial change.

4.2 Evidence from Occupational Choice

As a first piece of evidence, we want to test the key assumption that people place greater weight on information learned from those within their own communities, i. e., from people they feel socially close to. Here, we use geographic proximity in combination with ethnicity as a measure of social proximity. We test the hypothesis that young people will base their occupational choices on the observed occupation choices of older members of their local ethnic communities. We expect that a young person's choice to enter an occupation will be positively correlated with the extent to which people of their same ethnic group are overrepresented in that occupation in their PUMA. Next, we further test our assumption by examining how a "shock" to members of one's community impacts the behavior of the group. Lastly, we test the prediction that the effect on beliefs will be stronger when matching more demographic characteristics. We can reasonably assume that men and women of the same ethnic communities have overlapping social networks such that they are aware of the occupations across genders at similar rates. We include an additional measure of social proximity using gender and hypothesize that young members of an ethnic community will enter occupations in which elders of the same ethnicity *and* gender are overrepresented. This analysis helps disentangles the effects of information frictions.

4.2.1 Data and Variables

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 (since our goal is to measure what occupation a person chooses), 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,[15] 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.

Constructed Variables Using a combination of these variables, we measure the fraction of each ethnic group in each occupation in each PUMA each year (denoted as *Ethnic Ratio*). We do this by taking the number of people of primary working age (which we defined as 26 to 65) in each of these occupation/ethnic/PUMA/year groups and dividing it by the number of people of the ethnic

[15] 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.

group in the PUMA in the same year and weight the estimates appropriately using the weights provided by IPUMS.[16] Similarly, we then consider the fraction of all people in the PUMA in each occupation each year (denoted as *Local Ratio*).

Our main variable of interest is ethnic overrepresentation (denoted as *Ethnic ORep*) and is the difference between *Ethnic Ratio* and *Local Ratio* and it measures the extent to which the ethnic group is overrepresented in an occupation in each PUMA each year.

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

$$\textit{Ethnic ORep}_{i,j,k,t} = \textit{Ethnic Ratio}_{i,j,k,t} - \textit{Local Ratio}_{j,k,t}.$$

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

$$\textit{Youth Rate}_{i,j,k,t} = \frac{\textit{Occupation Size}_{i,j,k,t}}{\textit{Population Size}_{i,k,t}}.$$

In addition to ethnicity, we consider gender as another measure of social proximity. In order to do so, we re-estimate the variables *Ethnic Ratio* and ethnic overrepresentation to be gender-specific. Here, the variable *Ethnic Ratio (Same Gender)* measures the fraction of people aged 26-65 of a given ethnicity and gender in an occupation and PUMA during a certain year. Subtracting the number of fraction of people in the PUMA in the occupation, *Local Ratio*, creates the variable *Ethnic ORep (Same Gender)*. This variable measures overrepresentation of an ethnic-gender group. For example, how many more Vietnamese women are nail technicians compared to all nail technicians in New York.

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

[16] 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.

2010. We 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.

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

Summary Statistics To construct the baseline estimates of the *Ethnic Ratio* and *Local Ratio* variables, we use the ACS data from 2005 to 2020 including all employed people ages 26 to 65 for which we had occupation and ethnicity information. This data totals 16,060,861 observations. In order to ensure reliable estimates, we required that at least 100 observations of an ethnic group be present in a PUMA in order to include it in the data. With this restriction, one ethnic group, Australians/New Zealanders, drops out of the dataset and we are left with 16 unique ethnicities. We observe a total 1,054 PUMAs in the data. We observe all 32 occupations for each ethnic group.[17] Table 1 shows the average number of observations of each ethnicity in each occupation/PUMA/year group. Table 2 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 3 and 4 provide summary statistics of the constructed variables. Table 3 contains values for the variables dependent on ethnicity only, while Table 4 provides the gender-specific values.

4.2.2 Empirical Results

We want to test the assumption that people place greater weight on information learned from those within their own communities and social networks. The ideal experiment would be to vary the success of similar individuals in different occupations and then measure how this affects the individuals beliefs about their ability to succeed in different occupations. In practice, we can only rely on natural variation in individuals success and only observe individuals occupational choices, not beliefs of success. Based on observable characteristics, we use geographic proximity in combination with ethnicity as a measure of social proximity. We test the assumption that young people will base their occupational choices on their observations of the occupation choices of older members within their local, ethnic communities. We expect that a young person's choice to enter an occupation will be positively correlated with the extent to which people of their same ethnic group are overrepresented in that occupation in their PUMA as that provides information about there ability to succeed in that occupation.

To test this assumption, we run a linear regression of the youths rate of choosing an occupation within a local-ethnic group, the proportion of 18-to-22-year-olds in each occupation in each PUMA-

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

ethnicity pair each year, on ethnic overrepresentation, the percent amount by which the ethnic group is overrepresented in the occupation in the PUMA that year. Our baseline specification is

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

where $\gamma_{j,k}$ is an occupation-PUMA fixed effect and ε is the error term. β is the coefficient of interest on ethnic overrepresentation and has the interpretation that in a given year for a one percent increase in ethnic overrepresentation, youth choice of that occupation in that puma-ethnicity pair will increase by β percent. Note that we cannot include year fixed effects because then *Ethnic ORep* would be colinear. Note that by construction that within a occupation-PUMA-year group ethnic overrepresentation accross ethnicities must sum to 0. Likewise, conditional on an ethnicity-puma-year group the sum of the youth rates across occupations must sum to 1.

Table 5 displays the result. As expected, we see a positive and significant coefficient on ethnic over representation in Column 5, our baseline specification. Our results are robust to occupation multiple specifications of fixed effects, including most stringently occupation-PUMA-year or occupation-PUMA-ethnicity fixed effects. These results suggest that the more overrepresented an ethnicity is in an occupation, the more young people of the same ethnicity choose to enter the occupation. For instance, based on our baseline result, if in a given PUMA during a given year 1.2 percent more people of an ethnic group are employed in a given occupation than all people employed in that occupation (which is the innerquartile range of ethnic overrepresentation), we expect the percentage of young people of that ethnicity entering the given occupation to be 0.15 percent higher. Here, the interacted Occupation-PUMA fixed effects account for differences in the general popularity/accessibility of certain occupations based on location (some occupations are more common in certain PUMAs and will have more young entrants by virtue of their location). Column 6 includes the triple interaction Occupation-PUMA-Year fixed effects, controlling for the popularity of the occupation in a given location across years. This inclusion, therefore controls for job trends over the years.

