

NBER WORKING PAPER SERIES

IS IT THE MESSAGE OR THE MESSENGER?
EXAMINING MOVEMENT IN IMMIGRATION BELIEFS

Hassan Afrouzi
Carolina Arteaga
Emily K. Weisburst

Working Paper 31385
<http://www.nber.org/papers/w31385>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2023

We gratefully acknowledge support from ISERP seed grant from Columbia University, SSHRC Institutional Grant and REG from UCLA. We thank Jonathan Weisburst for sound editing and Michael Floyd for voice recording. We are also grateful to Atom Vayalinkal, Marco Tabellini, Pablo Querubin, Stefano Fiorin, Tianyi Wang, Eric Merkley, Arthur Blouin, Kristen Cornelson, Sandra Black, Sarah Reber, and Felipe Goncalves for helpful feedback as well as seminar participants at UC Merced, UCLA, Claremont McKenna, the University of Toronto, and the Canadian Economic Association (CEA). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Hassan Afrouzi, Carolina Arteaga, and Emily K. Weisburst. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Is it the Message or the Messenger? Examining Movement in Immigration Beliefs
Hassan Afrouzi, Carolina Arteaga, and Emily K. Weisburst
NBER Working Paper No. 31385
June 2023
JEL No. C90,D83

ABSTRACT

How do political leaders affect constituents' beliefs? Is it rhetoric, leader identity, or the interaction of the two that matters? Using a large-scale experiment we decompose the relative importance of partisan messages vs leader sources, in the context of beliefs about immigration. Participants listen to anti-immigrant and pro-immigrant speeches from both Presidents Obama and Trump. These treatments are benchmarked to versions of the speeches recorded by an actor to control for message content, and to non-ideological presidential speeches to control for leader priming. Our findings show that political leader sources influence beliefs beyond the content of their messages in a special case: when leaders deliver unanticipated messages to individuals in their own party. This evidence supports the hypothesis that individuals will "follow their leader" to new policy positions.

Hassan Afrouzi
Department of Economics
Columbia University
420 W 118th Street
New York, NY 10027
and NBER
ha2475@columbia.edu

Emily K. Weisburst
UCLA Luskin School of Public Affairs
Department of Public Policy
337 Charles E Young Dr E
Los Angeles, CA 90095
weisburst@ucla.edu

Carolina Arteaga
University of Toronto
150 St. George Street 300
Toronto, Ontario M5S3G7
Canada
carolina.arteaga@utoronto.ca

A data appendix is available at
<http://www.nber.org/data-appendix/w31385>
A randomized controlled trials registry entry is available at:
<https://www.socialscienceregistry.org/trials/6552>

Statements from political leaders directly influence their followers' beliefs and actions (Zaller (1992) and Lenz (2013)).¹ But, how much of this influence is due to the leaders' brand identity as opposed to the content of their political rhetoric, or to the interaction of the two? The answer to this question is important not only because it enhances our understanding of how political beliefs are formed, but also because it informs policies that aim to either (1) communicate new information effectively,² or (2) correct misconceptions in beliefs.³

Nonetheless, it is difficult to separately identify the role of political messages from the identity of their sources because what leaders say is endogenous to their position. Leaders might choose to cater to the underlying views of their party constituents in order to increase their chances of getting elected, simply echoing and amplifying the views of their electors. Conversely, party constituents might choose to follow the political preferences of their elected leaders.⁴ While these two cases are observationally equivalent, their implications for the roles of leaders differ. Consequently, the ideal experiment that would identify these two separate roles is one that independently varies the identity of the source and the content of a political message.

In this paper, we leverage a novel large-scale experiment with over 13,000 participants to isolate the impacts of political message content and leader sources on beliefs about immigration policy. Immigration is an ideal context to study the determinants of leader influence because it is both a meaningful policy topic for voters and voters' views on immigration vary substantially across party lines.⁵ In our experiment, we use audio recording treatments that are excerpts of *actual* speeches given by Presidents Barack Obama and Donald Trump and compare these treatments to exact anonymous replicas that were recorded by a voice actor. Our constructed speech treatments include both an *anti-immigrant* and a *pro-immigrant* speech from each president. Further, we include additional treatments of audio segments of non-ideological speeches for each president (ceremonial "turkey pardon" speeches on Thanksgiving⁶), for a

¹See also Beaman et al. (2009), Beaman et al. (2012), Bidwell et al. (2020) and Fujiwara and Wantchekon (2013).

²For example to effectively communicate the benefits of a new vaccine (Larsen et al., 2022).

³Research studying interventions relevant to the diffusion of fake news in online forums (Cinelli et al., 2021), beliefs about characteristics of migrants (Alesina et al., 2018), or support for female labor force participation (Bursztyn et al., 2018).

⁴Research in economics predicts that news media will cater to segmented consumer political preferences, potentially exacerbating polarization (Gentzkow and Shapiro, 2006; Baron, 2006; Mullainathan and Shleifer, 2005). Further, individuals who choose to voice opinions in opposition of their political party or group can face costly social sanctions (Bursztyn et al., 2022).

⁵Democrats are generally pro-immigrant while the majority of Republicans hold anti-immigrant views. 52% (70%) of voters in the 2020 (2016) presidential election characterized immigration as being "very important" to their vote (Doherty et al., 2020, 2016). Recent survey evidence shows that 50% of Democrats would like to see immigration levels increase, relative to only 13% of Republicans (Younis, 2020)

⁶The presidential tradition of a ceremonial turkey pardon typically consists of an event where a president gives a speech on the day before Thanksgiving, which includes the presentation of a live turkey who the president spares from being killed for a Thanksgiving dinner.

total of ten speeches and eleven treatment arms, including a control group.⁷ We embed these treatments in an online survey using a sample of Republican and Democrat participants, and stratify the treatment randomization within parties.

Political messages from leaders involve both the content of the message and the identity of the leader, creating a double treatment effect. To better understand this, we use a conceptual framework that decomposes the role of leader sources and partisan messages. In our main decomposition, we conduct an experiment that fixes the content of an immigration speech while varying the source of the statement. This disentangles the double treatment effect into two separate effects: the *anonymous* effect of the message, which captures the message content effect, and the leader source's *persuasion* effect for each message, which captures how the identity of the leader either amplifies or dampens the effect of the message. To create such variation in our experiment, we use exposure to the voice actor version of an immigration speech and the president's version of the same speech. In our experimental results, we find that both of these effects are present.

One possibility in our context is that leader priming, or pure exposure to a leader's brand identity, could contribute to leader-specific treatment effects. To further understand the interaction of leader identity with the content of a message, we conclude our framework with an alternative decomposition that fixes the identity of the president and varies the content of the message. This exercise decomposes the double treatment effect into the *identified* effect of the message coming from a particular leader and the *priming* effect of that leader, which captures how simple exposure to the identity of the leader affects beliefs about immigration. To implement this, we use variation in the exposure to a president delivering an immigration speech and a non-ideological "turkey pardon" speech. Although priming is a theoretical possibility in our setting, we do not find any empirical evidence of a priming effect.

The conceptual framework also provides several useful predictions regarding when particular sources increase the persuasive power of a particular message type. Here, persuasiveness can be characterized by two multiplicative factors: 1) how unexpected a message is when it comes from a particular source and 2) the agent's subjective view of the reliability of that source. Messages are unexpected when they differ from the known reputation of a source, and the subjective reliability of a source could be a function of how much an agent trusts that source and/or prefers to align with that source. Thus, the framework predicts that leaders will have the most influence on beliefs when they express surprising or unexpected messages (e.g. a pro-immigrant message from President Trump) to an audience of supporters who are likely to find the leader to be reliable

⁷Research in economics and psychology has found that cues about social identity can alter perception, beliefs, and actions (e.g. [Cohn and Maréchal, 2016](#)). This last treatment arm of the "turkey pardon" speech allows us to explicitly control for the possibility that exposure to a particular leader primes participants about their political identity and subsequently alters their views on immigration.

(e.g. Republican participants for President Trump messages).

Our decompositions of the determinants of leader influence illustrate that both political messages and sources matter. To provide a benchmark of their relative importance, we first estimate the total combined effect of political message and leader sources, by considering differences between participants who hear a speech from a president’s voice relative to participants who are not exposed to any audio message. We find that participants from both parties update their beliefs based on the total effect of presidential speeches, moving beliefs by $\approx 2 - 8\%$ in the direction of the message, statistically significant in seven out of eight cases, and with effects that are larger for speeches that oppose the party prior.

Implementing our main decomposition described above, we separate the total effect of the president’s speeches between the effect of an anonymous message and the source persuasion effect. We estimate that in nearly all cases, the effect of the anonymous message drives the majority of the total effect. Our findings on the importance of message content are notable given that the rhetoric used in the speeches for this experiment are emotional and political statements about immigration policy, and do not include any factual information or supporting evidence. These results suggest that the increasingly divergent political views we observe across parties may be partly a function of a lack of exposure to opposition views.

Turning to the role of leader persuasion, which captures the added effect of a specific source delivering a particular message, we find that this channel is only important in one specific and symmetric case: when a party leader delivers a message against party lines to members of his own party. That is, Obama (Trump) has an additional source persuasion effect only for Democrats (Republicans) when giving a speech that is anti-immigrant (pro-immigrant). For these treatments, we find that the source persuasion effect comprises 44-60% of the total movement in beliefs. These effects serve to reduce the distance between party beliefs (“partisan polarization”) in the treatment groups, relative to the control group, because the source persuasion effects are only present for speeches that oppose the party prior.

The overall implication of the results is that *supporters of a leader are willing to follow their leaders to a new position*. In our context, this leads to a 30% total reduction in partisan polarization in beliefs about immigration, with half of this effect attributable to leader persuasion. However, in our general framework, any surprising position taken by a trusted leader could serve to sway beliefs, which could alternatively increase political division in a different context. The source persuasion findings underscore the importance of particular messengers delivering partisan statements.

We take several steps to establish the robustness of our findings. First, in our experimental design, we follow recommendations of Haaland et al. (2020) to minimize experimenter demand effects. In particular, we present a neutral framing for recruitment after which respondents

complete the survey on their own technological devices without the physical presence of a researcher. This approach maximizes privacy and anonymity and potentially increase truthful reporting (Ong and Weiss, 2000). We also expose participants to a only single treatment type, which cannot reveal the overarching variation across treatments of the study or research question to a particular participant. Second, we follow the procedure outlined in List et al. (2021) to establish that our results are robust to multiple hypothesis testing concerns. Finally, we show that our findings are robust to changes in the empirical specification, sample restrictions, and construction of the outcome measures.

This study builds on a large literature in behavioral economics, political economy, and political science about the nature of bias in information sources and the impact that new information can have on beliefs. In particular, our findings on persuasion of political leaders is closely related to Chiang and Knight (2011) who find that newspaper endorsements of political candidates are most influential when they come from an unexpected source. In contemporaneous work, Larsen et al. (2022) document similar patterns by finding that a video of Donald Trump on Fox News encouraging the take up of the Covid vaccine increases vaccination rates. These findings are consistent with our results and model framework in the immigration context. We exploit the richness of our experimental design, to extend previous work and find that voter alignment or voter's trust, play a key role in explaining the influence of unexpected or counter-stereotypical information.

More broadly, our work builds on a large empirical and theoretical literature in political science (see Druckman (2022) for a review). These studies have examined the role of various factors in shaping the effects of party cues and leader persuasion, including the salience of party/group/elite identification, the context, and the media environment. Here, researchers have found that agents may be more skeptical of sources that do not align with their ideology (Chopra et al., 2022), use party cues as a heuristic to validate information in an efficient way (Cohen, 2003; Kam, 2005), and can engage in motivated reasoning, which posits that people distort their inference process in the direction of states they find more attractive (Taber and Lodge, 2006).⁸ This work also focuses on understanding the psychological mechanisms that underly the effects of party cues and leader persuasion, such as cognitive heuristics, affective reactions and motivated reasoning (see, e.g., Bullock, 2011; Nicholson, 2012).⁹ We contribute to this literature by documenting when and how leader identity matters for updating immigration beliefs. In particular, our experiment allows us to vary message ideology for two different

⁸See also (Berinsky, 2009; Lenz, 2013; Leeper and Slothuus, 2014; Bakker et al., 2020; Thaler, 2021; Tappin et al., 2023), among others.

⁹See, also, Merkley and Stecula (2021); Bakker et al. (2020); Bullock (2020); Barber and Pope (2019); Broockman and Butler (2017); Boudreau and MacKenzie (2014); Druckman et al. (2013); Goren et al. (2009); Gilens and Murakawa (2002).

leaders on opposite ends of the political spectrum, and to compare these treatments to identical messages with an ambiguous source, as well as to non-ideological messages from these leaders.

Furthermore, our work is also related to the economic literature that studies whether partisan rhetoric from leaders, rather than factual information, can sway beliefs. First, individuals are often aware of biases in new information (Gentzkow et al., 2018), perceive different biases from identical messages presented from different sources (Baum and Gussin, 2008), or interpret identical messages in different ways given their priors (Andreoni and Mylovanov, 2012).¹⁰ Second, a number of research studies have examined whether curated messages can move beliefs in an intended direction.¹¹ Work in this area has found that individuals have large misperceptions about facts relevant to policy issues, and that individuals may update their beliefs when provided with new *fact-based* information.¹² But this research may not translate to our research question testing the messaging effects of *emotion-based* partisan speeches from leaders. By using actual audio from presidential speeches in lieu of listed statistics or text narratives, we attempt to approximate the biased, incomplete, and sometimes inaccurate political messages that are commonly faced by voters.

In the context of immigration, there is a large literature that documents substantial misperceptions in terms of the size of the immigrant population, their origins, and immigrant crime rates (Abramitzky and Boustan, 2022). It has been extensively found that providing accurate information does indeed help to correct these beliefs (Grigorieff et al., 2020; Hopkins et al., 2019; Alesina et al., 2018); however, there is mixed evidence on whether this updating translates into changes in policy preferences and views (Hopkins et al., 2019; Alesina et al., 2018). The approach taken in our paper differs from this literature as we leverage emotional speeches from political leaders with distinct party affiliations and positions on immigration. Furthermore, our outcome variables are a mix of subjective views on immigrants and immigration policy preferences and not elicited guesses about statistics pertaining characteristics of the immigrant population. Our results suggest that immigration beliefs are indeed malleable, but that the extent of these updates depends on both the direction of the message and its messenger.

Finally, our treatments that use presidential speeches on “turkey pardons,” that are irrelevant to immigration, speak to the literature on priming effects in behavioral economics. Studies in this literature have found that exposure to cues which remind individuals of their ethnic, religious, and/or cultural identity can meaningfully alter beliefs, economic decisions, and risk preferences.¹³ Our null finding of this form of priming related to political party affiliation is

¹⁰See, also, Baysan (2021); Fryer Jr et al. (2019); Benoît and Dubra (2019).

¹¹e.g. Levy (2021); Song (2021); Durante et al. (2019); Kalla and Broockman (2018); Martin and Yurukoglu (2017); Adena et al. (2015); Enikolopov et al. (2011); DellaVigna and Kaplan (2007); Gerber et al. (2009).

¹²e.g. Haaland and Roth (2020); Bursztyn et al. (2018); Cruces et al. (2013).

¹³e.g. Akerlof and Kranton (2000); Callen et al. (2014); Cohn et al. (2015); Benjamin et al. (2016); Cohn et al. (2015).

interesting in the context of a broad literature that highlights the importance of identity cues in signaling group norms (Cohn and Maréchal, 2016).

1 Experiment Design

Our experiment is embedded within an online survey, where participants are exposed to different audio segments of presidential speeches. Participants are asked a series of background questions on demographics, political views, and news consumption prior to treatment. After treatment, participants are asked questions on their views on immigration. We explicitly stratify our sample by political party and recruited participants who identify as Republicans or Democrats, as we expect that respondents will interpret treatment through the lens of political identity. The experiment was pre-registered with the American Economic Association and was conducted during the period of the 2020 presidential election between October 16 - November 10, 2020.¹⁴

The experiment contains 11 treatment arms for each political party. These include 4 president immigration speech segments, consisting of one pro-immigrant and one anti-immigrant speech for *both* Presidents Donald Trump and Barack Obama, and replicate versions of these 4 speeches recorded by a voice actor. We additionally include 2 presidential speeches with no ideological or immigration content; these speeches are ceremonious “turkey pardon” addresses for Trump and Obama.¹⁵ The last arm of the study is a control group that is not exposed to any audio speech treatment. Figure 1 depicts the experiment design.¹⁶

The source material for the treatments are actual speeches delivered by each president. The audio clips used in the survey are extracted excerpts from these speeches. It is important to note that there is *no deception* used in this study, as speeches are always introduced to study participants as “an excerpt of a presidential speech made by President [Donald Trump or Barack Obama].” After the edited speeches were constructed, we hired a voice actor to record replicate versions of the speech segments. The actor versions of speeches are introduced to study participants as “an excerpt of a presidential speech read by an actor.” The actor treatments do not reveal any source for participants, and they likewise do not falsely attribute statements to an alternative source.

We use addresses given by Barack Obama on November 20, 2014 and by Donald Trump

¹⁴The pre-registration number is AEARCTR-0006552.

¹⁵As noted above, the holiday tradition of the turkey pardon typically consists of an event where a president gives a speech on the day before Thanksgiving, which includes the presentation of a live turkey who the president ceremoniously spares from being killed for a Thanksgiving dinner.

¹⁶We note that our experiment includes only one president source from each party by design; as a result, we are not able to separate any potential differences between Presidents Obama or Trump from other partisan figures within their parties. We use the term “leader” as a general term associated with the experimental effect of Presidents Obama and Trump, and note that the distinction of these figures versus other figures within the same ideology is an interesting avenue for future research.

on January 19, 2019 as source material. The purpose of the Obama speech was to introduce protections for undocumented immigrants, including the Deferred Action for Childhood Arrivals (DACA) program, which created new safeguards for individuals who arrived in the U.S. illegally as children. Coupled with these reforms, the Obama speech also outlined additional provisions for border security. The purpose of the Trump speech was to provide a proposal to end a government shutdown related to immigration policy negotiations with Congress; this proposal included both new border security programs as well as a concession to protect DACA recipients. Both speeches were televised from the White House, prepared in advance, and delivered using a teleprompter. We are able to extract both a pro-immigrant and an anti-immigrant segment from each speech because each address contains proposals to provide protections for immigrants as well as proposals to curb illegal immigration. Similarly, we use original “turkey pardon” speeches to compose a non-ideological treatment speech that contains no information about immigration for each president.¹⁷

1.1 Immigration Speech Treatments

Below, we include the text of the four immigration speech treatments. Supplemental Materials Appendix S2 includes text and audio web links for the original full-length speeches and each excerpted treatment speech.

Trump Anti-Immigrant Speech “We believe in a safe and lawful system of immigration, one that upholds our laws, our traditions and our most cherished values. Unfortunately, our immigration system has been badly broken for a very long time. Over the decades, many presidents and many lawmakers have come and gone, and no real progress has been made on immigration.

We are now living with the consequences – and they are tragic – brought about by decades of political stalemate, partisan gridlock, and national neglect. Illegal immigration reduces wages and strains public services. The lack of border control provides a gateway, and a very wide and open gateway, for criminals and gang members to enter the United States. I want this to end; it’s got to end now. These are not talking points. These are the heartbreaking realities that are hurting innocent, precious human beings every single day on both sides of the border.

The good news is these problems can all be solved, but only if we have the political courage to do what is just and what is right. Both sides in Washington must simply come together, listen to each other, put down their armor, build trust, reach across the aisle and find solutions. The proposal I will outline today is based, first and foremost, on input from our border agents and Homeland Security professionals – and professionals they are; they know what they’re doing. Our plan includes the following: \$805 million for drug detection technology to help secure our ports of entry. An additional 2,750 border agents and law enforcement professionals, 75 new immigration judge teams to reduce the court backlog of - believe it or not - almost 900,000 cases. This is a common-sense compromise both parties should embrace. If we are successful in this effort, we will then have the best chance in a very long time at real bipartisan immigration reform. Any reforms we make to our immigration system will be designed to improve your lives, make your communities safer, and make our nation more prosperous and secure for generations to

¹⁷We use the first Thanksgiving turkey pardon for each president, delivered on November 25, 2009 by Obama and on November 21, 2017 by Trump. These speeches are edited only for length.

come. Thank you, and God bless America. Thank you."