We also predict that the effects will be stronger when matching multiple demographic characteristics. We include an additional measure of social proximity using gender and hypothesize that young members of an ethnic community will enter occupations in which elders of the same ethnicity *and* gender are overrepresented. This analysis helps to disentangle the effects of information frictions as it is plausible that, within the same ethnic communities, the men and women of the same ethnic communities run in the same social networks and discuss their occupations with each other at similar rates. However, if people tend to follow elders of the same gender, this implies that outcomes of people from the same social groups resonate more deeply and drive behavior.

In order to add this additional dimension to the analysis, we re-estimate the variables *Ethnic Ratio* and *Ethnic ORep* to be gender-specific. 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, thus, providing evidence that information access is not the only mechanism at play in determining occupation choices.

Table 6 displays the effect of gender-specific ethnic overrepresentation. Table 6 uses the $\textit{Ethnic ORep (Same Gender)}_{i,l,j,k,t}$ variable, which describes the extent to which a certain ethnic-gender group was overrepresented in an occupation. We see that in all specifications, same-gender overrepresentation has a significant and positive effect on the fraction of young people entering an occupation. However, the opposite gender's overrepresentation effect is significant and negative, suggesting that information access is not the primary mechanism driving the results. People seem to be discouraged from following the path of opposite gender elders of their same ethnicity, even when they learn about similar outcomes of both genders. Based on Column (5) of Table 6, if a given ethnic group of a given gender in a given PUMA and year is employed in a given occupation at 0.8 percent more than the rate of all people in the given PUMA and year, we expect the fraction of young people of the given gender and ethnicity in the given occupation to increase by 0.35 percent in the given year. However, if members of the same ethnicity, but opposite gender are overrepresented by the same amount, we expect the percent of young people to decrease by 0.26.

Lastly, we test the hypothesis that a "shock" to members of one's community impacts the behavior of the group. To test this hypothesis, we consider the effect of the 2008 financial crisis on young peoples' entrance into occupations that experienced severe layoffs. Using a difference-in-differences model, we compare the effect of ethnic overrepresentation on young peoples' entrance into an occupation 2 years prior to and 2 years after the financial crisis.

For this analysis, we set the ethnic overrepresentation variable to its average in 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. We include an indicator variable indicating that that the observation is from the year 2010 or later (i.e., post-recession). Here, we use data from 2005 and 06 as our "pre-recession" years, and data from 2010 and 11 to study the "post-recession" effects.

Table 7 presents the results. We find a positive coefficient on the interaction of the high layoffs dummy and pre-recession ethnic overrepresentation, we see that prior to the recession, young people were following their coethnic elders into occupations where they were overrepresented. In particular, based on Column (5) of Table 7, we see that before the recession if ethnic overrepresentation in a job that would later experience high layoffs during the recession is 1.8 percent, meaning that that the ethnic group has 1.8 percent more people in the occupation relative to all people in the PUMA, we would expect the fraction of young people in the occupation to increase by 0.60 percent.[18] After the

[18]We focus on Column (5) as the primary measure, since, here, we have included Occupation-PUMA fixed effects, which are our baseline fixed effects and control for the overall popularity of an occupation in a PUMA. However,

recession, the effect is reversed. Young people are deterred from occupations in which their ethnic community experienced large job losses. Using the same example of a job with overrepresentation of 1.8 percent in 2005 and 06, after the recession, we expect the fraction of young people entering this occupation to decrease by .15 percent, one-fourth of the effect.

Using the fixed values, we are also able to run the regression separately for ethnic groups that were overrepresented pre-recession and ethnic groups that were underrepresented pre-recession separately. Our model suggests that having witnessed relatively more people in overrepresented professions be affected negatively by the recession (and thus have bad outcomes), you will be less likely to join those professions. In Table 8 and Table 9, we replicate Table 7 for the overrepresented and underrepresented occupation-PUMA-ethnicity groups. We find in the respective Column (5)s that pre-recession ethnic overrepresentation, which measures representation relative to the average, continues to increase youth occupation choice positively for underrepresented professions, but for overrepresented professions actually has a negative effect on the post-recession likelihood for youths to choose that job.

For robustness, we also expanded the post-crisis period to include all of 2010-2020, with very similar results (see Appendix-Table A.1). Additionally, we split this table between the positive and negative values of the ethnic overrepresentation variable (see Appendix-Tables A.2, A.3). Appendix-Table A.4 shows that we find similar results to Table 7when expanding the crisis shock analysis to be gender-specific.

5 Conclusion

This paper explores the difference between information access and information resonance. We note that in the modern world, with abundant data at our fingertips, most people do not lack access to information. Instead, they need information that speaks to them in a meaningful way. Neuroscience research suggests that information affects choices when the observer of that information connects to that information, in an emotional way. Emotional connections with anothers' experience is stronger when the observer and the information provider have similar characteristics.

By formally modeling this phenomenon, we learn about policies that could be more effective in inducing behavioral change. Information and personal experiences need to be conveyed by people who resonate with their communities. One well-placed local leader may be as effective as many, many outsiders and expensive advertising. The experiences of that central, trusted figure will inspire confidence in the information that a shift is actions will ultimately be beneficial. In contrast, redistribution only temporarily offsets the lower payoff to the community that chooses the suboptimal action. It does not solve the underlying problem that this community is stuck in a

we also see similar effects with the inclusion of Occupation-PUMA-Year fixed effects (Column (6)), as well as the exclusion of all fixed effects (Column (1)) and other fixed effects specifications.

bad equilibrium.

The process of information diffusion governs everything from technological adoption, to political attitudes, to financial trading strategies. We raise the possibility that this diffusion happens not only in geographic space, but in social characteristic space. This change in focus, from information access to information resonance, could change the way we organize social science data for analysis and the way we enact policy to improve our collective well-being.

A unique contribution of these insights and findings is the prescription that policy makers should target community leaders (in distinct communities) to maximize impact with minimal cost.