Obama Anti-Immigrant Speech "Today, our immigration system is broken, and everybody knows it. Families who enter our country the right way and play by the rules watch others flout the rules. Business owners who offer their workers good wages and benefits see the competition exploit undocumented immigrants by paying them far less. All of us take offense to anyone who reaps the rewards of living in America without taking on the responsibilities of living in America. Millions of us, myself included, go back generations in this country, with ancestors who put in the painstaking work to become citizens. So we don't like the notion that anyone might get a free pass to American citizenship. I know that some worry immigration will change the very fabric of who we are, or take our jobs, or stick it to middle-class families at a time when they already feel like they've gotten the raw end of the deal for over a decade. I hear these concerns.

It's been this way for decades. And for decades, we haven't done much about it. When I took office, I committed to fixing this broken immigration system. And I began by doing what I could to secure our borders. Today, we have more agents and technology deployed to secure our southern border than at any time in our history. And over the past six years, illegal border crossings have been cut by more than half. We'll build on our progress at the border with additional resources for our law enforcement personnel so that they can stem the flow of illegal crossings, and speed the return of those who do cross over. Even as we are a nation of immigrants, we are also a nation of laws. Undocumented workers broke our immigration laws, and I believe that they must be held accountable – especially those who may be dangerous. That's why, over the past six years, deportations of criminals are up 80 percent. And that's why we're going to keep focusing enforcement resources on actual threats to our security.

Now here's the thing: we expect people who live in this country to play by the rules. We expect that those who cut the line will not be unfairly rewarded. If you're a criminal, you'll be deported. If you plan to enter the U.S. illegally, your chances of getting caught and sent back just went up. Now, I continue to believe that the best way to solve this problem is by working together to pass that kind of common sense law. I want to work with both parties to pass a more permanent legislative solution. Because for all the back-and-forth of Washington, we have to remember that this debate is about something bigger. It's about who we are as a country, and who we want to be for future generations. Thank you, God bless you, and God bless this country we love."

Trump Pro-Immigrant Speech "Just a short time ago, I had the honor of presiding over the swearing in of five new great American citizens. It was a beautiful ceremony and a moving reminder of our nation's proud history of welcoming immigrants from all over the world into our national family. I told them that the beauty and majesty of citizenship is that it draws no distinctions of race or class or faith or gender or background. All Americans, whether first generation or tenth generation, are bound together in love and loyalty, friendship and affection. We're all equal. We are one team and one people proudly saluting one great American flag.

Unfortunately, our immigration system has been badly broken for a very long time. Over the decades, many presidents and many lawmakers have come and gone, and no real progress has been made on immigration. The good news is these problems can all be solved, but only if we have the political courage to do what is just and what is right. Both sides in Washington must simply come together, listen to each other, put down their armor, build trust, reach across the aisle and find solutions.

In order to build the trust and goodwill necessary to begin real immigration reform, there are two elements to my plan. Number one is three years of legislative relief for 700,000 DACA recipients brought here unlawfully by their parents at a young age many years ago. This extension will give them access to work permits, social security numbers, and protection from deportation, most importantly. Secondly, our

proposal provides a three-year extension of temporary protected status or TPS. This means that 300,000 immigrants whose protected status is facing expiration will now have three more years of certainty, so that Congress can work on a larger immigration deal, which everybody wants – Republicans and Democrats. And our farmers and vineyards won't be affected because lawful and regulated entry into our country will be easy and consistent. This is a common-sense compromise both parties should embrace. If we are successful in this effort, we will then have the best chance in a very long time at real bipartisan immigration reform. Any reforms we make to our immigration system will be designed to improve your lives, make your communities safer, and make our nation more prosperous and secure for generations to come. Thank you, and God bless America. Thank you."

Obama Pro-Immigrant Speech "For more than 200 years, our tradition of welcoming immigrants from around the world has given us a tremendous advantage over other nations. It's kept us youthful, dynamic, and entrepreneurial. It has shaped our character as a people with limitless possibilities – people not trapped by our past, but able to remake ourselves as we choose. But today, our immigration system is broken, and everybody knows it. It's been this way for decades. And for decades, we haven't done much about it. When I took office, I committed to fixing this broken immigration system. We need more than politics as usual when it comes to immigration; we need reasoned, thoughtful, compassionate debate that focuses on our hopes, not our fears.

So we're going to offer the following deal: If you've been in America for more than five years; if you have children who are American citizens or legal residents; if you register, pass a criminal background check, and you're willing to pay your fair share of taxes – you'll be able to apply to stay in this country temporarily, without fear of deportation. You can come out of the shadows and get right with the law. That's what this deal is. We'll take steps to deal responsibly with the millions of undocumented immigrants who already live in our country. After all, most of these immigrants have been here a long time. They work hard, often in tough, low-paying jobs. They support their families. They worship at our churches. Many of their kids are American-born or spent most of their lives here, and their hopes, dreams, and patriotism are just like ours.

Now, I continue to believe that the best way to solve this problem is by working together to pass that kind of common sense law. I want to work with both parties to pass a more permanent legislative solution. Because for all the back-and-forth of Washington, we have to remember that this debate is about something bigger. It's about who we are as a country, and who we want to be for future generations. Thank you, God bless you, and God bless this country we love."

How similar are the anti-immigrant (or pro-immigrant) speech treatments for the two presidents? While our experiment does not directly compare treatment effects across presidents, as the extracted segments are sourced from two separate original speeches, the Trump and Obama speeches are in fact similar in their rhetoric and ideology. As a direct test of the similarity of the statements across presidents, we asked participants how anti-immigrant (or pro-immigrant) they felt each speech was after treatment, on a scale of 0 to 100. Figure 3 shows the perceived degree of sentiment *strength* for the speeches, or the participant perception that an anti-immigrant (pro-immigrant) speech is anti-immigrant (pro-immigrant). We plot perceptions from the actor speech groups so that responses are not colored by the president's reputation.¹⁸ Strikingly, the Trump and Obama speeches are perceived nearly identically within message type, for both

¹⁸Figure A1 shows corresponding estimates for all treatment groups.

Republicans and Democrats. Further, the perceived strength of statements across parties is also similar, with both groups viewing the pro-immigrant speeches as somewhat stronger than the anti-immigrant speeches. Again, while the experiment design does not rely on any direct comparisons *across* president or *across* party, these results are reassuring in that they show that the constructed speeches are similar for both Presidents and message types.

2 Conceptual Framework

A statement from a political leader is a double treatment with a message's *content* and its *source*. This section provides two decompositions of this treatment that guide our experimental design.

Setting. Let Ω represent a binary sample space for an individual's beliefs about immigration on whether it is favorable or unfavorable. A belief system about immigration is a probability space on events that also incorporate messages and sources. We define the augmented outcome space $\tilde{\Omega}$ as $\Omega \times M \times S$, where M denotes all messages an individual might expect from a leader, and S represents all potential sources. An individual's subjective beliefs is a probability space on $\tilde{\Omega}$. Now, consider a treatment group hearing a political message m from a leader s . This treatment group's belief about an outcome $\omega \in \Omega$ relative to the control group is given by Equation 1:

$$\frac{1 + P(\omega|m, s)}{1 + P(\omega)} \quad (1)$$

where $P(\omega) \equiv \sum_{s' \in S, m' \in M} P(\omega, m', s')$ is the unconditional belief of the control group about ω .¹⁹ We now present two decompositions of this treatment effect that helps us disentangle the treatment effect of the message m from the source s .

2.1 Persuasion Decomposition: Fixing the Message

Consider an additional treatment group exposed to a political immigration message with content m from an anonymized source—i.e., a voice actor that reads a message from some $s' \in S$ but the identity of s' is not revealed independently. This group's belief about ω conditional on m can be expressed as $P(\omega|m) = \sum_{s' \in S} P(s'|m)P(\omega|s', m)$. We can then decompose Equation 1 as:

$$\ln\left(\frac{1 + P(\omega|m, s)}{1 + P(\omega)}\right) = \underbrace{\ln\left(\frac{1 + P(\omega|m)}{1 + P(\omega)}\right)}_{\beta_m \equiv \text{anonymous message}} + \underbrace{\ln\left(\frac{1 + P(\omega|m, s)}{1 + P(\omega|m)}\right)}_{\beta_{ms} \equiv \text{source persuasion}} \quad (2)$$

Here, the first term represents the *anonymous message* effect, which is the pure impact of the content without revealing the source. Then the residual in the second term, called the source

¹⁹We have added 1 to both the denominator and numerator to avoid dividing by zero when $P(\omega) = 0$, i.e. when an individual finds immigration fully favorable (or unfavorable). This allows us to keep all participants in the analysis, regardless of the value of their belief outcome after treatment. The important notion here is that both denominator and numerator are transformed symmetrically so that when $P(\omega|m, s) = P(\omega)$, the relative effect $\frac{1 + P(\omega|m, s)}{1 + P(\omega)}$ is also 1.

persuasion effect, measures how knowing the source's identity influences the audience's beliefs about immigration, fixing the message.

Bayesian Interpretation of Anonymous Message Effect. Using Bayes' law, we have:

$$P(\omega|m) = P(\omega) \times \underbrace{\frac{P(m|\omega)}{P(m)}}_{\text{odds ratio} \equiv \Theta(m|\omega)} \quad (3)$$

The left-hand side represents an individual's beliefs after receiving message m from an anonymous source, while $P(\omega)$ on the right-hand side is the control group's belief. The third term, the odds ratio of message m , shows the likelihood difference of the message in a specific state of the world ω . Beliefs of control and treatment groups differ only if the likelihood varies across states, i.e., $P(m|\omega) \neq P(m) = \sum_{\omega' \in \Omega} P(\omega')P(m|\omega')$. We can formalize this observation as follows.

Remark 1. For $\omega \in \Omega$, let $\neg\omega$ denote the event where the outcome is not ω , then the beliefs of the treatment and control groups in the experiment above are different if and only if the likelihood of the message m is different conditional on ω versus $\neg\omega$:

$$P(m) \neq P(m|\omega) \Leftrightarrow P(m|\neg\omega) \neq P(m|\omega)$$

We can show that the anonymous message effect, β_m , increases with the likelihood ratio $\Theta(m|\omega)$ and is non-zero only if this ratio differs from one. Thus, beliefs are updated if the message is perceived to be informative of the true state of the world regarding immigration.

Bayesian Interpretation of Source Persuasion Effect. Considering a treatment group receiving a message m from source s . This group's belief for ω , $P(\omega|m, s)$, relative to the belief of a group receiving the same message from an anonymized source, $P(\omega|m)$, represents the source's persuasion effect as in Equation 2. Using Bayes' law, we can write:

$$P(\omega|m, s) = P(\omega|m) \times \underbrace{\frac{P(s|\omega, m)}{P(s|m)}}_{\Theta(s|\omega, m) \equiv \text{odds ratio of } s|m} \quad (4)$$

which shows that the persuasion effect is related to the odd ratio of the source *conditional* on the message m . The following remark illustrates when the source will be persuasive.

Remark 2. For $s \in S$, let $\neg s$ denote the event that the source is not s . Then, there is source persuasion if (1) revealing the identity of the source is surprising *and* (2) the source is more reliable than other sources for the outcome ω :

$$\Theta(s|\omega, m) > 1 \Leftrightarrow \underbrace{(1 - P(s|m))}_{\text{surprise}} \times \underbrace{(P(\omega|s, m) - P(\omega|\neg s, m))}_{\text{(relative) reliability of source}} > 0 \quad (5)$$

While the surprise term is always non-negative, the reliability term can be positive or negative,

governing the direction of persuasion.²⁰ If either term is zero – i.e., revealing the source isn’t surprising for a particular message or the source isn’t more or less reliable than others – there is no persuasion.

Furthermore, to account for the possibility of having subjects with beliefs that assign zero probabilities to different events, we can map this prediction to a formulation in terms of $\ln(1 + P(\omega|m, s))$ by rewriting the source persuasion effect, $\beta_{m,s}$, as:

$$\beta_{ms} = \ln\left(\frac{1 + P(\omega|m, s)}{1 + P(\omega|m)}\right) = \underbrace{\ln\left(1 + \frac{P(\omega|m)}{1 + P(\omega|m)} \times (\Theta(s|\omega, m) - 1)\right)}_{\text{persuasion effect of } s \equiv \Delta(s|m, \omega)} \quad (6)$$

$\Delta(s|m, \omega)$ characterizes the source persuasion effect $\beta_{m,s}$: it increases with the odds ratio $\Theta(s|m, \omega)$ and equals zero when there’s no persuasion. Generally, $\Delta(s|m, \omega)$ allows the source to amplify the message’s effect on beliefs about ω through surprise and reliability channels.

2.2 Alternative Decomposition for Priming Effect: Fixing the Source

Now consider a second treatment group treated with an irrelevant message m^0 from a fixed source s that’s independent of the outcome $\omega \in \Omega$ (in our experiment, m^0 is a turkey pardon speech from a specific president). This group’s belief corresponds to $P(\omega|s, m^0)$, which, due to the independence assumption, equals $P(\omega|s)$.²¹ This treatment enables us to decompose the treatment effect in Equation 1 as:

$$\ln\left(\frac{1 + P(\omega|m, s)}{1 + P(\omega)}\right) = \underbrace{\ln\left(\frac{1 + P(\omega|s)}{1 + P(\omega)}\right)}_{\beta_s \equiv \text{source priming}} + \underbrace{\ln\left(\frac{1 + P(\omega|m, s)}{1 + P(\omega|s)}\right)}_{\beta_{sm} \equiv \text{identified message}} \quad (7)$$

In this equation, the first term under braces captures the unconditional effect of the source’s identity on the audience’s belief about ω . This source effect for an irrelevant message is equivalent to the isolated source effect across all possible messages and thus corresponds to the *priming* effect of the source. The second term measures the effect of the immigration message m on the audience’s beliefs about ω , knowing that m was delivered by source s . We call this the *identified message* effect of m from source s .

Bayesian Interpretation of Source Priming Effect. The elicited belief from a treatment group about $\omega \in \Omega$ that hears a message m^0 which is irrelevant to immigration from a source s is given by $P(\omega, m^0|s) = P(\omega|s)$. We can then apply Bayes’ rule to write this in terms of the odds ratio for

²⁰As shown in the results, the sign of the estimated coefficient is also a function of the outcome designation of ω , which we set as anti-immigrant (likelihood that immigration is “unfavorable”) for Republicans and pro-immigrant for Democrats (likelihood that immigration is “favorable”). For example, an anti-immigration treatment that persuades Democrats results in a negative persuasion effect because it makes participants more anti-immigrant, when the outcome for this group is degree pro-immigrant.

²¹Note that by independence we have $P(\omega, m^0|s) = P(\omega|s)P(m^0|s) \Rightarrow P(\omega|s, m^0) = P(\omega, m^0|s)/P(m^0|s) = P(\omega|s)$.

the source,

$$P(\omega|s) = P(\omega) \times \underbrace{\frac{P(s|\omega)}{P(s)}}_{\text{odds ratio} \equiv \Theta(s|\omega)} = P(\omega) \times \underbrace{\frac{P(s, \omega)}{P(\omega)P(s)}}_{\text{odds ratio} \equiv \Theta(s|\omega)} \quad (8)$$

where the left-hand side captures the beliefs of the treated individual, $P(\omega)$ on the right-hand side represents the belief of the control group (or prior), and the third term—the odds ratio of the source—captures the priming effect for the source. It is, useful in this case to think about the treatment effect in terms of independence of outcome and source:

Remark 3. The beliefs of the primed individual should be the same as the belief of a subject in the control group only if the identity of the source and the outcome ω are subjectively independent:

$$P(\omega, s) = P(\omega)P(s) \Leftrightarrow \Theta(s|\omega) = 1$$

Furthermore, to map this prediction to our formulation in terms of $\ln(1 + P(\omega|s))$, we can re-write the source priming effect, β_s , in Equation (7) as

$$\beta_s = \ln\left(\frac{1 + P(\omega|s)}{1 + P(\omega)}\right) = \ln\left(1 + \underbrace{\frac{P(\omega)}{1 + P(\omega)}}_{\text{priming effect of } s \equiv \Delta(s|\omega)} \times (\Theta(s|\omega) - 1)\right) \quad (9)$$

Thus, according to Bayes' law, the term $\Delta(s|\omega)$ characterizes the priming effect of the source: it is increasing in the odds ratio $\Theta(s|\omega)$ and is equal to zero if ω and s are perceived to be independent, i.e. if $\Theta(s|\omega) = 1$. Remark 3 shows that if individuals subjectively link a particular leader's identity to a specific state of the world, they'll be subject to priming effects when exposed to *any* message from that leader. However, if a person thinks exposure to an uninformative message from a president is unrelated to the actual state of immigration, they won't update their beliefs due to the source priming effect.

Bayesian Interpretations of Identified Message Effect. Similarly, to interpret the identified message effect, we can use Bayes' law to write

$$P(\omega|m, s) = P(\omega|s) \times \underbrace{\frac{P(m|\omega, s)}{P(m|s)}}_{\text{odds ratio} \equiv \Theta(m|\omega, s)} \quad (10)$$

where the left-hand side captures the beliefs of an individual treated with message m from source s , $P(\omega|s)$ on the right-hand side represents the belief of the control group that was treated with an irrelevant message m^0 from the same source (so that their belief is $P(\omega|m^0, s) = P(\omega|s)$), and the third term—the odds ratio of the message m conditional on the source s —captures how the likelihood of the message is different in the particular state of the world ω . It is then easy to see that the beliefs of control and treatment groups in this experiment are different only if this likelihood is different in different states of the world, i.e. $P(m|\omega, s) \neq P(m|s) =$

$\sum_{\omega' \in \Omega} P(\omega')P(m|\omega', s)$. In particular, we can derive the following remark.

Remark 4. For $\omega \in \Omega$, let $\neg\omega$ denote the event where the outcome is not ω , then the beliefs of the treatment and control groups in the experiment above are different if and only if the likelihood of the message m coming from source s is different conditional on ω versus $\neg\omega$:

$$P(m|s) \neq P(m|s, \omega) \Leftrightarrow P(m|s, \neg\omega) \neq P(m|s, \omega)$$

Finally, similar to our exercise above, we can show that the identified message effect, β_{sm} , is increasing in the likelihood ratio $\Theta(m|\omega, s)$ and is non-zero if and only if this likelihood ratio is different than one. We can interpret this effect as moving individuals when they believe that a president's statement provides new information about the true state of the world regarding immigration, relative to a benchmark of exposure to hearing that president's voice, or adjusted for any source priming effect.

3 Empirical Framework

This study uses a randomized controlled experiment, and the empirical approach will use simple comparisons of treatment arms that leverage this randomization.

Our first objective is to test the total impact of an anti-immigration or a pro-immigration statement from a party leader on participant beliefs about immigration. To do this, we compare the president recordings of a particular immigration message to the control group that received no audio treatment, separately by political party, p . The regression corresponds to:

$$\ln(P(y_i) + 1) = \beta_0 + \beta_t 1[Message_i = 1] \times 1[Source_i = 1] + \gamma X_i + \epsilon_i \quad (11)$$

The coefficient β_t comprises the *total* combined impact of a statement, including both the source and message effects. We denote $Message_i = 1$, as treatment arms which discuss immigration topics, and $Source_i = 1$; as treatment arms where we have the presidents delivering the speeches. Outcomes are constructed as $\ln(P(y_i) + 1)$, where the outcome, $P(y_i)$, is the index of probability anti-immigrant (pro-immigrant) for Republicans (Democrats). The intercept, β_0 , is the average belief in the control group. In our preferred specifications, we also include a vector of pre-treatment control variables selected from a double lasso procedure, X_i , to increase precision of the estimates (see Online Appendix A3).²² As discussed below and shown in Figure A5, the results are robust to excluding these controls.

²²The procedure consists of first including all potential control variables from baseline and converting categorical variables into sets of indicators. Our set of potential controls is quite rich given the large number of survey questions we ask to participants prior to the treatment. We keep the controls selected from the LASSO minimization of a regression which predicts treatment assignment. We then repeat this exercise for a regression that predicts the outcome variable and take the union of controls selected from both procedures as our X_i controls. The list of controls included in the models is outlined in Online Appendix A3.