References

W. C. Abraham, O. D. Jones, and D. L. Glanzman. Is plasticity of synapses the mechanism of long-term memory storage? *NPJ science of learning*, 4(1):1–10, 2019.

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. Alsan, F. C. Stanford, A. Banerjee, E. Breza, A. G. Chandrasekhar, S. Eichmeyer, P. Goldsmith-Pinkham, L. Ogbu-Nwobodo, B. A. Olken, C. Torres, et al. Comparison of knowledge and information-seeking behavior after general covid-19 public health messages and messages tailored for black and latinx communities: a randomized controlled trial. *Annals of internal medicine*, 174(4):484–492, 2021.

K. J. Arrow. Classificatory notes on the production and transmission of technological knowledge. *American Economic Review Papers and Proceedings*, 59(2):29–35, 1969.

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

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

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

A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.

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.

M. F. Bear. A synaptic basis for memory storage in the cerebral cortex. *Proceedings of the National Academy of Sciences*, 93(24):13453–13459, 1996.

D. Bergemann and J. Välimäki. Experimentation in Markets. *The Review of Economic Studies*, 67(2):213–234, 2000. ISSN 0034-6527.

T. V. Bliss and T. Lømo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of physiology*, 232(2):331–356, 1973.

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.

J. Cai, A. De Janvry, and E. Sadoulet. Social networks and the decision to insure. *American Economic Journal: Applied Economics*, 7(2):81–108, 2015.

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.

D. Comin, M. Dmitriev, and E. Rossi-Hansberg. The spatial diffusion of technology. 2013.

T. G. Conley and C. R. Udry. Learning about a new technology: Pineapple in ghana. *American Economic Review*, 100(1):35–69, 2010.

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

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

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

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

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

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

A. D. Foster and M. R. Rosenzweig. Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6):1176–1209, 1995.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

R. Malinow. Transmission between pairs of hippocampal slice neurons: quantal levels, oscillations, and ltp. *Science*, 252(5006):722–724, 1991.

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.

B. L. McNaughton, R. Douglas, and G. V. Goddard. Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. *Brain research*, 157(2):277–293, 1978.

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.

R. G. M. Morris, S. Davis, and S. Butcher. Hippocampal synaptic plasticity and nmda receptors: a role in information storage? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 329(1253):187–204, 1990.

K. Munshi. Social learning in a heterogeneous population: Technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1):185–213, 2004.

T. Otto, H. Eichenbaum, C. G. Wible, and S. I. Wiener. Learning-related patterns of ca1 spike trains parallel stimulation parameters optimal for inducing hippocampal long-term potentiation. *Hippocampus*, 1(2):181–192, 1991.

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. Torres, L. Ogbu-Nwobodo, M. Alsan, F. C. Stanford, A. Banerjee, E. Breza, A. G. Chandrasekhar, S. Eichmeyer, M. Karnani, T. Loisel, et al. Effect of physician-delivered covid-19 public health messages and messages acknowledging racial inequity on black and white adults' knowledge, beliefs, and practices related to covid-19: a randomized clinical trial. *JAMA Network Open*, 4(7):e2117115–e2117115, 2021.

Table 1: Number of Observations of Each Ethnic Group ($Freq_{i,j,k,t}$ (unweighted))

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

Table 2: Number of People in Each Ethnic Group ($Freq_{i,j,k,t}$ (weighted))

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

Table 3: Summary of Constructed Variables (Ethnicity)

Statistic	N	Mean	Standard Deviation	25th Percentile	Median	75th Percentile
Ethnic Ratio	1,203,328	0.031	0.046	0.003	0.014	0.039
Local Ratio	1,203,328	0.031	0.042	0.006	0.015	0.039
Ethnic $\text{ORep}_{i,j,k,t}$	1,203,328	−0.000	0.019	−0.007	−0.001	0.005
Ethnic $\text{ORep}_{i,j,k,(t-1)}$	1,122,624	−0.000	0.019	−0.007	−0.001	0.005
Ethnic ORep_{2005}	1,203,328	0.000	0.017	−0.005	−0.0004	0.004
Percent Layoffs	1,035,800	−0.172	1.176	−0.277	−0.004	0.220
Job Loss	1,098,528	−63.336	1,779.922	−489.000	−3.000	403.000

Table 4: Summary of Constructed Variables (Ethnicity, Gender)

Statistic	N	Mean	Standard Deviation	25th Percentile	Median	75th Percentile
Ethnic Ratio	2,108,064	0.031	0.054	0.000	0.009	0.037
Local Ratio	2,108,064	0.031	0.042	0.006	0.015	0.039
Ethnic $\text{ORep}_{i,l,j,k,t}$	2,108,064	0.000	0.034	−0.011	−0.002	0.007
Ethnic $\text{ORep}_{i,l,j,k,(t-1)}$	1,964,544	0.000	0.034	−0.011	−0.002	0.007
Ethnic ORep_{2005}	2,108,064	0.000	0.032	−0.009	−0.001	0.006
Percent Layoffs	1,833,188	−0.170	1.149	−0.274	−0.004	0.218
Job Loss	1,941,440	−64.108	1,803.196	−494.000	−3.000	407.000

Table 5: Occupation Choice: Ethnicity

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ethnic ORep	0.0786***	0.126***	0.126***	0.131***	0.123***	0.131***	0.0385***
	(0.0201)	(0.0101)	(0.0101)	(0.0109)	(0.00934)	(0.0107)	(0.0126)
Constant	0.0312***	0.0214***	0.0312***	0.0312***	0.0312***	0.0312***	0.0312***
	(0.000243)	(0.000520)	(0.000105)	(0.0000995)	(0.0000765)	(0.0000971)	(4.70e-14)
Occupation FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	No	Yes	No
Ethnicity FE	No	No	No	Yes	No	No	Yes
PUMA FE	No	No	No	No	Yes	Yes	Yes
Occupation*Year FE	No	No	Yes	No	No	No	No
Occupation*Ethnicity FE	No	No	No	Yes	No	No	No
Occupation*PUMA FE	No	No	No	No	Yes	No	No
Occupation*PUMA*Year FE	No	No	No	No	No	Yes	No
Occupation*PUMA*Ethnicity FE	No	No	No	No	No	No	Yes
Observations	1203328	1203328	1203328	1203328	1202976	1050624	1196736

Clustered standard errors in parentheses.