Next, we decompose this total effect, β_t , to investigate whether and how leaders may persuade constituents. First, we decompose the total effect, β_t , between the anonymous message effect, β_m , and the source's persuasion power for that message, β_{ms} . For this decomposition, we fix the message of a speech, and vary the source between the actor and the president. Each regression includes a president immigration speech, the voice actor version of the same speech, and the "no audio" control group. The regression equation is:

$$\ln(P(y_i) + 1) = \beta_0 + \beta_m 1[Message_i = 1] + \beta_{ms} 1[Message_i = 1] \times 1[Source_i = 1] + \gamma X_i + \epsilon_i \quad (12)$$

Where $1[Message_i = 1]$, includes both treatment groups with a particular message, or the actor and president speech versions, and $1[Message_i = 1] \times 1[Source_i = 1]$, includes only the president's version of the immigration speech. Here β_m captures the effect of an anonymous message, that is an immigration effect delivered by the actor; and β_{ms} captures the source persuasion effect, or the difference in beliefs from listening to the same words, delivered by the president relative to the actor.

In a second alternative decomposition, we test for any presence of source priming effects. Here, we decompose the effect between source priming, or the effect of a source when content is not relevant; and the effect of the immigration speech voiced by Obama or Trump, conditional on hearing the president. That is, we fix the president source as either Trump or Obama's voice, and vary the message between the turkey pardon and an immigration speech. Each regression includes a president immigration speech, a president turkey pardon speech, and the "no audio" control group. The relevant regression is:

$$\ln(P(y_i) + 1) = \beta_0 + \beta_s 1[Source_i = 1] + \beta_{sm} 1[Source_i = 1] \times 1[Message_i = 1] + \gamma X_i + \epsilon_i \quad (13)$$

Where the term $1[Source_i = 1]$, includes both treatment groups with a president's voice, or the turkey pardon and the immigration speech, and the second term, $1[Message_i = 1] \times 1[Source_i = 1]$, includes only the president's immigration speech. This approach decomposes β_t into β_s and β_{sm} . We interpret β_s as source priming, or the difference between beliefs after listening to the turkey pardon audio, relative to the control group; and β_{sm} , the identified message effect, as the effect of an immigration statement from an identified president, net of priming effects, or relative to the president turkey pardon message.

4 Experiment Setting and Structure

4.1 Survey Instrument

We recruited participants who are eligible voters and identify as either Republicans or Democrats through a company called Cloud Research. Cloud Research partners with a number of different online survey panels to compile a sample that targets particular demographic groups and ensures basic quality standards for participant responses. We targeted a participant pool that would mimic the demographic characteristics of the Democratic and Republican parties, such that our results would best represent responses of individuals in these parties.

Prior to treatment, we ask a series of demographic questions, including gender, race, age, education, employment status, party affiliation, and the candidate that the participant supported in the 2016 presidential election. We additionally ask a background question on four different political issues; immigration, gun control, abortion, healthcare, and taxes, as well as which of these issues is most important to the respondent's vote. We also ask how often participants receive their news from different modes (e.g. newspaper, TV, facebook), and different sources (e.g. Fox News, MSNBC). Lastly, we ask participants to state whether they are "fans of" a list of public figures and celebrities, which includes both Donald Trump and Barack Obama. Our questions about views on immigration and approval of Presidents Trump and Obama are embedded within larger lists of policy issues and public figures so as not to prime participants prior to treatment.

After treatment, we ask several questions about participant immigration views. We ask whether participants favor or oppose proposals to: expand construction of a wall on the U.S./Mexico border, hire more border patrol agents, require businesses to check the immigration status of workers, deport all immigrants living in the U.S. without legal status, deport the subset of this population with a criminal record, allow immigrants living in the U.S. without legal status to become citizens, or allow the subset of this population who came to the U.S. without legal status as children to become citizens (DACA recipients). We also ask whether participants view the following groups as positively or negatively contributing to U.S. society: immigrants working legally, immigrants not working legally, immigrants from English-speaking countries, immigrants from Spanish-speaking countries, and "dreamers" or undocumented immigrants brought to the U.S. by their parents as children (DACA recipients). Lastly, we ask participants whether immigrants benefit the economy, commit a disproportionate share of crimes, and if they feel immigration should be increased, decreased, or kept at present levels, as well as their overall perception of the contribution of immigrants to the U.S.

Each of these questions have responses that are collapsed into either a pro-immigrant response, an anti-immigrant response, or a neutral response. We scale each question to have responses that range from 0 to 1, where 0 is a pro-immigrant answer, 1 is an anti-immigrant

answer, and 0.5 is a neutral response. We then average across 16 questions to construct our anti-immigration views index. We use this anti-immigrant index outcome for Republicans. For Democrats, we construct a pro-immigrant index which is simply calculated as 1 minus the anti-immigrant index.

Our post-treatment questions also include several questions about the treatment itself. We ask participants who heard an immigration treatment speech how anti-immigrant or pro-immigrant they perceived the speech to be. We also ask participants which president they thought gave the original speech out of a choice of the four most recent presidents (and an “other” option). For individuals in the ambiguous source group, or those that received the actor versions of the speeches, this question solicits the participants’ best guess of the original source of the speech.²³

Throughout the survey, within blocks of questions, both the question order and possible response order is randomized, so that idiosyncratic features of the survey do not color participant response. The full survey instrument is available in Supplemental Materials Appendix S1.

4.2 Sample Restrictions

The initial recruitment criteria included: consenting to the survey, having audio capability, being a U.S. citizen, and declaring in the survey that they were affiliated with the same party they said they were affiliated with in the recruitment advertisement.²⁴ These initial criteria were pre-set in the survey design and failure resulted in automatic exit from the survey. Further, some individuals attempted the survey more than once, and we keep only the first attempts for these participants. Prior to treatment randomization, all participants listen to an audio clip of a weather forecast and then are asked attention questions about the topic of this weather clip. Participants who fail to answer the attention check correctly are also automatically exited from the survey, resulting in a sample of 14,356.²⁵

We apply four additional restrictions to increase the quality of survey responses. First, we geocode the locations of the IP addresses of survey-takers and keep only the respondents located in the U.S. Second, we remove individuals who took the survey exceptionally quickly or slowly, keeping those who completed the survey in 4 to 45 minutes. Third, we remove individuals

²³This question also serves as a non-binding attention check for individuals who received a speech with a revealed president source; and in fact, nearly all participants in these groups guess the president correctly in this case (Figure A1).

²⁴A sub-set of individuals state that they are Independents (or neither Republican or Democrat) after starting the survey; we keep these individuals in the primary sample and exclude them in a robustness check.

²⁵We play an audio segment of a weather forecast from a local San Diego TV News segment for all participants, that describes a sunny and warm weekend forecast. We then ask participants multiple choice questions about the topic discussed in the clip (weather, traffic, sports, crime story, stock market news, other) and then the type of weather that was discussed. If participants fail to identify the correct topic(s) of the clip, we exclude them from the study. The inclusion of the weather attention clip was part of the pre-analysis plan for this project.

who did not answer all of the pre-treatment demographic questions so that we can control for pre-treatment variables for all respondents. We also remove a small number of individuals who listened to an audio treatment clip and did not answer a comprehension question about the clip correctly (297 unique individuals). Across these four restrictions, we drop 1,257 or 8.7% of respondents in the 14,356 person baseline sample. Figure A10 shows that these restrictions do not materially affect the pattern of results, though they do increase precision given that they increase data quality.

The final sample contains 13,099 individuals, of which 7,125 are Democrats and 5,974 are Republicans. This translates to treatment group sizes of approximately 650 individuals in the Democrat sample and 545 individuals in the Republican sample. Figure A9 depicts the sample restrictions that were used to arrive at the final sample and Appendix A4 discusses each of these restrictions in greater detail.

4.3 Sample Characteristics and Balance

Political party affiliation is an increasingly important group identifier in the U.S., where the political climate is highly polarized. Given this environment, we deliberately constructed our experiment to separately measure effects by party affiliation. Figure 2 plots the distribution of the key outcome, the anti-immigration beliefs index, for the control group that received no audio treatment. Republicans clearly hold views that are more anti-immigrant than Democrats in our sample, with the mass of each party separated along the distribution of this outcome.

Table 1 further summarizes the outcome index and each of the question components for the control group which was not assigned to any audio speech treatment. There is notable variation in the views across the component questions. This pattern highlights the strength of our survey in capturing multi-dimensional views on immigration. We incorporate the views across questions by constructing an average index; however, we also explore patterns across questions in Appendix A6.

Figure A3 shows that our sample is quite similar to the make-up of the national Democratic and Republican parties. One of our pre-treatment questions is lifted from the national Gallup survey on immigration attitudes and asks whether participants think that immigration levels should be increased, decreased, or be kept at present levels. Panel A and B show that our sample is similar to but somewhat more moderate than the national Gallup survey data for each party, with slightly larger share of Republicans and Democrats who feel that immigration should be kept at present levels. Panels C and D show that the study sample is quite similar to the national parties in demographics, though the sample is slightly more likely to be young, white, female, and more educated. These differences may partly reflect the fact that the study pool was recruited for an online survey, which tends to attract individuals who have greater access and comfortability

with technology and may also be more likely to be young and educated.

The summary statistics for both samples are shown in Table 2. As noted above, the Republican and Democrat samples are designed to reflect the demographics of their respective national parties. Likewise, the Republican sample has a higher share of respondents who are white, older, and from the southern U.S., while a smaller fraction are college educated. 36% of the Republican sample regularly watches FOX News, compared to 15% of the Democrat sample. The parties are predictably stratified in their support of Trump and Obama; 92% (8.5%) of Republicans (Democrats) voted for Trump in 2016, while 78% (6%) of Republicans (Democrats) are fans of Donald Trump. Similarly, 89% of Democrats are fans of Obama, as compared to only 17% of Republicans. Our sample has a high voter participation and is likely engaged in politics, as over 80% of both sample groups voted in the 2016 election. 16% of Republicans view immigration as their top political issue, compared to 5% of Democrats. Consistent with anticipated party views, 45% of Republicans believe that immigration should be decreased, compared with only 14% of Democrats.²⁶

Table 2 also shows that the study sample is balanced across treatment groups. For each demographic characteristic and party sample, we regress the characteristic on indicators for the 11 treatment arms in the study and calculate the joint significance of these indicators. Successful randomization will be associated with a lack of joint significance for these treatment indicators. Nearly all tests pass balance and do not show statistical significance. Only 4 of 62 tests are significant at the 10% level and 3 at the 5% level, similar to the number we should expect to fail due to chance alone (≈ 6 and 3 tests, respectively).

4.4 Representativeness of Experiment

Lastly, we highlight ways in which our experiment setting and structure may offer insights into the broader context of understanding how political signals may meaningfully change beliefs in the real world.²⁷ First, by design, our study sample is well matched to the demographics of the Democratic and Republican parties in the United States, which provides reassurance that the effects we observe are representative of typical partisan voters. Second, as we discuss in Section 4.2 and Appendix A4, the sample was constructed to ensure high quality responses, but the applied restrictions did not substantively impact the pattern of the results, and the balance tests in the final sample provide evidence of successful randomization (Section 4.3). Third, we deliberately designed the treatment interventions to mimic natural or real-world political information by using audio speeches from actual presidents, rather than using artificial researcher-designed text statements embedded in a survey. Due to resource limitations, we study changes in respondent beliefs within our online survey, which occurs over a short time

²⁶These views are taken from the single pre-treatment question on immigration views asked of all participants.

²⁷This discussion is guided by the framework and discussion of external validity in List (2020).

horizon. In the real world, voters are inundated with a barrage of political information across multiple media, and the impact of any single political signal may have varying persistence. While we cannot address the persistence of changes in beliefs in our study, we note that understanding the mechanics of how beliefs respond to political messages and sources is foundational to understanding the longer-term dynamics of belief formation and voter behavior.

5 Results

5.1 Total Effect

Our first set of results relate to the total impact of partisan statements on immigration beliefs. Table 3 displays the differences between treatment groups that heard a president version of an anti-immigrant or pro-immigrant speech relative to the control group that did not hear any speech. Our results are consistent with findings from other work; partisan statements from political leaders do move beliefs. We measure immigration beliefs in the direction of each party's policy leaning. For Republicans we measure outcomes as anti-immigration beliefs, $\ln(P(Anti) + 1)$, and for Democrats to pro-immigration beliefs, $\ln(P(Pro) + 1) = \ln((1 - P(Anti)) + 1)$. To clearly illustrate the magnitude of all changes, we include a column "% Diff" in all tables which shows the implied percent change in the *untransformed probability* anti-immigrant (or pro-immigrant) rather than the log-transformed outcome.

Overall, these results show that participants from both parties update in the *direction* of the message that they hear. A pro-immigrant message increases pro-immigrant beliefs, while an anti-immigrant message increases anti-immigrant beliefs. Seven of the eight total message-by-party comparisons have a significant impact on updating beliefs. The only case without a significant effect is the Trump pro-immigrant message for Democrats. We find that movements *away from the party prior* appear to be larger and more significant; pro-immigrant effects are larger for Republicans and anti-immigrant effects are larger for Democrats, and we find effect sizes are symmetric across parties. For Republicans (Democrats), pro-immigrant (anti-immigrant) messages decrease their subjective probability that immigration is unfavorable (favorable) by 2-8%. In addition to estimating differences in the average effect across groups, we find similarly significant differences in the distributions of the immigration index across treatment groups, using Kolmogorov-Smirnov tests of equality.

5.2 Persuasion Decomposition: Fixing the Message

Our central decomposition separates the total effect of the partisan statements between the effect of an anonymous message (β_m) and the source persuasion effect (β_{ms}). In Table 4 and Figure 4, we estimate that in nearly all cases, the effect of the anonymous message β_m , drives a substantial portion of the total effect. Again, the results are strikingly parallel across parties. Both

Republicans and Democrats update their priors on immigration in the direction of a message from an anonymous source. Here again, seven of the eight total message-by-party comparisons yield significant effects, changing beliefs by 2-5%.

The findings for the importance of message content, given by β_m , are notable for two reasons. First, the rhetoric used in the messages for this experiment is emotional and political, as the speech statements focus on immigration views and general policy proposals and do not include any factual information or substantive content about immigration. Second, messages appear to move individuals more when the messages oppose an individual's prior; Democrats are more moved by anti-immigration messages and Republicans are moved more by pro-immigrant messages. Collectively, these results imply that individuals are capable of updating beliefs when presented with new viewpoints, even when these arguments are not supported by rigorous facts or documentation. These findings suggest that increasing partisan polarization may be partly caused by a lack of exposure to opposition views.

Turning to the role of source persuasion, which captures the added effect of a specific source delivering a particular message, we find that this channel is only important in one symmetric case: when a party leader delivers a message against party lines to members of his party. That is, Obama (Trump) has a persuasion effect only for Democrats (Republicans) when giving a speech that is anti-immigrant (pro-immigrant). In these cases, source persuasion effects explain a substantial share of the total updating of participants: 60% of the total 6.9% change in beliefs for Republicans, and 44% of the total 7.6% change in beliefs for Democrats (Figure 6). Figure A4 plots the sample group outcomes for the cases with source persuasion effects, and shows that the entire distribution shifts as a result of the leader source delivering these statements in both symmetrical cases.

The source persuasion findings underscore the importance of particular messengers delivering partisan statements. Our results show that party members will *follow their leader* and update their policy views when presented with new and surprising information from a leader that they support. The findings show that leaders have the power to affect the beliefs of followers, and that party affiliates do not simply support leaders when they mirror their pre-existing beliefs.

5.3 Mechanisms for the Source Persuasion Effect

The conceptual framework predicts that the strength of leader persuasion will be a function of two factors, the surprise of the message coming from a particular source and the subjective reliability of the source. Surprise of the message is the participant's belief that the actual source was unlikely given the message, or whether the participant *does not expect* that a particular source would deliver a message of a certain type. Specifically, this is characterized as 1 minus the perceived probability of the actual or true source conditional on the message, or $(1 - P(s|m))$.

Ex ante, it is not clear that participants will find any message to be surprising, as it could be the case that participants think that all politicians can easily change positions and any message type from any leader is equally likely.

In our setting, we can examine actual data on the surprise of the message using participant guesses about the identity of the true source when they heard an actor message. Panel B of Figure 3 shows our measure of message surprise: the share of respondents in the actor groups who guessed the true president *incorrectly* when asked which president they thought gave the speech (out of four recent presidents and an “other” option).²⁸ The plot shows that the most surprising messages are those where a leader delivers a message that opposes their reputation on immigration. The anti-immigrant Obama speech and the pro-immigrant Trump speech have more surprising sources from the perspective of participants in the study.

The conceptual framework illustrates that the second factor that drives leader persuasion is the subjective reliability of the source. Again, prior to conducting this experiment, it is not clear that participants will view some leaders as more reliable messengers of information than others. If individuals are persuaded solely by the information in messages, it could be the case that they prefer particular leaders only because those leaders have platforms that reflect what the individuals already believe.

We are not able to directly measure subjective reliability in our data, but we are able to illustrate differences in the favorability of leaders across parties which are likely correlated with reliability. Panel C of Figure 3 summarizes pre-treatment support for presidents across parties. Republicans overwhelmingly voted for Trump in 2016, with a share of 91.5%, in contrast to only 8.5% of Democrats. Similarly, when participants are asked whether they are fans of Obama and Trump, the vast majority of Republicans state that they are fans of Trump (78%) and a minority are fans of Obama (17%). A similar story is present for Democrats, with 89% stating they are fans of Obama, versus only 6% that are fans of Trump. This pattern suggests that Republicans may find Trump to be more subjectively reliable than Obama, and vice versa for Democrats.

The source persuasion effects that we find are consistent with the predicted importance of these two factors in the conceptual framework. The persuasion effect is a joint or multiplicative function of surprise and subjective reliability; this implies that a statement must be both surprising (coming from the particular source) and the source must be perceived as reliable for an effect to be present. We find this pattern in our results: participants are only persuaded by counter-reputational messages from a leader that they support, i.e. Trump’s pro-immigrant speech for Republicans and Obama’s anti-immigrant speech for Democrats. We see null effects in source persuasion when a statement is surprising, or counter-reputational, but is delivered by

²⁸We replicate this figure to show the share of respondents who guess the true political party of the president speaker in Figure A2 for illustrative purposes.

the opposition leader, who may be perceived as less reliable, i.e. Trump’s pro-immigrant speech for *Democrats* or Obama’s anti-immigrant speech for *Republicans*. Likewise, while it is likely that each party finds its own leader to be reliable, statements that are consistent with leader reputation are unsurprising and leaders do not persuade in this case, i.e. Trump’s *anti-immigrant* speech for Republicans and Obama’s *pro-immigrant* speech for Democrats.

5.4 Alternative Decomposition Test for Priming Effect: Fixing the Source

When designing our experiment, we wanted to explicitly investigate whether any source effects in this study could be attributable to priming, rather than persuasion. We use an alternative decomposition which fixes the president source to test for the presence of such effects.

Table 5 and Figure 5 show the results that decompose the impact of a political statement between source priming (β_s) and an identified message (β_{sm}). We find there is no priming effect in any of the treatment groups. This null finding is present in both the average differences (coefficient significance) and the Kolmogorov-Smirnov tests of distributional equality. Neither party changes their immigration views after listening to a non-ideological (turkey pardon) message from either leader, when compared to the control group that does not hear any speech. This means that, for both parties, the bulk of the updating we observe from the total effect loads on the coefficient of the identified message (β_{sm}). These effects conform with the direction the message, and are significant for all messages that oppose party priors, as well as for Obama pro-immigrant message for Democrats. Figure 7 plots decomposition between these two channels. The light blue and red bars, which represent the identified message effect (β_{sm}), drive both the direction and the size of the total effect, whereas the darker bars, which represent the source priming effect are always negligible and insignificant.

5.5 Economic Significance and Implications for Polarization

Evaluating the effects of the intervention on political behaviors, is out of the scope of this paper; however, the mechanisms revealed in our results are likely to have important effects on political outcomes and policy choices. For example, anti-immigration views have been found to independently increase right-wing vote and subsequent policy changes (Lubbers et al., 2002; Arzheimer, 2009; Hooghe and Dassonneville, 2018). Furthermore, the magnitudes of our treatment effects are economically significant in terms of their implied changes in the distance between Democrats and Republicans. Leveraging the randomization of the experiment, we can explore how the distance between Republican’s and Democrat’s immigration beliefs, or partisan polarization, would be affected under different alternative scenarios.²⁹ In particular, we are

²⁹Throughout this paper, we use the term “polarization” to refer specifically to partisan/party sorting or partisan polarization, or the distance in beliefs between individuals of different political parties.

interested in how the anonymous message and source persuasion effects we identify in our central decomposition may translate to implications for polarization.