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

Table 6: Occupation Choice: Ethnicity, Gender

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethnic ORep (Same Gender)	0.474***	0.429***	0.429***	0.442***	0.432***	0.441***	0.393***	0.773***
	(0.0107)	(0.00572)	(0.00573)	(0.00606)	(0.00540)	(0.00594)	(0.00668)	(0.00686)
Ethnic ORep (Wrong Gender)	-0.242***	-0.326***	-0.326***	-0.311***	-0.321***	-0.319***	-0.361***	
	(0.0108)	(0.00557)	(0.00557)	(0.00582)	(0.00526)	(0.00584)	(0.00647)	
Constant	0.0312***	0.0208***	0.0312***	0.0312***	0.0312***	0.0312***	0.0312***	0.0312***
	(0.000252)	(0.000453)	(0.0000952)	(0.0000897)	(0.0000677)	(0.0000732)	(3.78e-19)	(2.90e-19)
Occupation FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	No	Yes	No	Yes
Ethnicity FE	No	No	No	Yes	No	No	Yes	Yes
PUMA FE	No	No	No	No	Yes	Yes	Yes	Yes
Occupation*Year FE	No	No	Yes	No	No	No	No	No
Occupation*Ethnicity FE	No	No	No	Yes	No	No	No	No
Occupation*PUMA FE	No	No	No	No	Yes	No	No	No
Occupation*PUMA*Year FE	No	No	No	No	No	Yes	No	No
Occupation*PUMA*Ethnicity FE	No	No	No	No	No	No	Yes	No
Occupation*PUMA*Ethnicity*Year FE	No	Yes	No	No	No	No	No	Yes
N	1964544	1964544	1964544	1964544	1964384	1916192	1961056	1683840

Clustered standard errors in parentheses.

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

Table 7: Impact of a Shock on Occupation Choice (2005/06 and 2010/11 with Ethnic $ORep_{2005}$)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ethnic $ORep_{2005}$	0.0351	0.0619*	0.0424	0.162***	0.0339	0.0413	
	(0.0506)	(0.0354)	(0.0354)	(0.0390)	(0.0356)	(0.0387)	
POST	0.000768	0.000755		0.000725	0.000706		0.000719
	(0.000479)	(0.000472)		(0.000472)	(0.000495)		(0.000472)
(POST)*(Ethnic $ORep_{2005}$)	0.0733	0.0342	0.0698	0.0341	0.0237	0.0316	0.0268
	(0.0519)	(0.0503)	(0.0502)	(0.0501)	(0.0522)	(0.0555)	(0.0503)
High Layoffs	-0.0174***	0.00157***	0.000105	0.00139***			
	(0.000681)	(0.000525)	(0.000522)	(0.000522)			
(High Layoffs) *(Ethnic $ORep_{2005}$)	0.156*	0.209***	0.230***	0.145*	0.299***	0.296***	
	(0.0939)	(0.0750)	(0.0746)	(0.0752)	(0.0754)	(0.0804)	
(POST)*(High Layoffs)	-0.00267***	-0.00249***	0.000466	-0.00239***	-0.00235***		-0.00240***
	(0.000760)	(0.000752)	(0.000747)	(0.000752)	(0.000790)		(0.000757)
(POST)*(High Layoffs)*(Ethnic $ORep_{2005}$)	-0.272***	-0.270***	-0.309***	-0.275***	-0.273**	-0.296***	-0.289***
	(0.105)	(0.104)	(0.103)	(0.104)	(0.109)	(0.115)	(0.105)
Constant	0.0384***	0.0192***	0.0331***	0.0327***	0.0332***	0.0328***	0.0331***
	(0.000461)	(0.000991)	(0.000220)	(0.000316)	(0.000245)	(0.000193)	(0.000186)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	275427	275427	275427	275427	274986	242493	273211

Clustered standard errors in parentheses.

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

Table 8: Impact of a Shock on Occupation Choice among Over-represented Ethnic Groups (Ethnic $\text{ORep}_{2005} > 0$)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ethnic ORep_{2005}	2.010***	0.172***	0.234***	0.209***	0.0179	0.0823	
	(0.0883)	(0.0552)	(0.0580)	(0.0587)	(0.0768)	(0.104)	
POST	-0.000268	-0.000790		-0.000766	-0.000925		-0.000750
	(0.000923)	(0.000795)		(0.000792)	(0.000850)		(0.000786)
(POST)*(Ethnic ORep_{2005})	0.143*	0.167**	0.0450	0.162**	0.170**	0.0453	0.158**
	(0.0858)	(0.0724)	(0.0825)	(0.0720)	(0.0773)	(0.143)	(0.0722)
High Layoffs	-0.00683***	0.000267	0.000313	-0.000302			
	(0.00152)	(0.00116)	(0.00117)	(0.00116)			
High Layoffs*(Ethnic ORep_{2005})	-0.468***	0.265**	0.143	0.301**	0.286	0.129	
	(0.172)	(0.134)	(0.134)	(0.134)	(0.200)	(0.260)	
(POST)*(High Layoffs)	-0.000126	0.000251	0.000168	0.000205	0.000550		0.000410
	(0.00170)	(0.00160)	(0.00162)	(0.00160)	(0.00176)		(0.00160)
(POST)*(High Layoffs)*(Ethnic ORep_{2005})	-0.532***	-0.516***	-0.273	-0.503***	-0.532***	-0.326	-0.530***
	(0.192)	(0.182)	(0.182)	(0.182)	(0.201)	(0.344)	(0.184)
Constant	0.0171***	0.0150***	0.0319***	0.0320***	0.0339***	0.0328***	0.0348***
	(0.000947)	(0.00159)	(0.000530)	(0.000652)	(0.000763)	(0.000752)	(0.000270)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	125389	125389	125389	125387	125049	78745	124486

Clustered standard errors in parentheses.

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

Table 9: Impact of a Shock on Occupation Choice among Under-represented Ethnic Groups (Ethnic $ORep_{2005} <= 0$)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ethnic $ORep_{2005}$	-2.799***	-0.0779	-0.246***	0.0280	-0.0134	-0.0241	
	(0.0936)	(0.0820)	(0.0858)	(0.0853)	(0.117)	(0.150)	
POST	-0.000716	-0.000921		-0.000843	-0.00111		-0.000932
	(0.000884)	(0.000884)		(0.000880)	(0.000937)		(0.000871)
(POST)*(Ethnic $ORep_{2005}$)	-0.0687	-0.150	0.179	-0.141	-0.183	-0.125	-0.156
	(0.109)	(0.109)	(0.126)	(0.109)	(0.116)	(0.237)	(0.109)
High Layoffs	-0.00600***	0.000758	0.000956	-0.0000257			
	(0.00118)	(0.000919)	(0.000915)	(0.000928)			
(High Layoffs)*(Ethnic $ORep_{2005}$)	0.976***	0.0502	0.287**	-0.0258	0.677***	0.949***	
	(0.167)	(0.133)	(0.134)	(0.135)	(0.198)	(0.270)	
(POST)*(High Layoffs)	-0.0000994	-0.000151	-0.000560	-0.000257	-0.000240		-0.000370
	(0.00131)	(0.00131)	(0.00130)	(0.00131)	(0.00139)		(0.00130)
(POST)*(High Layoffs)*(Ethnic $ORep_{2005}$)	0.0125	0.0327	-0.447**	0.00194	0.0116	-0.487	-0.00819
	(0.189)	(0.188)	(0.193)	(0.188)	(0.202)	(0.411)	(0.190)
Constant	0.0119***	0.0160***	0.0305***	0.0319***	0.0331***	0.0327***	0.0318***
	(0.000773)	(0.00210)	(0.000592)	(0.000732)	(0.000825)	(0.000907)	(0.000260)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	150038	150038	150038	150036	149647	104511	148725

Clustered standard errors in parentheses.

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

A Appendix: Proofs

A.1 Proof of Proposition 1

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Simplifying yields the expression in proposition 1.

A.2 Proof of Proposition 3

Proposition 3. When no agents in a community choose action 1, and both actions suffer a negative payoff shock, an agent is more likely to choose 1 next period.

Proof of Proposition 3:
Step 1: Action 1 is more likely to be chosen at t than it was at date $t-1$, if action 1 looks equally attractive, but action zero looks worse: $E[z_1|\mathcal{I}_{it}] = E[z_1|\mathcal{I}_{i,t-1}]$ and $E[z_0|\mathcal{I}_{it}] < E[z_0|\mathcal{I}_{i,t-1}]$.

Step 2: If no agents choose action 1, then $E[z_1|\mathcal{I}_{it}] = E[z_1|\mathcal{I}_{i,t-1}]$.

Step 3: If there is a negative payoff shock to both actions at time $t-1$, then, on average (for an average realization of the payoff shock ϵ), beliefs about the payoff to action 0 will be more pessimistic: $\int E[z_0|\mathcal{I}_{i,t}]d\Phi(\epsilon_{i,t}) < E[z_0|\mathcal{I}_{i,t-1}]$.

To form conditional expectations, we use our belief-formation rule (2):

$$E[z_0|\mathcal{I}_{it}] = Var[z_0|\mathcal{I}_{it}]\left\{\sigma_0^{-2}\mu_0 + \sum_{j=1}^{N}\sum_{t'=1}^{t-1}\tilde{\psi}_{ij}1_{a_{j,t-1}=a_0}\omega_{ij}(z_0 + \epsilon_{0i,t-1})\right\}.$$

Since we have assumed that all agents take action 0 at date $t-1$, $1_{a_{j,t-1}=a_0} = 1$ for all agents j. Since the mean of $\epsilon_{0i,t-1} = 0$, the average realization of this date t belief is

$$\int\left[E[z_0|\mathcal{I}_{it}]d\Phi(\epsilon_{i,t}) = Var[z_0|\mathcal{I}_{it}]\left\{\sigma_0^{-2}\mu_0 + \sum_{j=1}^{N}\sum_{t'=1}^{t-1}\tilde{\psi}_{ij}\omega_{ij}z_0\right\}\right].$$

This is decreasing in the payoff z_0.

A.3 Proof of Proposition 2

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

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

To form conditional expectations, we use our belief-formation rule (2). When an agent as access to the geographically-unconstrained information set $\mathcal{I}^U_{it}$, with N_U regular observations and 1 observed celebrity action, their beliefs are

$$\begin{aligned} E[z_k|\mathcal{I}^U_{it}] &= Var[z_k|\mathcal{I}^U_{it}]\Big\{\sigma_k^{-2}\mu_k + \tilde{\psi}_{i,celeb}1_{a_{celeb,t}=a_k}\omega_{i,celeb}(z_{celeb}+\epsilon_{celeb,t}) \\ &\quad + \textstyle\sum_{j=1}^{N_U}\sum_{t'=1}^{t}\tilde{\psi}_{ij}1_{a_{jt}=a_k}\omega_{ij}(z_k+\epsilon_{kit})\Big\}. \end{aligned} \tag{9}$$

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

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

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

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

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

$$Var[z_k|\mathcal{I}^C_{it}] = \left(\sigma_k^{-2} + 1_{a_{celeb,t}=a_k}\omega_{i,celeb}\sum_{j=1}^{N_C}\sum_{t'=1}^{t}1_{a_{jt}=a_k}\omega_{ij}\right)^{-1} \tag{10}$$

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

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

$$Var[z_k|\mathcal{I}^U_{it}]^{-1} = \sigma_k^{-2} + 1_{a_{celeb,t}=a_k}\omega_{i,celeb}\sum_{j=1}^{N_U}\sum_{t'=1}^{t}1_{a_{jt}=a_k}\omega_{ij} \tag{11}$$

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

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

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

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

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

A.4 Proof of Lemma 1

Does there exist a variance-covariance matrix of payoffs such that if $z \sim N(\mu_z \Sigma^R)$, then the Bayesian sensitivity of beliefs to any action realization is the same as it would be under the true distribution of beliefs, but with the resonance adjustment? by Bayes Law, the sensitivity of Bayesian beliefs to observing a set of payoffs $\psi(z+\epsilon)$, if $z \sim N(\mu_z \Sigma^R)$ and $\epsilon \sim i.i.d.\, N(0, \sigma_\epsilon^2)$ would be

$$\left(\psi_{it}\Sigma^R\psi_{it} + \sigma_\epsilon^2 I\right)^{-1}\psi_{it}\Sigma^R \nVdash_i$$

where $\nVdash_i$ is an $N \times 1$ vector of zeros with a 1 in the ith place.