We calculate the change in polarization by estimating a regression using the anti-immigration index as the outcome that includes both parties in the sample. Each regression includes a treatment group for Democrats and Republicans relative to a corresponding control group for both parties, and records the change in distance between the two groups in the treatment groups versus the control group. Figure 8 and Table A1 display the results of this exercise. Given our findings that people move in the direction of the messages that they hear, we show that polarization increases when participants hear messages consistent with their priors and decreases when participants hear messages that oppose their priors.

First, we consider a counterfactual that is most likely to arise in the real world, the case when party leaders deliver statements to their respective followers which are consistent with the priors of their followers (Own Leader, With Prior). These messages increase polarization by 11%, and this effect is the result of the anonymous message effect, or message content. Statements consistent with party priors that come from the opposition leader also increase polarization, but these point estimates are smaller and not significant.

Partisan statements that are contrary to the priors of participants decrease polarization. The against-prior anonymous message effects from *either* the party leader or the opposition leader both decrease polarization by 12-14%. As discussed above, the source persuasion effect amplifies the impact of the against-prior messages only for party leaders, leading to an extra decrease in polarization of 16%, or a total decrease in polarization of $\approx 30\%$. In contrast, against-prior statements are not more persuasive when they come from the opposition leader, and here, the anonymous message effect dominates. In sum, these results imply that party leaders have the capacity to move their followers in a new alternative direction, and in this setting, yields the potential to reduce political polarization. Relative to other papers in the literature, the reduction in polarization we estimate are meaningful. Alesina et al. (2018), provide information treatments that correct misperceptions on the share of immigrants and this reduces the left-right gap by 6.5%, sharing stories of a hard-working immigrants reduces this gap by 11%. In their work this update was also accompanied by changes in preferences for redistribution policies.

5.6 Heterogeneity

Our experiment was designed to specifically address differences in reactions to treatment along political party lines, which we hypothesized would be the critical dimension of heterogeneity for our study. In line with our hypothesis, the results are symmetrical and mirrored by party, as party affiliation appears to be the key characteristic that predicts both prior beliefs about immigration as well as responses to different leaders.

We explore multiple additional dimensions of heterogeneity, including strength of party affiliation, strength of prior beliefs about immigration, degree of political engagement, news consumption habits, and additional demographic characteristics in Appendix A5. In sum, we find very limited variation in our findings across several dimensions of heterogeneity *within* party. This result could be attributable to a lack of power to detect these kinds of subgroup differences, as we did not design our experiment to discern these effects given resource constraints. At the same time, while the standard errors do moderately increase with the smaller sub-groups, the point estimates for each cut of the data are quite consistent in practice. This stability stands in contrast to prior work in political science using either motivated reasoning models, which posits that individuals with stronger party affiliation will have stronger effects, or dual processing models, which posits that individuals who are less informed about policy issues will have stronger effects (Bullock, 2020). Instead, it appears that the shifts in beliefs that we observe are present for all segments of the distribution within party. The pattern of findings supports the hypothesis that *party affiliation* is the most important factor that determines responsiveness to treatment.

6 Robustness

6.1 Multiple Hypothesis Testing

This study includes several treatment arms and multiple research hypotheses; thus, one concern could be whether some of our results might be the artifact of multiple hypothesis testing. Throughout the main body of this paper, we display our estimates separately by party affiliation, message type (pro- or anti-immigrant), and president, in order to make the exposition of our results as clear as possible to the reader. However, the effective number of research hypotheses in this study are fewer than the number of pairwise tests we present in our main tables, as the study hypotheses only leverage ideological alignment between participants and messages or participants and leaders. In practice, this means that we can group our study hypotheses into two dimensions: (1) whether a message is in line with or against the respondent’s party prior on immigration, and (2) whether the leader is a representative of the respondent’s party or the opposition party. Along each dimension, we measure whether and how the treatments induce changes in immigration beliefs. We use these effective hypotheses to examine the potential influence of multiple hypothesis testing on our results.

This structure produces six hypotheses for each decomposition in our paper, which we group into four testing families in Table A2. Table A2 displays the results of regressions that include both parties of respondents and assign treatment groups based on the underlying hypothesis of the study. We apply the multiple hypothesis adjustment developed in List et al. (2021), which uses a bootstrap algorithm to flexibly permit the use of researcher defined sets of hypotheses, adjustments for controls, and multiple outcomes within a family of tests. While the regressions

are altered in structure from our baseline analysis, the adjusted p-values for these estimates functionally illustrate which of our hypotheses are robust to multiple hypothesis testing. For the central Persuasion Decomposition, Panel A shows that both anonymous message effects, messages aligned with and against the party prior, are highly significant and move individuals toward the direction of the message. Panel B highlights the significance of the central persuasion result, that conditional on a message that is against the party prior, a leader source from the respondent's party will yield an additional persuasive effect. The results in the Alternative Decomposition for Priming also accord with our baseline estimates in Table 5; with the most notable result being that there continues to be no significant effect of leader priming on beliefs.

6.2 Specification Tests

In this section, we probe the robustness of the baseline findings for the central Persuasion Decomposition. Figures A5, A6 and A8 show that the results are stable across various specifications and sample restrictions. The first specification in Figure A5 drops the demographic covariates that are included in the baseline model using the double Lasso procedure (See Appendix A3) and shows very similar results. In the second specification, we adjust the sample to exclude observations collected on election day or later (November 3, 2020) to adjust for idiosyncratic features of the political environment at this time, and again find similar results.³⁰ In our baseline sample, we include the first attempts of the survey for individuals who attempted to take the survey multiple times. As a third test, we check that the results are consistent when dropping the first attempts of all duplicate responders, and again find very similar results.

Next, we vary the construction of the immigration index outcome. Some of the questions in the survey have 5 option responses (e.g. Strongly Agree, Agree, Disagree, Strongly Disagree, No Opinion) and others have 3 option responses (e.g. Yes, No, No Opinion). In our baseline index, we translate all questions into 3 options; a pro-immigrant answer, an anti-immigrant answer, or a neutral response, collapsing the variation in the 5 option questions. In the fourth test in Figure A5, we show the results when we consider the full variation from 5 option questions and find very similar results. Lastly, when we initially pre-registered the experiment, we specified an index that would use only 14 of the 16 post-treatment questions in the survey, because we believed that some of the questions might produce ambiguous responses, given that they related to potentially more favored immigrant populations. These two questions asked whether respondents thought that immigrants from English-speaking countries or immigrants living in the U.S. legally contribute positively or negatively to society. In practice, these questions did not produce anomalous results and thus we added them to our baseline index to increase its

³⁰We purposefully conducted our experiment during the period leading up to the national election in order to capture political beliefs and attitudes at this time, beginning on October 16, 2020. Due to a slower recruitment pace than anticipated, the study ran until November 10, 2020, and our baseline sample includes all survey dates.

information content. However, using only the 14 question version also produces similar results to the baseline estimates (final test in Figure A5).³¹

In Figure A6, we consider alternate formulations of the outcome. In our baseline model, we utilize the $\ln(P(y) + 1)$ transformation in our baseline regressions so as not to exclude individuals with purely anti-immigrant or pro-immigrant views, where $P(y) = 0$. As a first check, we estimate the model using a simpler outcome of $\ln(P(y))$ and find consistent results. As additional checks, we include a specification using the inverse hyperbolic sine of $P(y)$ and a simple linear regression using $P(y)$ as the dependent variable. In Figure A7, we standardize our baseline probability index using values in the “no audio” control group, and construct an inverse covariance weighted version of the standardized index following Anderson (2008). The results of each of these alternative specifications produces results that are similar in direction and significance to the baseline estimates, given that nearly all specifications are a monotonic transformation of the key outcome. Likewise, results using these alternative formulations do not statistically differ from the estimates using an analogous baseline specification. However, the scale of the values differs in these models as a result of the difference of arguments in the expressions.

Our baseline regressions consider the log of an immigration index that has a continuous support of $[0,1]$, given that it is an average of multiple questions each of which can assume values of 0, 0.5, or 1. This construction has the benefit of incorporating gradations in beliefs about immigration for each participant. However, it is also useful to consider alternative specifications that discretize this outcome and employ logit or probit models (discretized using a value of 0.5). Figure A8 finds results that are quite similar in direction and significance to the baseline estimates using a logit and probit model, as well as when using a linear probability model.

6.3 Results By Question

Our central results rely on an index composed of 16 different questions to incorporate multiple dimensions of immigration attitudes. In Appendix A6, we test the robustness of the index to its components and explore effects for individual questions and sub-sets of questions. First, we find that the results are robust to re-calculating the index leaving out one question at a time in Figure A15 and are broadly similar (though less precise) when estimating the models separately for each question in Figure A16. We likewise find similar effects when we split our index into two sub-indices of (1) questions focused on immigration policy and (2) questions focused on immigration attitudes in Figure A19. When we explore effects for individual questions we find a suggestive pattern whereby Republicans move most on questions that involve modest policy

³¹In our initial survey design, we also attempted to measure two out-of-sample outcomes, voting decisions in the 2020 election and interest in donating to either an anti- or pro-immigrant charity. Unfortunately, we found limited variation in these outcomes in practice; due in part to the fact that the sample collection period overlapped with the election (See Appendix A7 for additional detail).

changes for individuals who may be viewed as “model” immigrants, and Democrats move most on questions that relate to penalizing immigrants who may be viewed as the most egregious “rule-breakers.”

Overall, the “by question” results underscore a crucial methodological strength of the study: the fact that we are able to aggregate responses of multiple dimensions of immigration attitudes into a summary index. Immigration is a complex policy topic, and partisans may change their beliefs on some aspects of immigration and not others. Additionally, using information from multiple distinct questions increases the total precision of our outcome. Had we taken the approach of earlier work and considered only 1-2 outcome questions with a general focus, we very well could have missed the robust and symmetric effects that we observe for both parties in our study.

7 Discussion

How do leaders change the beliefs of their constituents? We find evidence that leaders can persuade individuals as messengers of information. While we do not find that participants change their beliefs through simple exposure to the identity of a leader (or source priming), we do find that certain messages are more persuasive when voiced by particular leaders. These persuasion effects are measured as the effect of a leader source, holding fixed the substance of a message. We find that a leader is most persuasive when expressing statements that are unexpected to an audience who finds the leader to be credible. Specifically, President Obama is most persuasive when voicing an anti-immigrant speech to Democrats and President Trump is most persuasive when voicing a pro-immigrant speech to Republicans.

This pattern of results illustrates that supporters will *follow their leader* to new political positions. In our context, these effects lead to a reduction in the distance in beliefs about immigration across parties, or partisan polarization, suggesting that leaders have the potential to play an important role in reducing partisan division.

While we find that leader persuasion reduces polarization, the implications of our results are more general. Our framework implies that surprising or new positions for leaders will be most persuasive to followers, which may or may not increase longer-term polarization depending on the circumstance. These impacts also have the potential to be meaningful in situations where a new issue arises and expectations about policy positions have yet to be set, such as during the early stages of a public health crisis like the recent COVID pandemic (Larsen et al., 2022). A more abstract implication of our general findings is that a leader has the potential to create a “cult of personality” and use his or her persuasive power to shape policy decisions.

Our work credibly and comprehensively isolates the determinants of the leader influence, focusing on the mechanisms that drive changes in beliefs. Future research should continue

to dissect the ways in which prominent partisan figures may shape, alter, or disrupt public opinion; as well as the persistence of a partisan leader's influence over time. Understanding these dynamics could provide insights into the strengths and fragilities of democratic governments, which rely on their ability to build consensus among voters and their elected officials.

References

- Abramitzky, R. and L. Boustan (2022). *Streets of gold: America's untold story of immigrant success*. Hachette UK.
- Adena, M., R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya (2015). Radio and the Rise of the Nazis in Prewar Germany. *The Quarterly Journal of Economics* 130(4), 1885–1939.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and Identity. *The Quarterly Journal of Economics* 115(3), 715–753.
- Alesina, A., A. Miano, and S. Stantcheva (2018). Immigration and Redistribution. NBER Working Paper.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association* 103(484), 1481–1495.
- Andreoni, J. and T. Mylovanov (2012). Diverging Opinions. *American Economic Journal: Microeconomics* 4(1), 209–32.
- Arzheimer, K. (2009). Contextual factors and the extreme right vote in western europe, 1980–2002. *American Journal of Political Science* 53(2), 259–275.
- Bakker, B. N., Y. Lelkes, and A. Malka (2020). Understanding Partisan Cue Receptivity: Tests of Predictions from the Bounded Rationality and Expressive Utility Perspectives. *The Journal of Politics* 82(3), 1061–1077.
- Barber, M. and J. C. Pope (2019). Does Party Trump Ideology? Disentangling Party and Ideology in America. *American Political Science Review* 113(1), 38–54.
- Baron, D. P. (2006). Persistent Media Bias. *Journal of Public Economics* 90(1-2), 1–36.
- Baum, M. A. and P. Gussin (2008). In the Eye of the Beholder: How Information Shortcuts Shape Individual Perceptions of Bias in the Media. *Quarterly Journal of Political Science* 3(1), 1–31.
- Baysan, C. (2021). Persistent Polarizing Effects of Persuasion: Experimental Evidence from Turkey. *American Economic Review*. Forthcoming.
- Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009). Powerful Women: Does Exposure Reduce Bias? *The Quarterly journal of economics* 124(4), 1497–1540.
- Beaman, L., E. Duflo, R. Pande, and P. Topalova (2012). Female Leadership Raises Aspirations and Educational Attainment for Girls: A Policy Experiment in India. *Science* 335(6068), 582–586.
- Benjamin, D. J., J. J. Choi, and G. Fisher (2016). Religious Identity and Economic Behavior. *Review of Economics and Statistics* 98(4), 617–637.
- Benoît, J.-P. and J. Dubra (2019). Apparent Bias: What Does Attitude Polarization Show? *International Economic Review* 60(4), 1675–1703.
- Berinsky, A. J. (2009). *In Time of War*. University of Chicago Press.
- Bidwell, K., K. Casey, and R. Glennerster (2020). Debates: Voting and Expenditure Responses to Political Communication. *Journal of Political Economy* 128(8), 2880–2924.
- Boudreau, C. and S. A. MacKenzie (2014). Informing the Electorate? How Party Cues and Policy Information Affect Public Opinion about Initiatives. *American Journal of Political Science* 58(1), 48–62.
- Broockman, D. E. and D. M. Butler (2017). the Causal Effects of Elite Position-Taking on Voter Attitudes: Field Experiments with Elite Communication. *American Journal of Political Science* 61(1), 208–221.
- Bullock, J. G. (2011). Elite Influence on Public Opinion in an Informed Electorate. *American Political Science Review* 105(3), 496–515.
- Bullock, J. G. (2020). Party cues. In *The Oxford Handbook of Electoral Persuasion*, pp. 129. Oxford University Press, USA.

- Bursztyn, L., G. Egorov, I. Haaland, A. Rao, and C. Roth (2022). Justifying Dissent. NBER Working Paper.
- Bursztyn, L., A. L. González, and D. Yanagizawa-Drott (2018). Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia. NBER Working Paper.
- Callen, M., M. Isaqzadeh, J. D. Long, and C. Sprenger (2014). Violence and Risk Preference: Experimental Evidence from Afghanistan. *American Economic Review* 104(1), 123–48.
- Chiang, C.-F. and B. Knight (2011). Media Bias and Influence: Evidence from Newspaper Endorsements. *The Review of Economic Studies* 78(3), 795–820.
- Chopra, F., I. Haaland, and C. Roth (2022). Do People Demand Fact-Checked News? Evidence from US Democrats. *Journal of Public Economics* 205.
- Cinelli, M., G. De Francisci Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini (2021). The Echo Chamber Effect on Social Media. *Proceedings of the National Academy of Sciences* 118(9).
- Cohen, G. L. (2003). Party over policy: The dominating impact of group influence on political beliefs. *Journal of personality and social psychology* 85(5), 808.
- Cohn, A., J. Engelmann, E. Fehr, and M. A. Maréchal (2015). Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review* 105(2), 860–885.
- Cohn, A. and M. A. Maréchal (2016). Priming in Economics. *Current Opinion in Psychology* 12, 17–21.
- Cohn, A., M. A. Maréchal, and T. Noll (2015). Bad Boys: How Criminal Identity Salience Affects Rule Violation. *The Review of Economic Studies* 82(4), 1289–1308.
- Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment. *Journal of Public Economics* 98, 100–112.
- DellaVigna, S. and E. Kaplan (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics* 122(3), 1187–1234.
- Doherty, C., J. Kiley, and N. Asheer (2020). In Changing U.S. Electorate, Race and Education Remain Stark Dividing Lines. *Pew Research Center*.
- Doherty, C., J. Kiley, N. Asheer, and C. Jordan (2020). Election 2020: Voters are Highly Engaged, but Nearly Half Expect Difficulties to Voting. *Pew Research Center*.
- Doherty, C., J. Kiley, and B. Johnson (2016). 2016 Campaign: Strong Interest, Widespread Dissatisfaction. *Pew Research Center*.
- Druckman, J. N. (2022). A framework for the study of persuasion. *Annual Review of Political Science* 25, 65–88.
- Druckman, J. N., E. Peterson, and R. Slothuus (2013). How Elite Partisan Polarization Affects Public Opinion Formation. *American Political Science Review* 107(1), 57–79.
- Durante, R., P. Pinotti, and A. Tesei (2019). The Political Legacy of Entertainment TV. *American Economic Review* 109(7), 2497–2530.
- Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and Political Persuasion: Evidence from Russia. *American Economic Review* 101(7), 3253–85.
- File, T. (2018). Characteristics of Voters in the Presidential Election of 2016. *U.S. Census*.
- Fryer Jr, R. G., P. Harms, and M. O. Jackson (2019). Updating Beliefs When Evidence Is Open to Interpretation: Implications for Bias and Polarization. *Journal of the European Economic Association* 17(5), 1470–1501.
- Fujiwara, T. and L. Wantchekon (2013). Can an Informed Public Deliberation Overcome Clientelism? Experimental Evidence from Benin. *American Economic Journal: Applied Economics* 5(4), 241–55.
- Gentzkow, M. and J. M. Shapiro (2006). Media Bias and Reputation. *Journal of Political Economy* 114(2), 280–316.

- Gentzkow, M., M. B. Wong, and A. T. Zhang (2018). Ideological Bias and Trust in Information Sources. *Working Paper*.
- Gerber, A. S., D. Karlan, and D. Bergan (2009). Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions. *American Economic Journal: Applied Economics* 1(2), 35–52.
- Gilens, M. and N. Murakawa (2002). Elite Cues and Political Decision-Making. *Research in Micropolitics* 6, 15–49.
- Goren, P., C. M. Federico, and M. C. Kittilson (2009). Source Cues, Partisan Identities, and Political Value Expression. *American Journal of Political Science* 53(4), 805–820.
- Grigorieff, A., C. Roth, and D. Ubfal (2020). Does Information Change Attitudes toward Immigrants? *Demography* 57(3), 1117–1143.
- Haaland, I. and C. Roth (2020). Labor Market Concerns and Support for Immigration. *Journal of Public Economics* 191, 104256.
- Haaland, I., C. Roth, and J. Wohlfart (2020). Designing Information Provision Experiments. CESifo Working Paper.
- Hooghe, M. and R. Dassonneville (2018). Explaining the trump vote: The effect of racist resentment and anti-immigrant sentiments. *PS: Political Science & Politics* 51(3), 528–534.
- Hopkins, D. J., J. Sides, and J. Citrin (2019). The Muted Consequences of Correct Information about Immigration. *The Journal of Politics* 81(1), 315–320.
- Igielnik, R. and A. Budiman (2020). The Changing Racial and Ethnic Composition of the U.S. Electorate. *Pew Research Center*.
- Jones, J. (2019). New High in U.S. Say Immigration Most Important Problem. *Gallup*.
- Kalla, J. L. and D. E. Broockman (2018). the Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments. *American Political Science Review* 112(1), 148–166.
- Kam, C. D. (2005). Who toes the party line? cues, values, and individual differences. *Political Behavior* 27, 163–182.
- Larsen, B., T. J. Ryan, S. Greene, M. J. Hetherington, R. Maxwell, and S. Tadelis (2022). Counter-Stereotypical Messaging and Partisan Cues: Moving the Needle on Vaccines in a Polarized Us. NBER Working Paper.
- Leeper, T. J. and R. Slothuus (2014). Political parties, motivated reasoning, and public opinion formation. *Political Psychology* 35, 129–156.
- Lenz, G. S. (2013). *Follow the leader?: how voters respond to politicians' policies and performance*. University of Chicago Press.
- Levy, R. (2021). Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *American Economic Review* 111(3), 831–70.
- List, J. A. (2020). Non est disputandum de generalizability? a glimpse into the external validity trial. Technical report, National Bureau of Economic Research.
- List, J. A., A. M. Shaikh, A. Vayalinal, et al. (2021). Multiple testing with covariate adjustment in experimental economics. Technical report, The Field Experiments Website.
- Lubbers, M., M. Gijsberts, and P. Scheepers (2002). Extreme right-wing voting in western europe. *European journal of political research* 41(3), 345–378.
- Martin, G. J. and A. Yurukoglu (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review* 107(9), 2565–99.
- Merkley, E. and D. A. Stecula (2021). Party Cues in the News: Democratic Elites, Republican Backlash, and the Dynamics of Climate Skepticism. *British Journal of Political Science* 51(4), 1439–1456.