According to equation (2), the sensitivity of resonant beliefs to observing a set of payoffs $\psi(z+\epsilon)$ is $\overline{\omega}_i \sum_{j=1}^{N}\sum_{t'=1}^{t}\psi_{i,j}\omega_{ij}\beta_{ij}$. Substituting in the formula for the Bayesian sensitivity under the correct distribution of payoffs, β_{ij}, we get the sensitivity vector

$$\overline{\omega}_i(\omega'_{i,:}\psi).*\left(\psi_{it}\Sigma\psi_{it} + \sigma_\epsilon^2 I\right)^{-1}\psi_{it}\Sigma^R \nVdash_i$$

where $(\omega'_{i,:}\psi_{it}).*$ denotes the element-by-element product of i's resonance with all observed informative actions of other agents j, with the argument that follows. The expression that follows is simply the Bayesian updating weight under the true payoff distribution.

The lemma is that we can equate these two weights for some variance-covariance matrix Σ^R.

$$\left(\psi_{it}\Sigma^R\psi_{it} + \sigma_\epsilon^2 I\right)^{-1}\psi_{it}\Sigma^R \nVdash_i = \overline{\omega}_i(\omega'_{i,:}\psi).*\left(\psi_{it}\Sigma\psi_{it} + \sigma_\epsilon^2 I\right)^{-1}\psi_{it}\Sigma^R \nVdash_i$$

Simple algebraic manipulation reveals that

$$\psi_{it}\Sigma^R\psi_{it} = \sigma_\epsilon^2\left(I - \bar{\omega}_i\left(\psi_{it}\Sigma\psi_{it} + \sigma_\epsilon^2 I\right)^{-1}\psi_{it}\Sigma^R\psi_{it}\mathrm{diag}(\omega_i)\right)^{-1} - \sigma_\epsilon^2 I.$$

Since this is a symmetric, positive semi-definite matrix, it is a valid variance-covariance matrix.

A.5 Crisis as a Time of Re-Invention

We begin by describing the general theoretical underpinnings of this thought experiment.

Proposition 3. *When no agent in a community chooses action 1 and both actions suffer a negative payoff shock, an agent in that community is more likely to choose 1 next period.*

The proof is in Appendix A.2. The idea is that, when the payoff to action 0 declines, agents see more negative outcomes associated with action 0. So their expected payoff $E[z_0|\mathcal{I}_{it}]$ declines. Agents get no signals about the action payoff z_1. So their expected payoff of action 1 stays the same. The relative decline in expected payoff of action 0 makes action 1 more likely to be chosen.

Simulation details. We return to our baseline assumption that action 1 yields a higher average payoff than action 0, $z_{0t} - c_0 < z_{1t} - c_1$, with $z_0 = 5$ and $z_1 = 10$, costs $c_0 = 0$ and $c_1 = 1$, and individual random shocks $\epsilon_{ait} \sim N(0, 100)$. However, a community of agents initially has a lower expected value of action 1, with prior beliefs about $z_0 \sim N(3, 10)$, and prior beliefs about $z_1 \sim N(1, 10)$, and thus $E[z_0 - c_0|\mathcal{I}_{i0}] > E[z_1 - c_1|\mathcal{I}_{i0}]$. (See Appendix A.2 for details.) At the start of the first period, unknown to the agents, we introduce a one-time, permanent payoff shock. The shock results in very low payoffs to *both* actions. We lower $z_0 - c_0$ from 5 to 1, and $z_1 - c_1$ from 9 to 5. That is the expected return to both actions diminishes by the same amount.

However, since community members do not take action 1, they do not see the low payoff of action 1; they only see the low payoff to 0. That new knowledge lowers the expected utility of 0, and prompts some to switch to action 1. There are heterogeneous shocks to payoffs (ϵ_{it}) that make some switch, but not others. Figure 1 illustrates this simulation and the transition that follows, in geographic space and characteristic space.

For example, we might have in mind a situation when the payoffs resulting from the preferred action in a community suffer a negative shock – such as the COVID pandemic or the financial crisis affecting earnings from a type of job in certain industries (food, hospitality) or depending on educational attainment (e. g., below college). The same shock might have affected the alternative action, such as jobs in other industries or in jobs that require college education. In that sense, models without information resonance would typically not generate differential updating. Here, however, it is the personal experience socially close community members that generates the switching behavior. And once some community members switch, others in the community learn from them and there is a quick transition of the whole community to action 1. They learn that 1 was actually the better action and choose it going forward.

B Simulation Details

There are three sets of simulations in the results section.

B.1 The Power of Role Models (Figure 2)

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

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

Simulation 1: Every node on the (20 x 20) grid has a type equal to their location. Note that this implies a uniform distribution of characteristics. Social distance is the Euclidean distance between nodes/agents on this characteristic grid.

There is a local leader in the center, who is given a subsidy of $\tau = 5$ to switch actions. This local leader is also more visible than other agents. Instead of $\theta = 0.2$, the local leader has $\theta = 1$. As a result of our parameter choices, everyone in the 20x20 grid starts with action 0, except for one agent in the center (local leader) who chooses action 1.

We simulate 20 periods of choices by every agent. Every agent updates beliefs each period, after observing others' payoffs. Since random payoff shocks matter greatly for the evolution path, we simulate 100 times and record the fraction of simulations that result in action 1 at that location. The plots are a colorbox plot, with only black and white for colors. Black boxes represent action 0. The simulation starts from the 20x20 grid being black, and the 1x20 squares white. We plot the fraction of simulations where action 1 was taken, for each node at $t = 0$, $t = 10$, $t = 20$.

Simulation 2: As before, each node's type is represented by their location, and the grid is 20x20. In addition, we leave 5 empty box lengths above the grid and then have a row of 20 boxes (a 1x20 row) lying above the rest of the grid. This row represents the neighboring community. We give the agents in the 1x20 band of neighboring community nodes above the grid a subsidy of 5 to take action 1: $z_1 - c_1 \sim N(6.5, 10)$. These beliefs and payoffs induce them to all choose action 1, while everyone on the 20x20 grid starts by taking action 0 given the parameter choices.

Note that the neighboring community offers more access to information than the local leader in Simulation 1 because there are twenty neighbors for the community to observe, rather than one local leader. We then simulate and let the agents learn from each other, according to the belief formation rule.We plot the same type of black and white colorbox plot as before for $t = 0$, $t = 10$, and $t = 20$.

B.2 Crisis as a Time of Re-Invention (Figure 1)

Start by assigning every agent to a node on a (20 x 20) grid. This represents their location in characteristic space. In this representation, social distance is the Euclidean distance between nodes/agents. Then, randomly assign each of these nodes to one geographic location on a 20×20 grid as well.

The parameters we use are: Action 0 has payoff $z_0 = 5$, with cost 0; action 1 has payoff $z_1 = 10$, with cost 1. The individual payoff shocks are distributed $\epsilon_{ait} \sim N(0, 100)$ for both actions, independently drawn for actions, agents and time. Prior beliefs about each of the action payoffs are: Project 0 $N(2.5, 10)$, meaning that an agent has prior belief that $E[z_0] = 2.5$, with prior variance 10. Prior beliefs for action 1 are $z_1 \sim N(0, 10)$. With these parameters, all agents initially choose action 0.

After the initial period, we introduce a negative shock to the payoffs of both actions. We permanently change z_0 from 5 to 3, and change z_1 from 10 to 8.

Then, we simulate the choices of all agents for 20 periods, 100 times, and plot the average action in periods 0, 5 and 10. For the first five periods, we control the shocks and hold them fixed across simulations to ensure a crisis-like dynamic. After period 5, shocks are drawn independently across simulations to provide a range of possible future paths. The same average outcomes are plotted in geographic space and characteristic space.

B.3 Digital Social Networks and the Decline of Experts (Figure 3)

In this exercise, there are no choices or payoffs, only agents with geographic locations and characteristics.

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

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

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

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

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

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

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

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

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

C Data Description

C.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 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,[19] 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."

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 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.

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

C.2 Constructed Variables

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

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 $\textit{Occupation Size}_{i,j,k,t}$ and the total number of people of an ethnic group in a PUMA each year is $\textit{Population Size}_{i,j,k,t}$. Therefore, we have:

$$\textit{Ethnic Ratio}_{i,j,k,t} = \frac{\textit{Occupation Size}_{i,j,k,t}}{\textit{Population Size}_{i,k,t}}$$

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

$$\textit{Local Ratio}_{j,k,t} = \frac{\textit{Occupation Size}_{j,k,t}}{\textit{Population Size}_{k,t}}$$

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

$$\textit{Ethnic ORep}_{i,j,k,t} = \textit{Ethnic Ratio}_{i,j,k,t} - \textit{Local Ratio}_{j,k,t}$$

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

$$\textit{Youth Rate}_{i,j,k,t} = \frac{\textit{Occupation Size}_{i,j,k,t}}{\textit{Population Size}_{i,k,t}}$$

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

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

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

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

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

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

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

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

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

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

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

D Appendix-Tables

Table A.1: Impact of a Shock on Occupation Choice (2005/06 and 2010-20)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ethnic ORep_{2005}	0.0351	0.0726**	0.0424	0.132***	0.0735**	0.0413	
	(0.0506)	(0.0356)	(0.0354)	(0.0367)	(0.0352)	(0.0386)	
POST	0.000557	0.000710*		0.000673*	0.000617*		0.000605*
	(0.000375)	(0.000365)		(0.000365)	(0.000371)		(0.000367)
(POST)*(Ethnic ORep_{2005})	0.110***	0.0908**	0.125***	0.0950**	0.0922**	0.129***	0.0872**
	(0.0400)	(0.0387)	(0.0384)	(0.0387)	(0.0390)	(0.0428)	(0.0388)
High Layoffs	-0.0174***	0.00174***	0.000105	0.00154***			
	(0.000681)	(0.000529)	(0.000522)	(0.000525)			
(High Layoffs)*(Ethnic ORep_{2005})	0.156*	0.201***	0.230***	0.163**	0.253***	0.296***	
	(0.0939)	(0.0753)	(0.0746)	(0.0750)	(0.0751)	(0.0802)	
(POST)*(High Layoffs)	-0.00268***	-0.00237***	-0.000408	-0.00226***	-0.00207***		-0.00205***
	(0.000585)	(0.000572)	(0.000562)	(0.000572)	(0.000582)		(0.000575)
(POST)*(High Layoffs)*(Ethnic ORep_{2005})	-0.219***	-0.251***	-0.283***	-0.255***	-0.269***	-0.333***	-0.269***
	(0.0820)	(0.0802)	(0.0792)	(0.0802)	(0.0817)	(0.0865)	(0.0822)
Constant	0.0384***	0.0214***	0.0331***	0.0326***	0.0331***	0.0327***	0.0331***
	(0.000461)	(0.000674)	(0.000144)	(0.000328)	(0.000257)	(0.000118)	(0.000241)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	829690	829690	829690	829690	829580	717246	829220

Clustered standard errors in parentheses.