- Mullainathan, S. and A. Shleifer (2005). The Market for News. *American Economic Review* 95(4), 1031–1053.
- Nicholson, S. P. (2012). Polarizing Cues. *American Journal of Political Science* 56(1), 52–66.
- Ong, A. D. and D. J. Weiss (2000). The Impact of Anonymity on Responses to Sensitive Questions 1. *Journal of Applied Social Psychology* 30(8), 1691–1708.
- Song, L. (2021). The Heterogeneous Effects of Social Media Content on Racial Attitudes. Working Paper.
- Taber, C. S. and M. Lodge (2006). Motivated skepticism in the evaluation of political beliefs. *American journal of political science* 50(3), 755–769.
- Tappin, B. M., A. J. Berinsky, and D. G. Rand (2023). Partisans' receptivity to persuasive messaging is undiminished by countervailing party leader cues. *Nature Human Behaviour*, 1–15.
- Thaler, M. (2021). The Fake News Effect: Experimentally Identifying Motivated Reasoning Using Trust in News. Available at SSRN 3717381.
- Younis, M. (2020). Americans Want More, Not Less, Immigration for First Time. *Gallup*.
- Zaller, J. R. (1992). *The nature and origins of mass opinion*. Cambridge university press.

Table 1: Immigration Index

	Republicans		Democrats	
	<i>Probability</i>		<i>Probability</i>	
<i>Control Group (No Audio)</i>	<i>Anti-Immigrant</i>	<i>Mean</i>	<i>Pro-Immigrant</i>	<i>Mean</i>
		<i>S.D.</i>		<i>S.D.</i>
Immigration Index	0.577	(0.195)	0.729	(0.196)
Components of Index				
Overall View of Immigrants	0.371	(0.470)	0.827	(0.368)
Immigrant Crime Share	0.596	(0.420)	0.797	(0.349)
Immigrant Economic Impact	0.611	(0.411)	0.782	(0.350)
Ideal Level of Immigration	0.673	(0.344)	0.610	(0.332)
Contribution: Undocumented Immigrants	0.828	(0.299)	0.537	(0.386)
Contribution: Legal Immigrants	0.188	(0.325)	0.914	(0.230)
Contribution: English-speaking Immigrants	0.288	(0.311)	0.845	(0.253)
Contribution: Spanish-speaking Immigrants	0.445	(0.352)	0.791	(0.294)
Contribution: Dreamers	0.433	(0.358)	0.833	(0.277)
Path to Citizenship: Dreamers	0.227	(0.397)	0.895	(0.282)
Path to Citizenship: All Immigrants	0.397	(0.473)	0.859	(0.323)
Deport Immigrants with Crime Record	0.913	(0.259)	0.380	(0.452)
Deport All Immigrants	0.655	(0.442)	0.788	(0.383)
Check Worker Immigration Status	0.912	(0.258)	0.393	(0.451)
Border Patrol	0.875	(0.289)	0.604	(0.449)
Border Wall	0.821	(0.356)	0.800	(0.368)
N	565		669	

Notes: This table summarizes the components of the immigration index outcome for the "no audio" control group. 16 questions are used to construct the index and their topics are listed in the table. Each question typically has 5 option responses (e.g. whether a respondent strongly agrees, agrees, disagrees, strongly disagrees or has no opinion about a statement). Each response is coded as either "pro-immigrant" (e.g. groups strongly agree and agree for a particular statement), "anti-immigrant" (e.g. groups strongly disagree and disagree for a particular statement), or neutral (e.g. "no opinion" for a particular statement). We construct the probability anti-immigrant for each question as "1" for an anti-immigrant response, "0" for a pro-immigrant response, and "0.5" for a neutral response. The complete survey and question wording is available in the Supplemental Materials for this paper. The total anti-immigrant index is the average across all question responses. In regressions using Democrat respondents we use the outcome of a "pro-immigrant index"; the pro-immigrant index is simply 1 minus the "anti-immigrant" index, or $P(\text{Pro-Immigrant}) = 1 - P(\text{Anti-Immigrant})$.

Table 2: Summary Statistics and Balance Tests

	Republicans				Democrats			
	Mean	S.D.	F-Test	P-Value	Mean	S.D.	F-Test	P-Value
Female	0.469	(0.499)	0.901	0.531	0.512	(0.500)	0.538	0.865
White	0.892	(0.311)	0.646	0.775	0.646	(0.478)	0.711	0.715
Black	0.030	(0.170)	0.745	0.682	0.188	(0.391)	1.096	0.361
Hispanic	0.040	(0.195)	1.228	0.267	0.084	(0.277)	0.574	0.837
Asian	0.029	(0.167)	0.409	0.943	0.065	(0.247)	0.551	0.854
18-24 years	0.081	(0.272)	0.633	0.786	0.153	(0.360)	0.918	0.515
25-34 years	0.184	(0.388)	1.046	0.401	0.256	(0.436)	1.226	0.268
35-44 years	0.220	(0.414)	0.932	0.502	0.219	(0.414)	0.862	0.569
45-64 years	0.324	(0.468)	1.105	0.354	0.234	(0.424)	0.800	0.629
65+ years	0.191	(0.393)	0.205	0.996	0.137	(0.344)	0.497	0.893
Northeast	0.166	(0.372)	0.657	0.765	0.219	(0.414)	0.523	0.875
Midwest	0.241	(0.427)	0.615	0.803	0.212	(0.409)	1.508	0.130
South	0.432	(0.495)	0.862	0.569	0.365	(0.481)	2.093	0.022
West	0.162	(0.368)	0.666	0.757	0.204	(0.403)	1.065	0.386
College or More	0.461	(0.499)	0.612	0.805	0.519	(0.500)	1.010	0.432
Full-time Employed	0.425	(0.494)	1.341	0.202	0.446	(0.497)	0.934	0.500
News (Weekly+): Facebook	0.398	(0.490)	0.963	0.473	0.386	(0.487)	0.553	0.853
News (Weekly+): Twitter	0.176	(0.381)	0.730	0.697	0.262	(0.440)	1.476	0.141
News (Weekly+): TV	0.585	(0.493)	0.860	0.571	0.602	(0.490)	0.751	0.677
News (Weekly+): Newspaper	0.250	(0.433)	2.039	0.026	0.319	(0.466)	1.214	0.276
News (Weekly+): FOX News	0.362	(0.481)	1.154	0.317	0.148	(0.355)	0.690	0.735
News (Weekly+): MSNBC	0.083	(0.276)	0.911	0.522	0.192	(0.394)	0.732	0.695
Party: Independent	0.076	(0.266)	1.606	0.098	0.053	(0.224)	2.178	0.016
Polarized (Estimated)	0.242	(0.428)	1.085	0.369	0.357	(0.479)	1.201	0.285
Voted (2016)	0.842	(0.365)	0.722	0.705	0.836	(0.371)	0.751	0.676
Voted for Trump (2016)	0.916	(0.278)	1.363	0.191	0.085	(0.280)	0.568	0.842
Fan of Trump	0.780	(0.414)	1.502	0.132	0.060	(0.237)	1.072	0.380
Fan of Obama	0.169	(0.375)	0.903	0.529	0.889	(0.314)	0.770	0.658
Immigration: Top Issue	0.156	(0.363)	1.104	0.355	0.051	(0.220)	1.050	0.398
Immigration: Should Increase	0.114	(0.318)	1.026	0.418	0.364	(0.481)	0.744	0.684
Immigration: Should Decrease	0.451	(0.498)	0.894	0.537	0.142	(0.349)	1.467	0.145
N	5974				7125			

Notes: Table summarizes demographic characteristics of sample from questions asked prior to intervention. Balance tests use separate regressions of each demographic characteristic on the full set of treatment group indicators, F-tests refer to joint significance of treatment assignment. “Party: Independent” refers to individuals recruited to the survey as Democrats or Republicans who indicate within the survey that they are Independents. “Polarized (Estimated)” is a probability of being in the top 25th percentile of anti-immigrant (pro-immigrant) views for Republicans (Democrats) constructed in the following way: predict this outcome in the post-treatment immigration index in the no audio control group using only pre-treatment characteristics and then use the coefficients to predict this outcome for the full sample. “Fan of Obama/Trump” comes from a question where we ask whether participants are a fan of these presidents. Prior to treatment, we ask participants to identify the top issue relevant to their vote (“Immigration: Top Issue”) and whether they think immigration should be increased, decreased or remain constant (“Immigration: Should Increase (Decrease)”).

Table 3: **Total Effect:** Combined Message and Source

Republicans						
Message	Source	β_t	(S.E.)	% Diff.	$P(Dist.)$	N
Anti	Trump	0.012**	(0.006)	3.37%	0.062	1120
Anti	Obama	0.016***	(0.006)	4.38%	0.002	1079
Pro	Trump	-0.025***	(0.006)	-6.85%	0.000	1126
Pro	Obama	-0.013**	(0.006)	-3.56%	0.083	1101
Control Group Mean: $\ln(P(Anti)+1)= 0.448$, $P(Anti)= 0.577$						
Democrats						
Message	Source	β_t	(S.E.)	% Diff.	$P(Dist.)$	N
Anti	Trump	-0.010*	(0.005)	-2.30%	0.180	1288
Anti	Obama	-0.032***	(0.005)	-7.61%	0.000	1303
Pro	Trump	-0.003	(0.005)	-0.70%	0.247	1321
Pro	Obama	0.008*	(0.005)	1.98%	0.003	1357
Control Group Mean: $\ln(P(Pro)+1)= 0.540$, $P(Pro)= 0.729$						

Notes: *p<0.1, **p<0.05, ***p<0.01. This table presents the estimates of the different president speech treatments, relative to the "no audio" control group. The relevant regression is $\ln(P(y) + 1) = \beta_0 + \beta_1 1[Message_i = 1] \times 1[Source_i = 1] + \gamma X_i + \epsilon$. The effects shown comprise the total combined impact of a president message, including both source and message effects. Outcomes are constructed as $\ln(P(y) + 1)$, where the outcome is the index probability anti-immigrant (pro-immigrant) for Republicans (Democrats). All regressions additionally include covariates selected by the double lasso procedure to improve precision (See Appendix A3). "% Diff." calculates the implied change in outcome probability, or the untransformed probability index, due to treatment, relative to the mean for the control group. " $P(Dist.)$ " is the p-value from a Kolmogorov-Smirnov test of equality of the distributions of the control and treatment groups.

Table 4: **Persuasion Decomposition:** Impact of Source within Fixed Message

Republicans										
ln(Probability Anti-Immigrant + 1)										
Message	Source	β_m	Anonymous Message			β_{ms}	Source Persuasion			N
			(S.E.)	%Diff.	P(Dist.)		(S.E.)	%Diff.	P(Dist.)	
Anti	Trump	0.012**	(0.006)	3.26%	0.049	0.001	(0.006)	0.33%	0.583	1648
Anti	Obama	0.010*	(0.006)	2.85%	0.108	0.006	(0.006)	1.65%	0.529	1618
Pro	Trump	-0.010*	(0.006)	-2.74%	0.042	-0.015**	(0.006)	-4.07%	0.067	1672
Pro	Obama	-0.019***	(0.006)	-5.09%	0.019	0.006	(0.006)	1.80%	0.320	1688
Control Group Mean: ln(P(Anti)+I)= 0.448, P(Anti)= 0.577										
Democrats										
ln(Probability Pro-Immigrant + 1)										
Message	Source	β_m	Anonymous Message			β_{ms}	Source Persuasion			N
			(S.E.)	%Diff.	P(Dist.)		(S.E.)	%Diff.	P(Dist.)	
Anti	Trump	-0.015***	(0.005)	-3.66%	0.022	0.005	(0.005)	1.31%	0.385	1913
Anti	Obama	-0.018***	(0.005)	-4.18%	0.002	-0.014***	(0.005)	-3.41%	0.000	1912
Pro	Trump	0.003	(0.005)	0.77%	0.238	-0.006	(0.005)	-1.46%	0.329	1972
Pro	Obama	0.012**	(0.005)	2.84%	0.058	-0.004	(0.005)	-0.98%	0.100	2021
Control Group Mean: ln(P(Pro)+I)= 0.540, P(Pro)= 0.729										

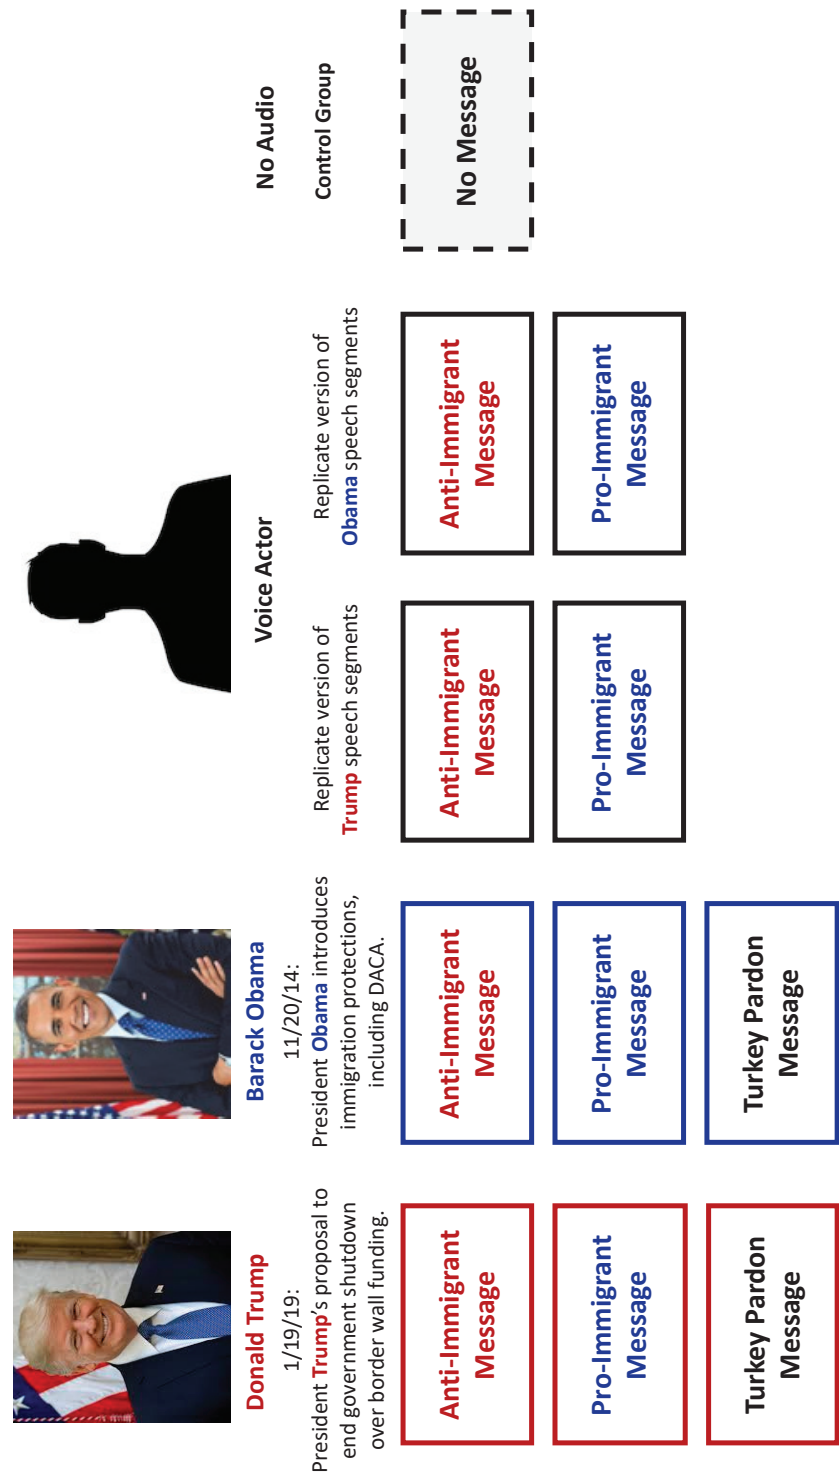
Notes: *p<0.1, **p<0.05, ***p<0.01. This table presents the estimates of the second decomposition of message and source effects of a speech treatment. This decomposition fixes the message of a speech, and varies the source between the actor and the president. Each regression includes a president immigration speech, the voice actor version of the speech, and the "no audio" control group. The relevant regression is $ln(P(y) + 1) = \beta_0 + \beta_m 1[Message_i = 1] + \beta_{ms} 1[Message_i = 1] \times 1[Source_i = 1] + \gamma X_i + \epsilon$, where the first regressor, $1[Message_i = 1]$, includes both treatment groups with a particular message, or the actor and president speech versions, and the second regressor, $1[Message_i = 1] \times 1[Source_i = 1]$, includes only the president's version of the immigration speech. Outcomes are constructed as $ln(P(y) + 1)$, where the outcome is the index probability anti-immigrant (pro-immigrant) for Republicans (Democrats). All regressions additionally include covariates selected by the double lasso procedure to improve precision (See Appendix A3). "% Diff." calculates the implied change in outcome *probability*, or the untransformed probability index, due to treatment, relative to the mean of the comparison group; where the comparison group is the "no audio" group for β_m and the actor speech group for β_{ms} . The p-values for the "% Diff." calculations correspond to the relevant coefficient estimates. "*P(Dist.)*" is the p-value from a Kolmogorov-Smirnov test of equality of the distributions of the treatment and comparison groups.

Table 5: **Alternative Decomposition for Source Priming:** Impact of Message within Fixed Source

Republicans									
<i>ln(Probability Anti-Immigrant + 1)</i>									
Message	Source	β_s	Source Priming		$P(Dist.)$	β_{sm}	Identified Message		N
			(S.E.)	%Diff.			(S.E.)	%Diff.	
Anti	Trump	0.008	(0.006)	2.31%	0.400	0.004	(0.006)	1.13%	1655
	Obama	0.008	(0.006)	2.29%	0.148	0.008	(0.006)	2.15%	1587
Pro	Trump	0.008	(0.006)	2.16%	0.400	-0.032***	(0.006)	-8.86%	1661
	Obama	0.008	(0.006)	2.23%	0.148	-0.021***	(0.006)	-5.65%	1609
Control Group Mean: $ln(P(Anti)+1)=0.448$, $P(Anti)=0.577$									
Democrats									
<i>ln(Probability Pro-Immigrant + 1)</i>									
Message	Source	β_s	Source Priming		$P(Dist.)$	β_{sm}	Identified Message		N
			(S.E.)	%Diff.			(S.E.)	%Diff.	
Anti	Trump	-0.001	(0.005)	-0.21%	0.759	-0.009*	(0.005)	-2.02%	1974
	Obama	-0.001	(0.005)	-0.15%	0.716	-0.031***	(0.005)	-7.27%	1931
Pro	Trump	-0.001	(0.005)	-0.28%	0.759	-0.002	(0.005)	-0.37%	2007
	Obama	-0.001	(0.005)	-0.17%	0.716	0.009*	(0.005)	2.10%	1985
Control Group Mean: $ln(P(Pro)+1)=0.540$, $P(Pro)=0.729$									


Notes: * $p<0.1$, ** $p<0.05$, *** $p<0.01$. This table presents the estimates of the first decomposition of message and source effects of a speech treatment. This decomposition fixes the president source as either Trump or Obama's voice, and varies the message between the turkey pardon and an immigration speech. Each regression includes a president immigration speech, a president turkey pardon speech, and the "no audio" control group. The relevant regression is $ln(P(y)+1) = \beta_0 + \beta_s 1[Source_i = 1] + \beta_{sm} 1[Message_i = 1] + \gamma X_i + \epsilon$, where the first regressor, $1[Source_i = 1]$, includes both treatment groups with a president's voice, or the turkey pardon and the immigration speech, and the second regressor, $1[Message_i = 1] \times 1[Source_i = 1]$, includes only the president's immigration speech. Outcomes are constructed as $ln(P(y)+1)$, where the outcome is the index probability anti-immigrant (pro-immigrant) for Republicans (Democrats). All regressions additionally include covariates selected by the double lasso procedure to improve precision (See Appendix A3). "% Diff." calculates the implied change in outcome *probability*, or the untransformed probability index, due to treatment, relative to the mean of the comparison group; where the comparison group is the "no audio" group for β_s and the turkey pardon group for β_{sm} . The p-values for the "% Diff." calculations correspond to the relevant coefficient estimates. " $P(Dist.)$ " is the p-value from a Kolmogorov-Smirnov test of equality of the distributions of the treatment and comparison groups.

Figure 1: Experiment Design



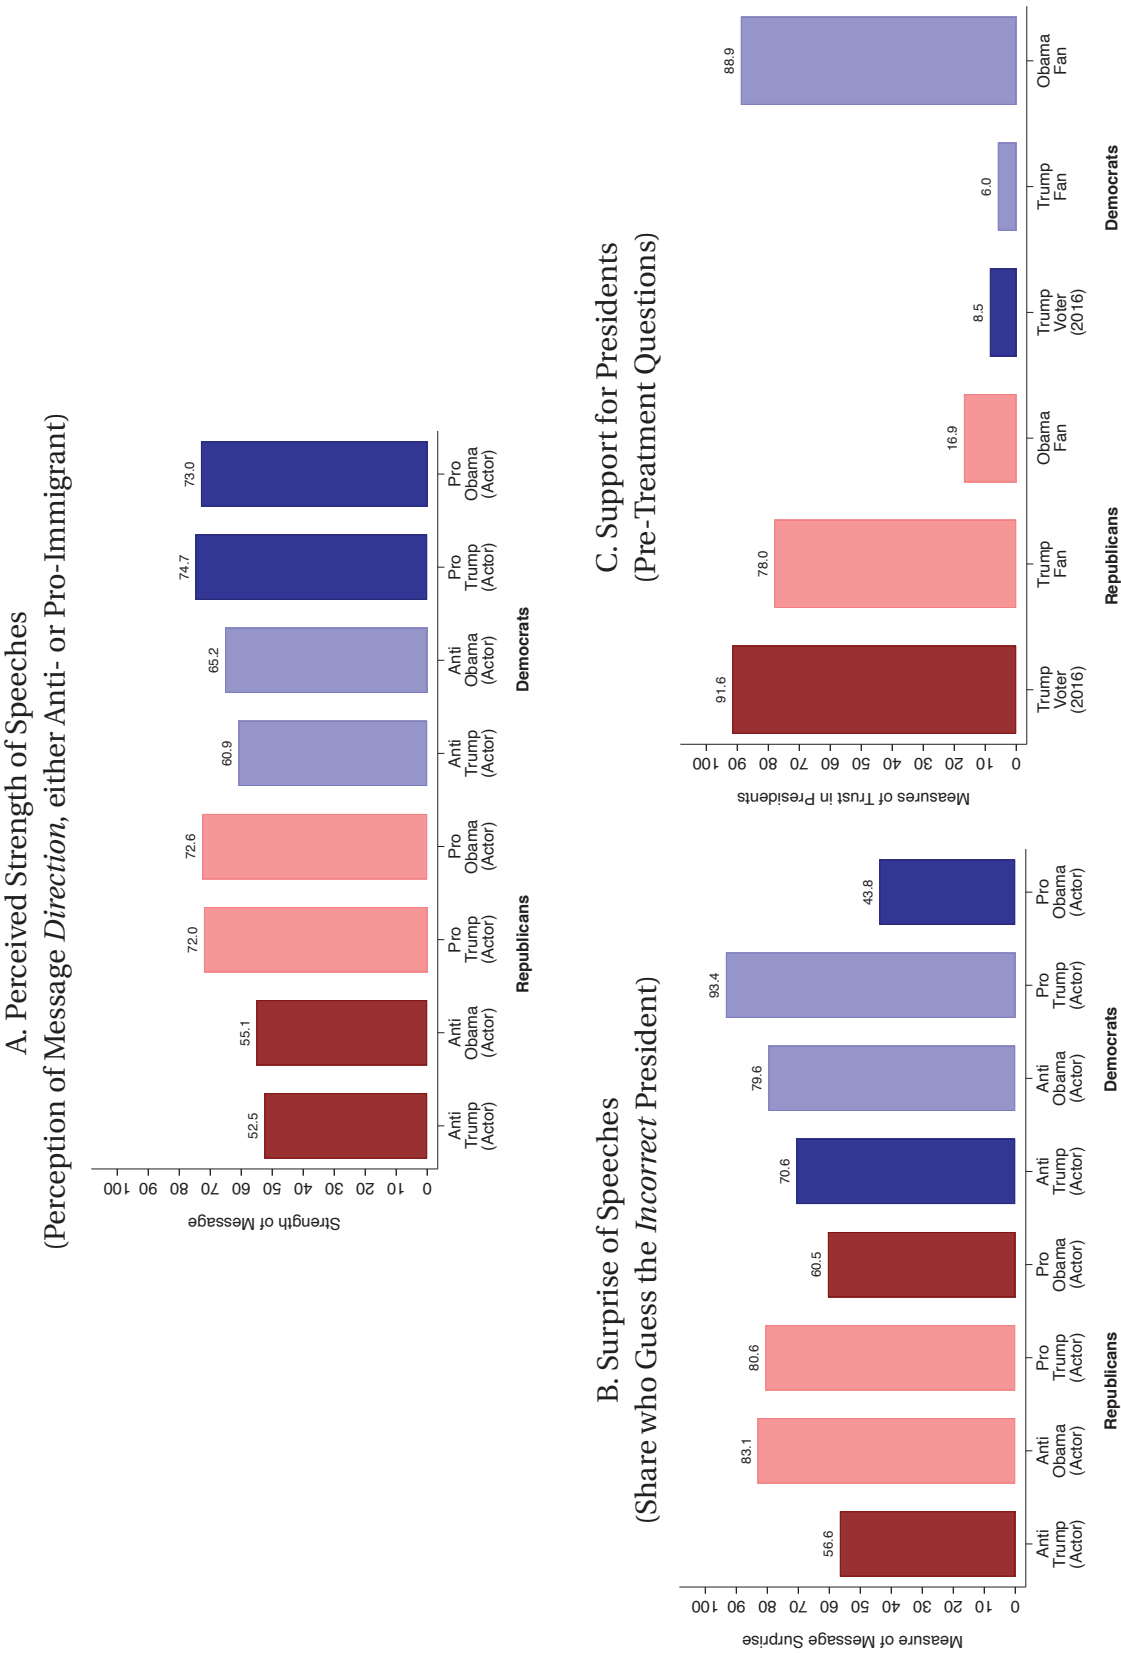
Notes: This figure depicts the treatment arms of the experiment. The sample is stratified by party, Republican or Democrat, and then within party all participants are randomized into the 11 treatment arms shown above.

Figure 2: Anti-Immigration Beliefs in No Audio Control Group



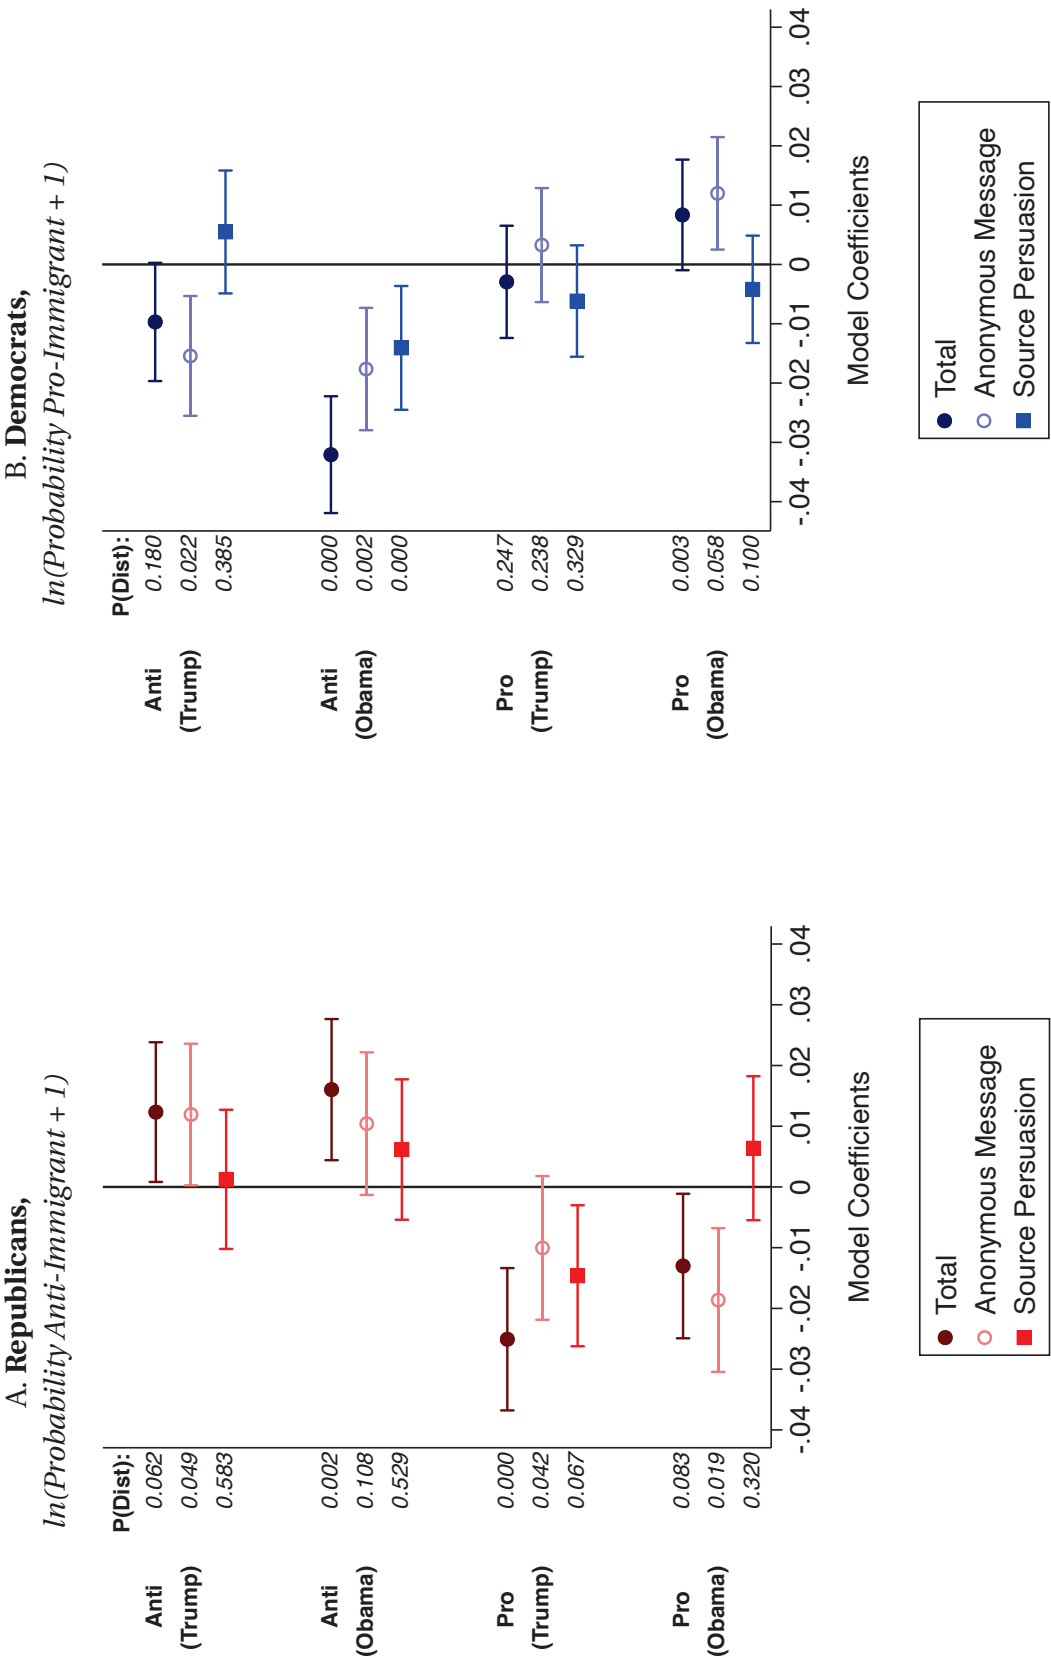
Notes: This plot shows the baseline distribution of the Anti-Immigration index or $P(\text{Anti} - \text{Immigration})$ in the no audio control group. This index is composed of 16 questions about immigration beliefs asked in the second part of the survey, described in Table 1.

Figure 3: Measures of Message Similarity, Surprise of Messages, and Source Favorability



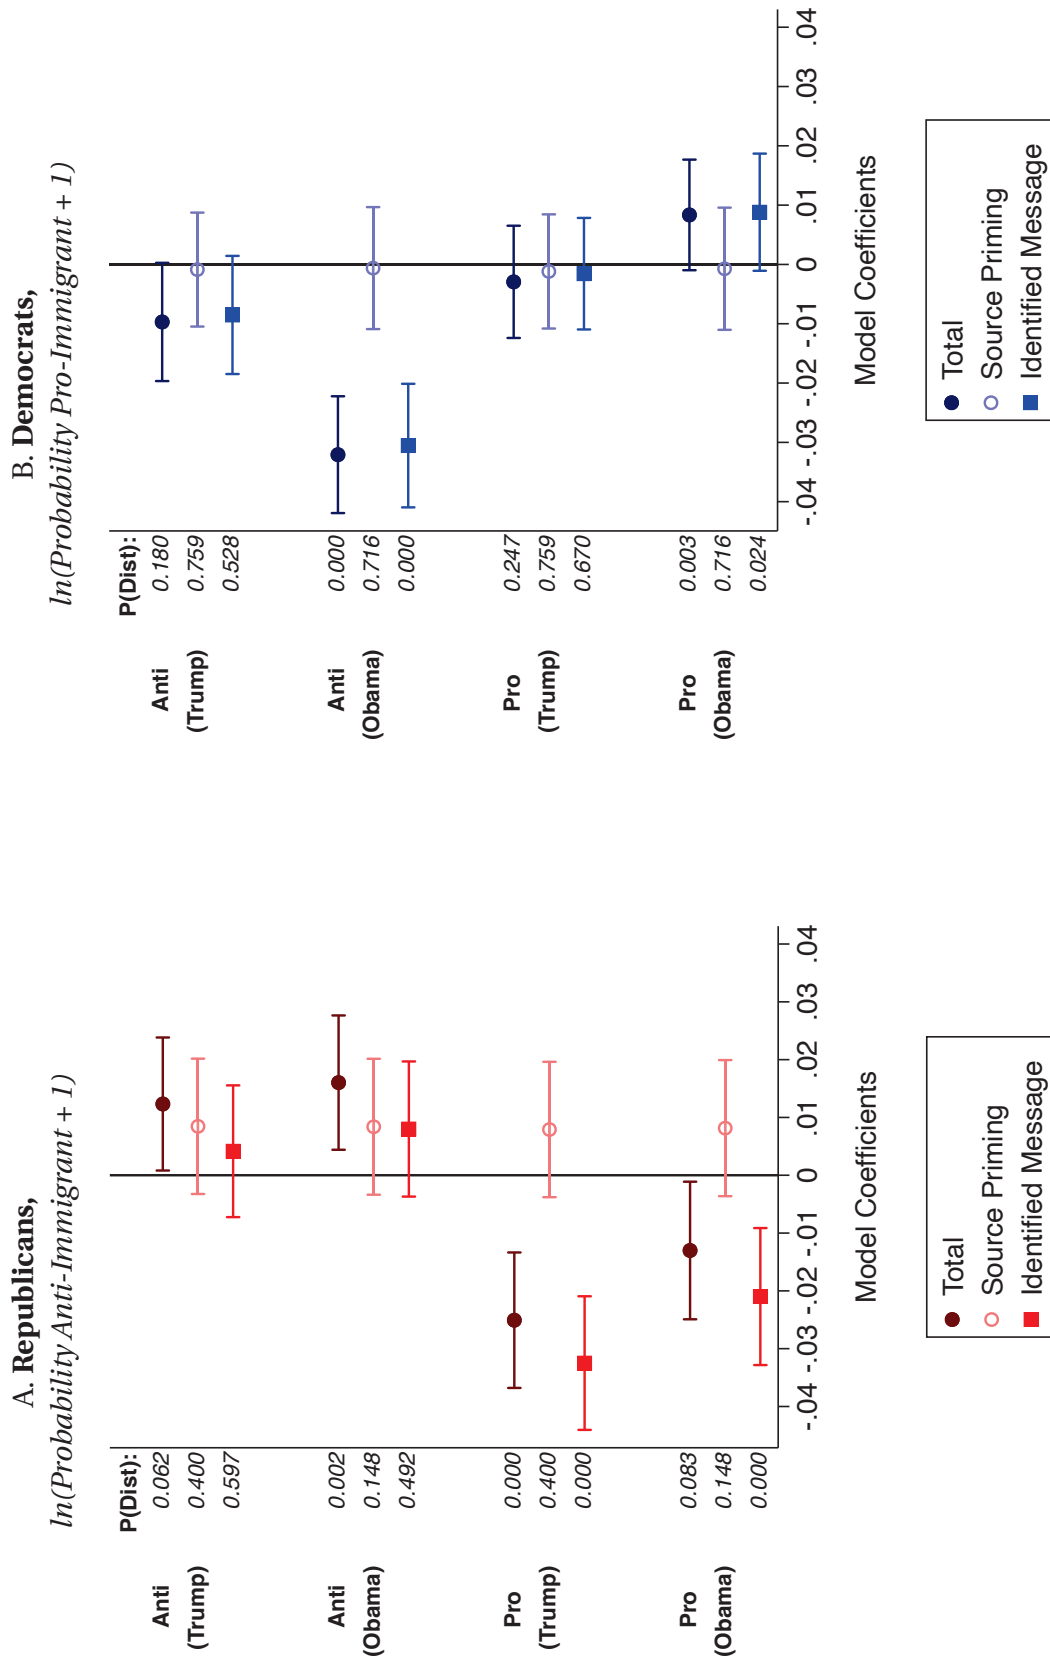
Notes: Panel A compiles participant responses to a question that asks participants how pro- or anti- immigrant they think the speech was after treatment. Specifically, Panel A plots the degree pro-immigrant for pro-immigrant speeches and the degree anti-immigrant for anti-immigrant speeches. The actor speeches shown in this figure do not reveal a stated source and can be compared across presidents to measure the “strength” of different speech messages. Panel B contains responses from a multiple choice question asked post-intervention about which president the participant thought was the source of the speech, from a list of the four most recent presidents (and an “Other” option). This plot shows the share of participants who guessed the *incorrect* president for each actor version of the treatments. This measure is the share of participants who would be surprised by the source of the speech. Panel C includes pre-treatment responses from questions asking whether participants voted for Trump in 2016 and whether participants are a fan of Trump and Obama.