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

Table A.2: Positive Difference Only (2005/06 and 2010-20)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Ethnic ORep_{2005})	2.010***	0.183***	0.234***	0.183***	0.0682	0.0823	
	(0.0883)	(0.0528)	(0.0580)	(0.0544)	(0.0571)	(0.104)	
POST	0.000821	-0.000530		-0.000542	-0.000672		-0.000675
	(0.000705)	(0.000617)		(0.000617)	(0.000632)		(0.000620)
(POST)*(Ethnic ORep_{2005})	0.0669	0.178***	0.117*	0.181***	0.180***	0.156	0.182***
	(0.0661)	(0.0567)	(0.0645)	(0.0566)	(0.0576)	(0.120)	(0.0568)
High Layoffs	-0.00683***	0.000186	0.000313	-0.000400			
	(0.00152)	(0.00115)	(0.00117)	(0.00115)			
(High Layoffs)*(Ethnic ORep_{2005})	-0.468***	0.284**	0.143	0.325**	0.244	0.129	
	(0.172)	(0.134)	(0.134)	(0.134)	(0.149)	(0.259)	
(POST)*(High Layoffs)	-0.000770	0.0000736	-0.0000738	0.0000280	0.000368		0.000585
	(0.00133)	(0.00123)	(0.00125)	(0.00123)	(0.00128)		(0.00125)
(POST)*(High Layoffs)*(Ethnic ORep_{2005})	-0.450***	-0.459***	-0.290**	-0.441***	-0.451***	-0.335	-0.475***
	(0.152)	(0.141)	(0.142)	(0.142)	(0.149)	(0.266)	(0.146)
Constant	0.0171***	0.0175***	0.0317***	0.0322***	0.0335***	0.0327***	0.0348***
	(0.000947)	(0.00112)	(0.000341)	(0.000601)	(0.000562)	(0.000438)	(0.000349)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	376802	376802	376802	376801	376692	234144	376600

Clustered standard errors in parentheses.

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

Table A.3: Negative Difference Only (2005/06 and 2010-20)

	Dependent variable: Youth Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Ethnic ORep_{2005})	-2.799***	-0.0779	-0.246***	0.0280	-0.0134	-0.0241	
	(0.0936)	(0.0820)	(0.0858)	(0.0853)	(0.117)	(0.150)	
POST	-0.000716	-0.000921		-0.000843	-0.00111		-0.000932
	(0.000884)	(0.000884)		(0.000880)	(0.000937)		(0.000871)
(POST)*(Ethnic ORep_{2005})	-0.0687	-0.150	0.179	-0.141	-0.183	-0.125	-0.156
	(0.109)	(0.109)	(0.126)	(0.109)	(0.116)	(0.237)	(0.109)
High Layoffs	-0.00600***	0.000758	0.000956	-0.0000257			
	(0.00118)	(0.000919)	(0.000915)	(0.000928)			
(High Layoffs)*(Ethnic ORep_{2005})	0.976***	0.0502	0.287**	-0.0258	0.677***	0.949***	
	(0.167)	(0.133)	(0.134)	(0.135)	(0.198)	(0.270)	
(POST)*(High Layoffs)	-0.0000994	-0.000151	-0.000560	-0.000257	-0.000240		-0.000370
	(0.00131)	(0.00131)	(0.00130)	(0.00131)	(0.00139)		(0.00130)
(POST)*(High Layoffs)*(Ethnic ORep_{2005})	0.0125	0.0327	-0.447**	0.00194	0.0116	-0.487	-0.00819
	(0.189)	(0.188)	(0.193)	(0.188)	(0.202)	(0.411)	(0.190)
Constant	0.0119***	0.0160***	0.0305***	0.0319***	0.0331***	0.0327***	0.0318***
	(0.000773)	(0.00210)	(0.000592)	(0.000732)	(0.000825)	(0.000907)	(0.000260)
Occupation FE	No	Yes	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes
Observations	376802	376802	376802	376801	376692	234144	376600

Clustered standard errors in parentheses.

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

Table A.4: Impact of a Shock on Occupation Choice (2005/06 and 2010/11): Ethnicity,Gender

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethnic ORep (Same Gender)$_{2005}$	0.715***	0.585***	0.583***	0.645***	0.584***	0.598***	0.986***	1.030***
	(0.0251)	(0.0185)	(0.0184)	(0.0199)	(0.0185)	(0.0197)	(0.0195)	(0.0210)
POST	0.000918**	0.000887**		0.000864**	0.000845**		0.000849**	
	(0.000432)	(0.000419)		(0.000419)	(0.000430)		(0.000422)	
(POST)*(Ethnic ORep$_{2005}$)	-0.120***	-0.127***	-0.125***	-0.126***	-0.131***	-0.149***	-0.126***	-0.185***
	(0.0256)	(0.0249)	(0.0247)	(0.0249)	(0.0256)	(0.0273)	(0.0251)	(0.0293)
(High Layoffs) *(Ethnic ORep$_{2005}$)	-0.0939**	-0.0353	-0.0316	-0.0532	-0.0318	-0.0395	-0.0695*	-0.108***
	(0.0392)	(0.0356)	(0.0356)	(0.0356)	(0.0366)	(0.0389)	(0.0392)	(0.0418)
(POST)*(High Layoffs)	-0.00320***	-0.00291***	-0.0000729	-0.00284***	-0.00277***		-0.00279***	
	(0.000671)	(0.000658)	(0.000647)	(0.000658)	(0.000678)		(0.000668)	
(POST)*(High Layoffs) *(Ethnic ORep$_{2005}$)	-0.0993**	-0.0916*	-0.0984**	-0.0933*	-0.0890*	-0.0810	-0.0984**	-0.0394
	(0.0505)	(0.0491)	(0.0490)	(0.0491)	(0.0506)	(0.0535)	(0.0497)	(0.0571)
Ethnic ORep (Wrong Gender)$_{2005}$	-0.230***	-0.395***	-0.396***	-0.337***	-0.396***	-0.396***		
	(0.0203)	(0.0121)	(0.0121)	(0.0137)	(0.0118)	(0.0127)		
Constant	0.0385***	0.0184***	0.0330***	0.0326***	0.0331***	0.0331***	0.0331***	0.0331***
	(0.000451)	(0.000875)	(0.000195)	(0.000279)	(0.000212)	(0.000145)	(0.000165)	(0.00000587)
Occupation FE	No	Yes	No	No	No	No	No	No
Occupation-Year FE	No	No	Yes	No	No	No	No	No
Occupation-Ethnicity FE	No	No	No	Yes	No	No	No	No
Occupation-PUMA FE	No	No	No	No	Yes	No	No	No
Occupation-PUMA-Year FE	No	No	No	No	No	Yes	No	No
Occupation-PUMA-Ethnicity FE	No	No	No	No	No	No	Yes	No
Occupation-PUMA-Ethnicity-Year FE	No	Yes	No	No	No	No	No	Yes
Observations	490838	490838	490838	490838	490536	480486	489190	430822

Clustered standard errors in parentheses.

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