Figure 4: Persuasion Decomposition: Impact of Source within Fixed Message



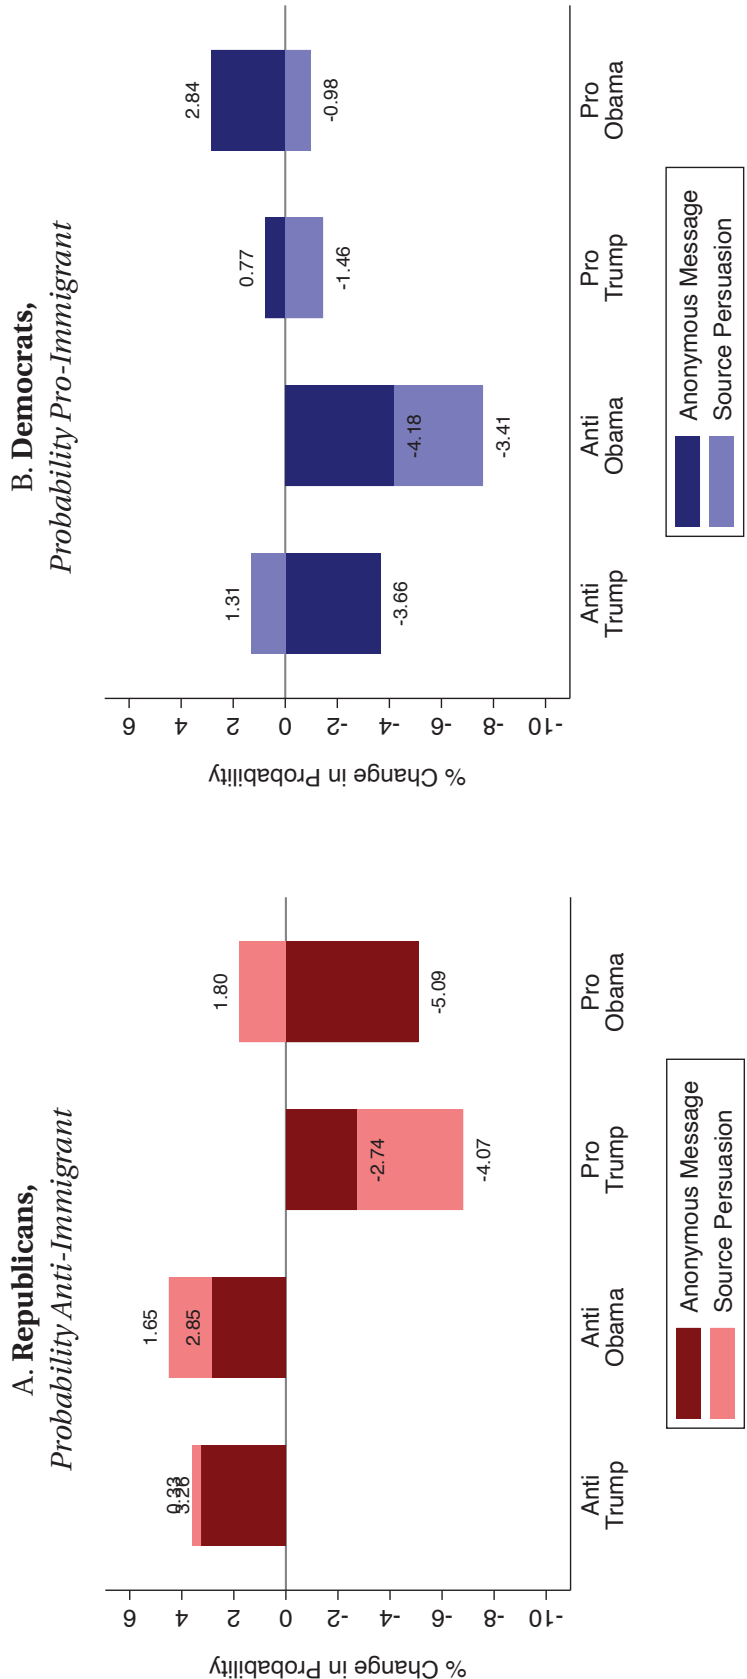
Notes: This figure presents the coefficient estimates and 95% confidence intervals of the first decomposition of message and source effects of a speech treatment (analogous to Table 4). This decomposition fixes the speech message, and varies the speaker source as either the actor or the president. The total estimate corresponds to estimates from Table 3, and the Anonymous Message and Source Persuasion effects correspond to β_m and β_{ms} from Table 4, respectively. Outcomes are constructed as $\ln(P(y) + 1)$, where the outcome is the index probability anti-immigrant (pro-immigrant) for Republicans (Democrats). All regressions additionally include covariates selected by the double lasso procedure to improve precision (See Appendix A3). “ $P(Dist.)$ ” is the p-value from a Kolmogorov-Smirnov test of equality of the distributions of the treatment and comparison groups.

Figure 5: Alternative Decomposition for Source Priming: Impact of Message within Fixed Source



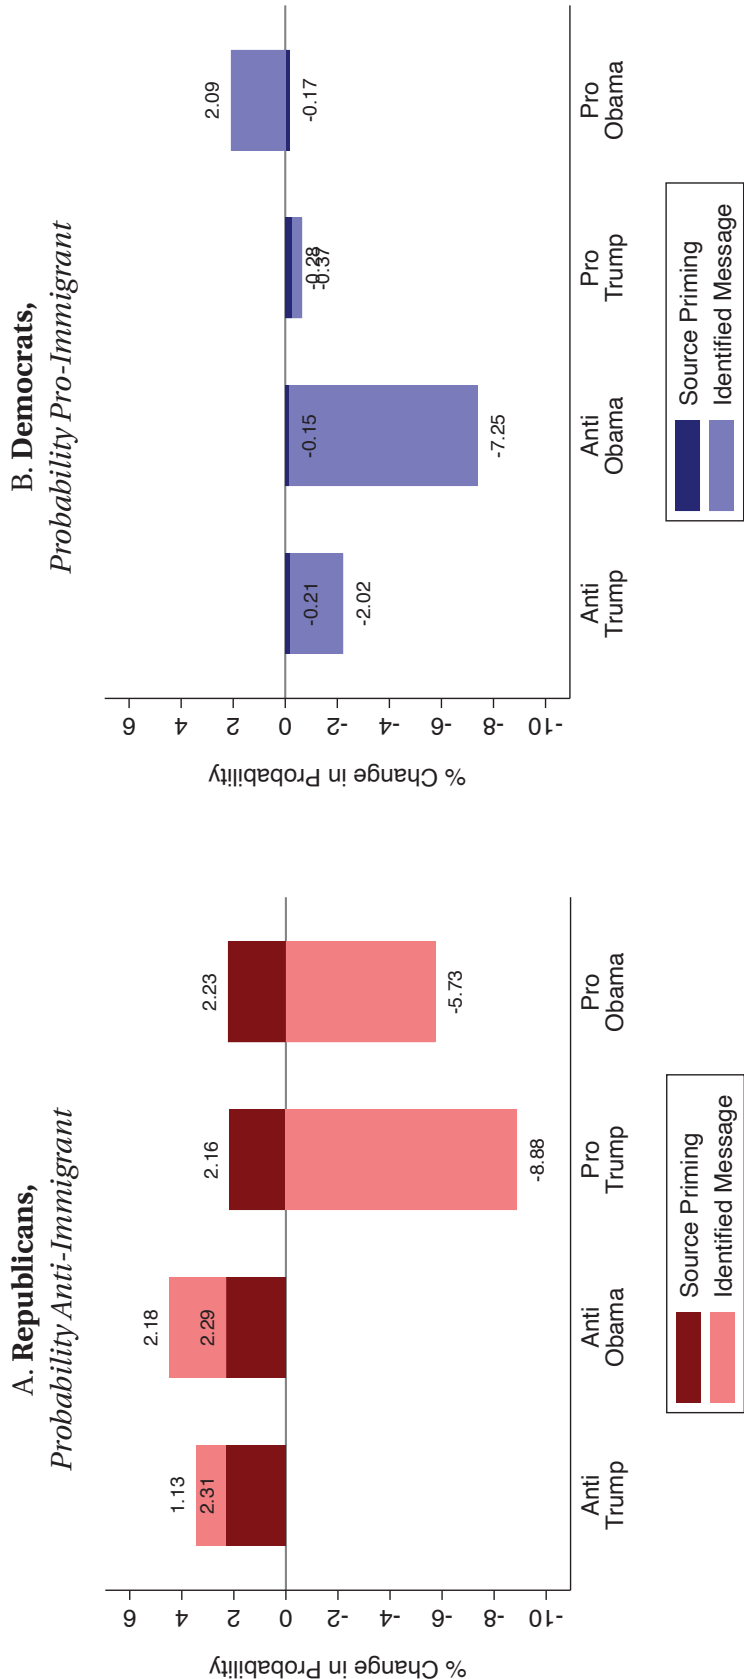
Notes: This figure presents the coefficient estimates and 95% confidence intervals of the first decomposition of message and source effects of a speech treatment (analogous to Table 5). This decomposition fixes the president source as either Trump or Obama's voice, and varies the message between the turkey pardon and an immigration speech. The total estimate corresponds to estimates from Table 3, and the Source Priming and Identified Message effects correspond to β_s and β_{sm} from Table 5, respectively. Outcomes are constructed as $\ln(P(y) + 1)$, where the outcome is the index probability anti-immigrant (pro-immigrant) for Republicans (Democrats). All regressions additionally include covariates selected by the double lasso procedure to improve precision (See Appendix A3). " $P(Dist)$," is the p-value from a Kolmogorov-Smirnov test of equality of the distributions of the treatment and comparison groups.

Figure 6: **Persuasion Decomposition:** Impact of Source within Fixed Message
(%Change in Probability)




Notes: This figure presents the implied percent changes in probability anti-immigrant or pro-immigrant, given the first decomposition, or the “%Diff.” estimates from Table 4. This decomposition fixes the speech message and varies the speech source as the voice actor or the president’s voice. The Anonymous Message and Source Persuasion effects correspond to β_m and β_{ms} from Table 4, respectively. The p-values for these estimates likewise correspond to the p-values for these coefficients in Table 4. Adding the two effects corresponds to the total change in probability for a particular president speech.

Figure 7: Alternative Decomposition for Source Priming: Impact of Message within Fixed Source
(%Change in Probability)



Notes: This figure presents the implied percent changes in probability anti-immigrant or pro-immigrant, given the first decomposition, or the “%Diff.” estimates from Table 5. This decomposition fixes the president source as either Trump or Obama’s voice, and varies the message between the turkey pardon and an immigration speech. The Source Priming and Identified Message effects correspond to β_s and β_{sm} from Table 5, respectively. The p-values for these estimates likewise correspond to the p-values for these coefficients in Table 5. Adding the two effects corresponds to the total change in probability for a particular president speech.

Figure 8: Polarization Change in Counterfactual Scenarios
 %Change in Party Difference in *Probability Anti-Immigrant*



Notes: This plot corresponds to the coefficient estimates in Table A1, and displays implied changes in differences in *Probability Anti-Immigrant*. The baseline party Difference in $P(Anti) = 0.308$. The estimates represent percent changes in the distance in the anti-immigrant index probability between Republicans and Democrats in the treatment versus control group. The treatments for each party vary according to constructed counterfactuals. “Own Leaders” are treatments that are Trump for Republicans and Obama for Democrats, while “Opposition Leaders” are the converse. “With Prior” messages are anti-immigrant speeches for Republicans and pro-immigrant speeches for Democrats, and “Against Prior” messages are the converse. Likewise, actor control groups correspond to replicate messages for a given treatment. Specifically, the corresponding regression is $\ln(P(Anti) + 1) = \beta_0 + \beta_{m1}1[Message_i = 1] + \beta_{m2}1[Message_i = 1] \times 1[Republican_i = 1] + \beta_{ms1}1[Message_i = 1] \times 1[Source_i = 1] + \beta_{ms2}1[Message_i = 1] \times 1[Source_i = 1] \times 1[Republican_i = 1] + \gamma X + \epsilon$, where the interaction coefficient β_{m2} represents the change in polarization from the Anonymous Message effect and β_{ms2} represents the change in polarization from the Source Persuasion effect. The covariates X correspond to those in Appendix A3 for both party groups. A positive estimate is an increase in polarization.

ONLINE APPENDIX

A1 Appendix Tables & Figures

Table A1: Polarization Change in Counterfactual Scenarios

		Polarization Change <i>ln(Probability Anti)</i> <i>Change Diff.: Republican-Democrat</i>			
Leader	Effect	β	(S.E.)	%Diff.	N
<i>Message With Prior</i>					
Own Leader	Total	0.023***	(0.009)	10.90%	2477
Own Leader	Anonymous Message	0.029***	(0.009)	13.26%	
Own Leader	Source Persuasion	-0.007	(0.009)	-2.59%	3669
Opposing Leader	Total	0.014	(0.009)	7.88%	2400
Opposing Leader	Anonymous Message	0.011	(0.008)	5.22%	
Opposing Leader	Source Persuasion	0.003	(0.009)	2.45%	3590
<i>Message Against Prior</i>					
Own Leader	Total	-0.066***	(0.009)	-29.72%	2429
Own Leader	Anonymous Message	-0.031***	(0.009)	-13.64%	
Own Leader	Source Persuasion	-0.035***	(0.009)	-15.97%	3584
Opposing Leader	Total	-0.025***	(0.009)	-11.58%	2389
Opposing Leader	Anonymous Message	-0.037***	(0.009)	-17.14%	
Opposing Leader	Source Persuasion	0.012	(0.009)	5.59%	3601

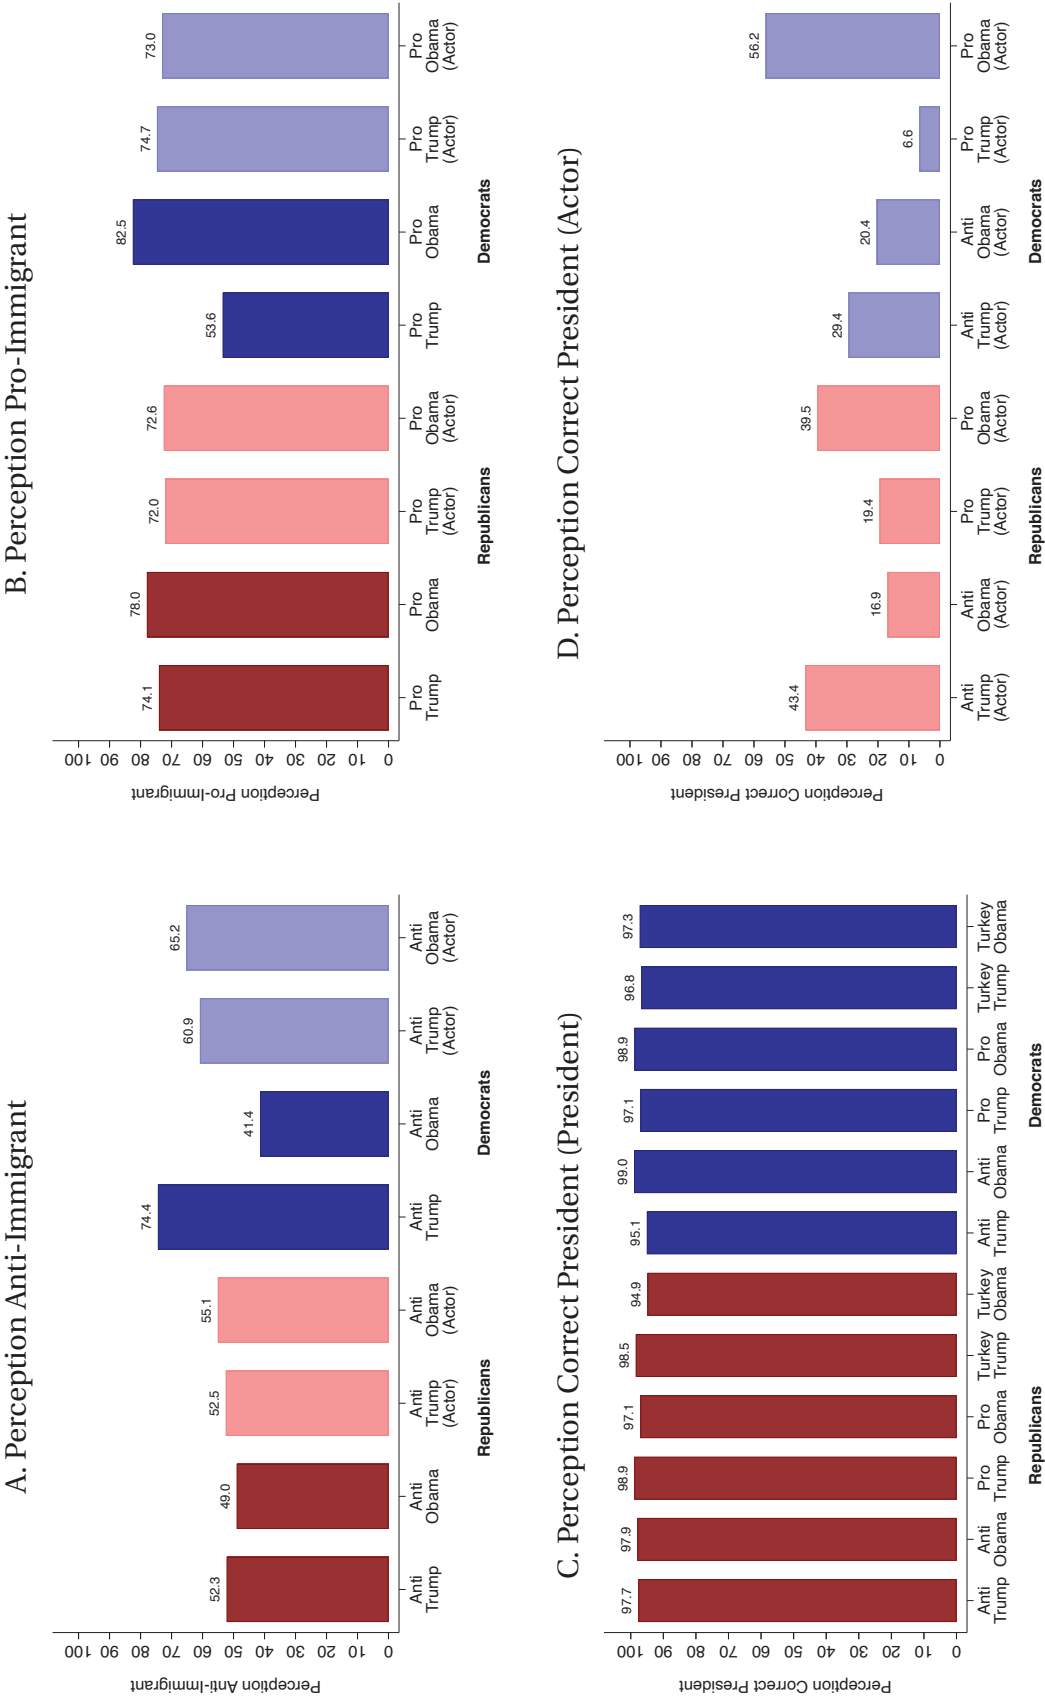
Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table corresponds to the probability changes in Figure 8, and displays coefficient estimates for differences in $\ln(\text{Probability Anti} - \text{Immigrant} + 1)$. The baseline party Difference in $P(\text{Anti}) = 0.308$. The estimates represent percent changes in the distance in the anti-immigrant index probability between Republicans and Democrats in the treatment versus control group. The treatments for each party vary according to constructed counterfactuals. "Own Leaders" are treatments that are Trump for Republicans and Obama for Democrats, while "Opposition Leaders" are the converse. "With Prior" messages are anti-immigrant speeches for Republicans and pro-immigrant speeches for Democrats, and "Against Prior" messages are the converse. Likewise, actor control groups correspond to replicate messages for a given treatment. Specifically, the corresponding regression is $\ln(P(\text{Anti}) + 1) = \beta_0 + \beta_{m1}1[\text{Message}_i = 1] + \beta_{m2}1[\text{Message}_i = 1] \times 1[\text{Republican}_i = 1] + \beta_{ms1}1[\text{Message}_i = 1] \times 1[\text{Source}_i = 1] + \beta_{ms2}1[\text{Message}_i = 1] \times 1[\text{Source}_i = 1] \times 1[\text{Republican}_i = 1] + \gamma X + \epsilon$, where the interaction coefficient β_{m2} represents the change in polarization from the Anonymous Message effect and β_{ms2} represents the change in polarization from the Source Persuasion effect. The covariates X correspond to those in Appendix A3 for both party groups. A positive estimate is an increase in polarization. All regressions additionally include covariates selected by the double lasso procedure to improve precision for both party groups (See Appendix A3). "% Diff." calculates the implied change in the difference of outcome *probability* across party, or the untransformed probability index, due to treatment, relative to the baseline difference across party in the control group.

Table A2: Multiple Hypothesis Testing Results,
Hypotheses Mirrored by Party (Grouped)

Persuasion Decomposition				
<i>Outcome: $\ln(\text{Prob}(\text{Convinced})+1)$</i>	β	Adj. P-Value	Y-Mean	N
A. Anonymous Message, β_m				
Message with Prior	0.0110***	0.0093	0.498	3616
Message against Prior	0.0173***	0.0003	0.281	3601
B. Persuasion, β_{ms}				
Own Leader Message with Prior	-0.0045	0.526	0.512	2435
Own Leader Message against Prior	0.0149**	0.0237	0.300	2350
Opposition Leader Message with Prior	-0.0078	0.2756	0.507	2356
Opposition Leader Message against Prior	-0.0048	0.4067	0.298	2367
Alternative Decomposition for Priming				
<i>Outcome: $\ln(\text{Prob}(\text{Convinced})+1)$</i>	β	Adj. P-Value	Y-Mean	N
C. Priming, β_s				
Own Leader	-0.0009	0.8470	0.498	2397
Opposition Leader	0.0028	0.7870	0.498	2428
D. Identified Message, β_{sm}				
Message with Prior Own Leader	0.0096*	0.0973	0.497	2406
Message against Prior Own Leader	0.0344***	0.0003	0.282	2358
Message with Prior Opposition Leader	0.0150**	0.0230	0.273	2349
Message against Prior Opposition Leader	-0.0035	0.474	0.505	2360


Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table produces estimates of regressions that correspond to the set of underlying hypotheses in this study. All hypotheses are grouped in the sense that we expect results to be mirrored by party affiliation, respective leader alignment with party, and message alignment with the presumed priors of a party. Thus, we can group our data into regressions that include both parties and use an outcome of $\ln(\text{Prob}(\text{Convinced}) + 1)$, where $\text{Prob}(\text{Convinced})$ is a collapsed outcome in each test family. In family A, B and D, $\text{Prob}(\text{Convinced})$ is $P(\text{Anti})$ for anti-immigration messages and $P(\text{Pro})$ for pro-immigration messages, meaning this outcome measures movement towards the message direction. In family C, the $\text{Prob}(\text{Convinced})$ corresponds to $P(\text{Anti})$ for Republicans and $P(\text{Pro})$ for Democrats. Stars and p-values correspond to multiple hypothesis testing adjustment within families of tests, using the bootstrap algorithm developed in [List et al. \(2021\)](#) and corresponding stata command `mhtexp2`. All regressions include controls for covariates defined by the lasso procedure (the union of controls for used in the separate party regressions in the main body of the paper). $Y - \text{Mean}$ corresponds to the mean of the outcome in the control group for each regression. This algorithm flexibly permits the use of researcher defined sets of hypotheses, adjustments for controls, and multiple outcomes within a family of tests.

Figure A1: Perception Treatment Content and Source



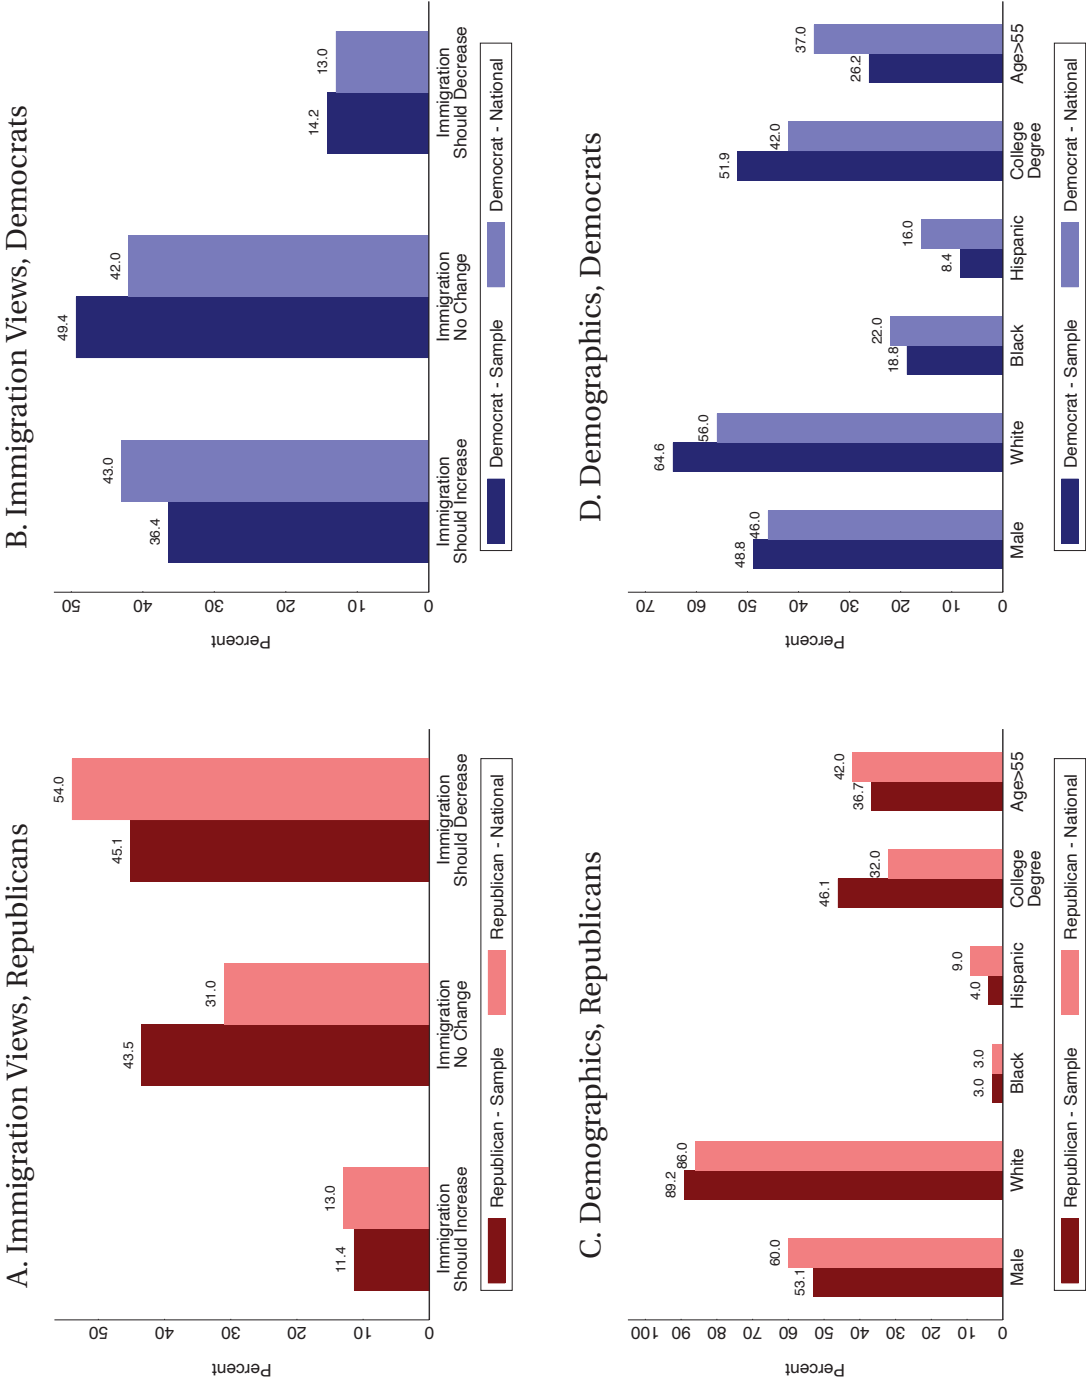
Notes: Panels A and B compile participant responses to a question that asks participants how pro- or anti- immigrant they think the speech was after treatment. Specifically, Panels A and B plot the degree pro-immigrant for pro-immigrant speeches and the degree anti-immigrant for anti-immigrant speeches. The actor speeches shown in this figure do not reveal a stated source and can be compared across presidents to measure the “strength” of different speech messages. Panels C and D plot responses from a multiple choice question asked post-intervention about which president the participant thought was the source of the speech, from a list of the four most recent presidents (and an “Other” option). This plot shows the share of participants who guessed the *correct* president for each of the treatments. Panel C shows the results for the revealed source president versions of speeches and Panel D shows the results for the actor versions of the treatments, without the revealed president source. This is the inverse of the share of participants who would be surprised by the source (share that guess the *incorrect* president) in Panel B of Figure 3.

Figure A2: Surprise of Speeches (Party)
(Share who Guess the *Incorrect* Party)



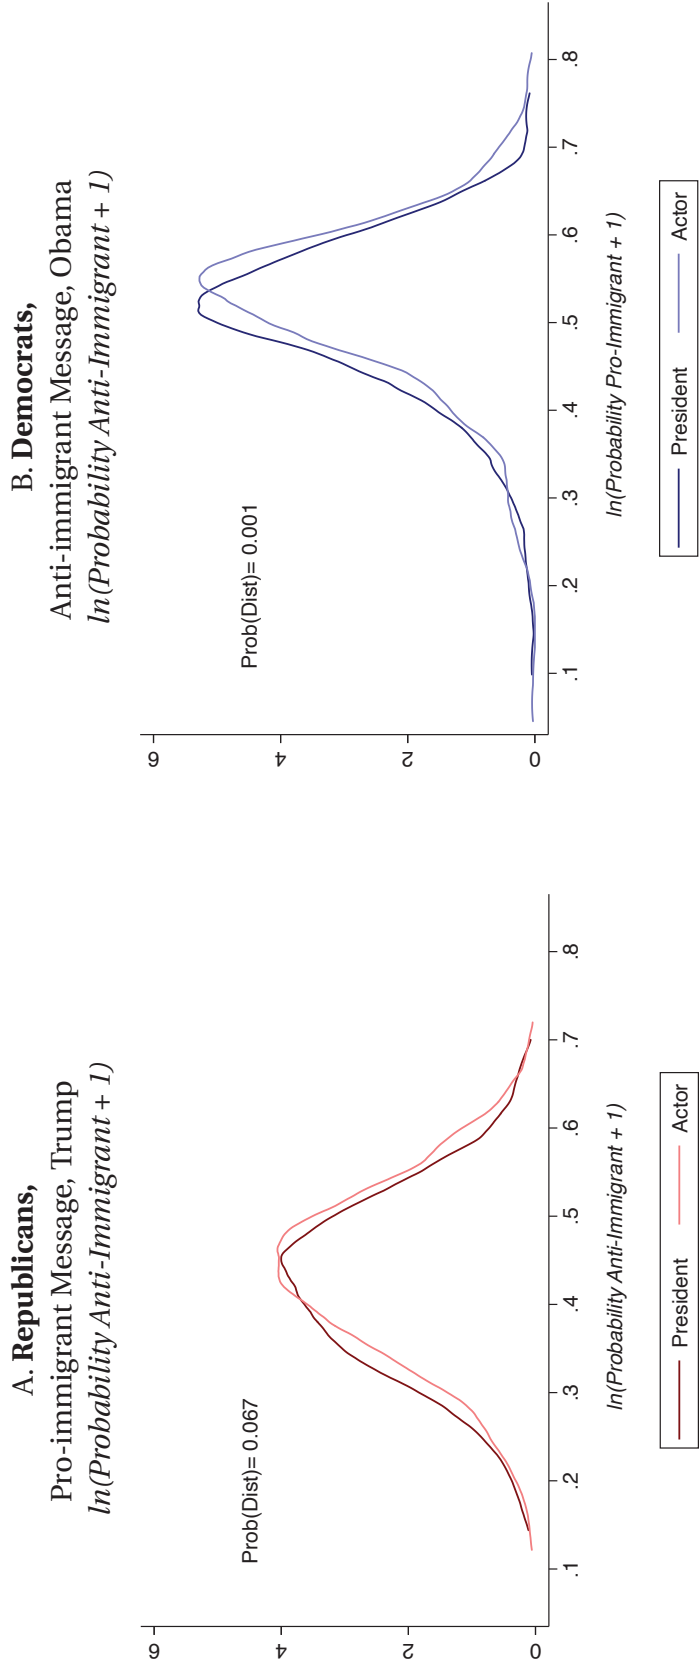
Notes: This figure replicates Panel B of Figure 3 using guesses of political party rather than president source. The figure is constructed from a post-intervention question which asked about which president the participant thought was the source of the speech, from a list of the four most recent presidents (and an “Other” option). This plot shows the share of participants who guessed the *incorrect* political party for each actor version of the treatments. This measure is the share of participants who would be surprised by the source of the speech.

Figure A3: Sample Comparison to Party Demographics



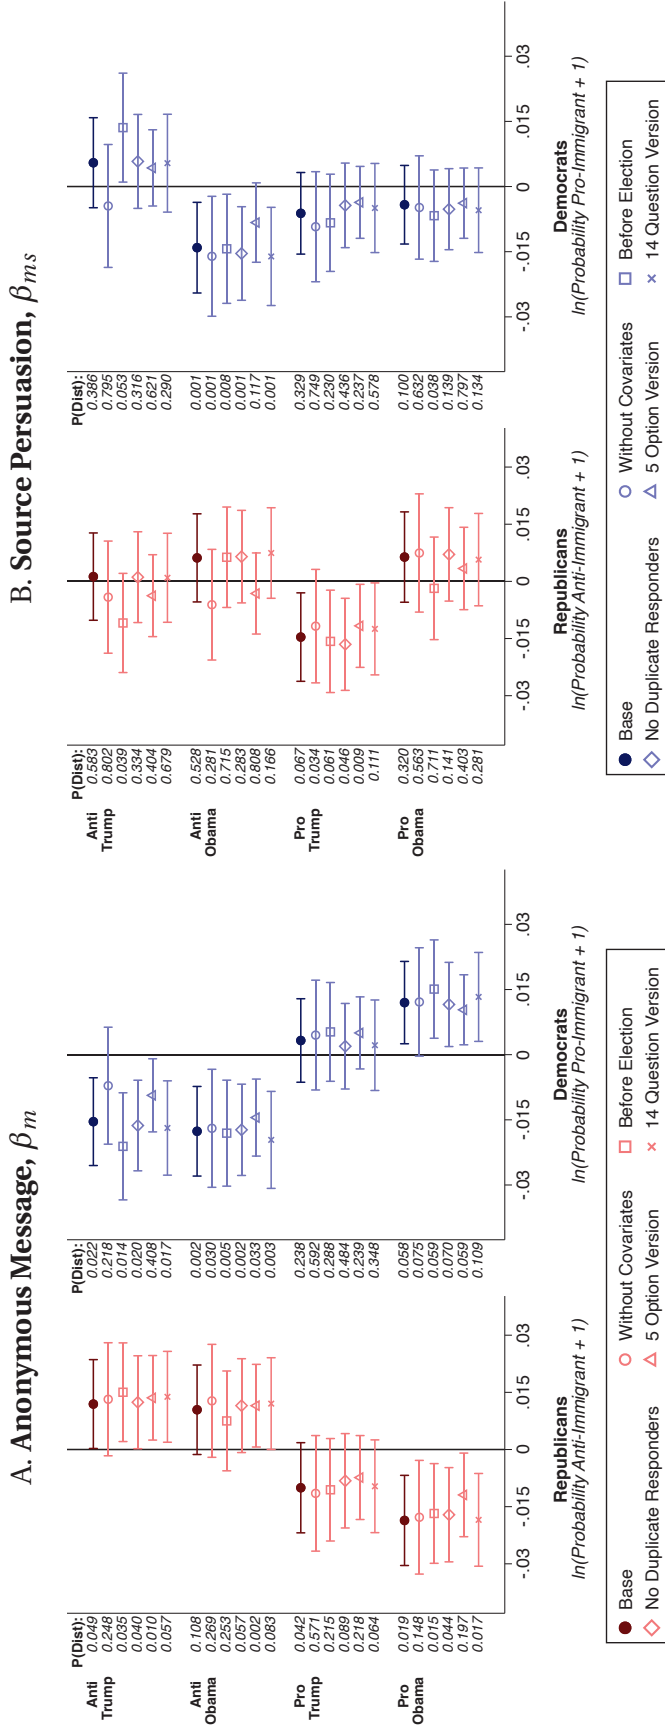
Notes: This set of plots compares characteristics of the study sample to national data for political parties. The top panel compares pre-treatment survey responses to a Gallup Survey question about whether immigration should be increased, decreased or stay the same (Jones, 2019). The national comparison is compiled from survey evidence and data on eligible voters from the U.S. Census, and Pew Research (File, 2018; Igielnik and Budiman, 2020; Doherty et al., 2020).

Figure A4: Source Persuasion Effect
Distribution of Outcome by Treatment Group (Corresponds to β_{ms})



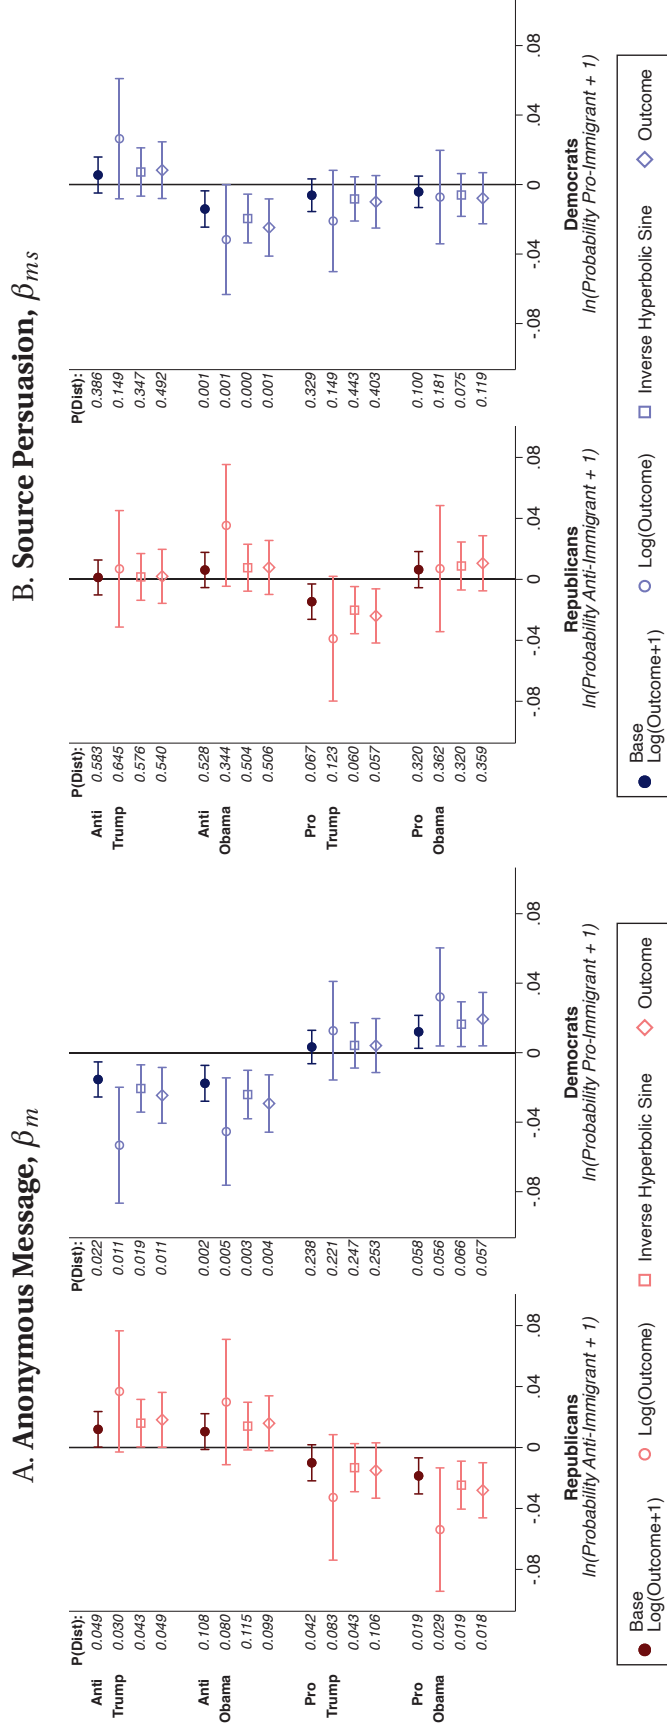
Notes: This figure plots the distribution of the outcomes for the two president treatments with Source Persuasion effects, the Pro-Immigrant Trump message for Republicans and the Anti-Immigrant Obama message for Democrats, relative to the replicate versions recorded by the actor. Outcomes are measured as $\ln(y + 1)$, and are first residualized by the set of control variables from the double lasso procedure described in Appendix A3 before being plotted. “Prob(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and control groups for each estimate.

Figure A5: Robustness Specifications



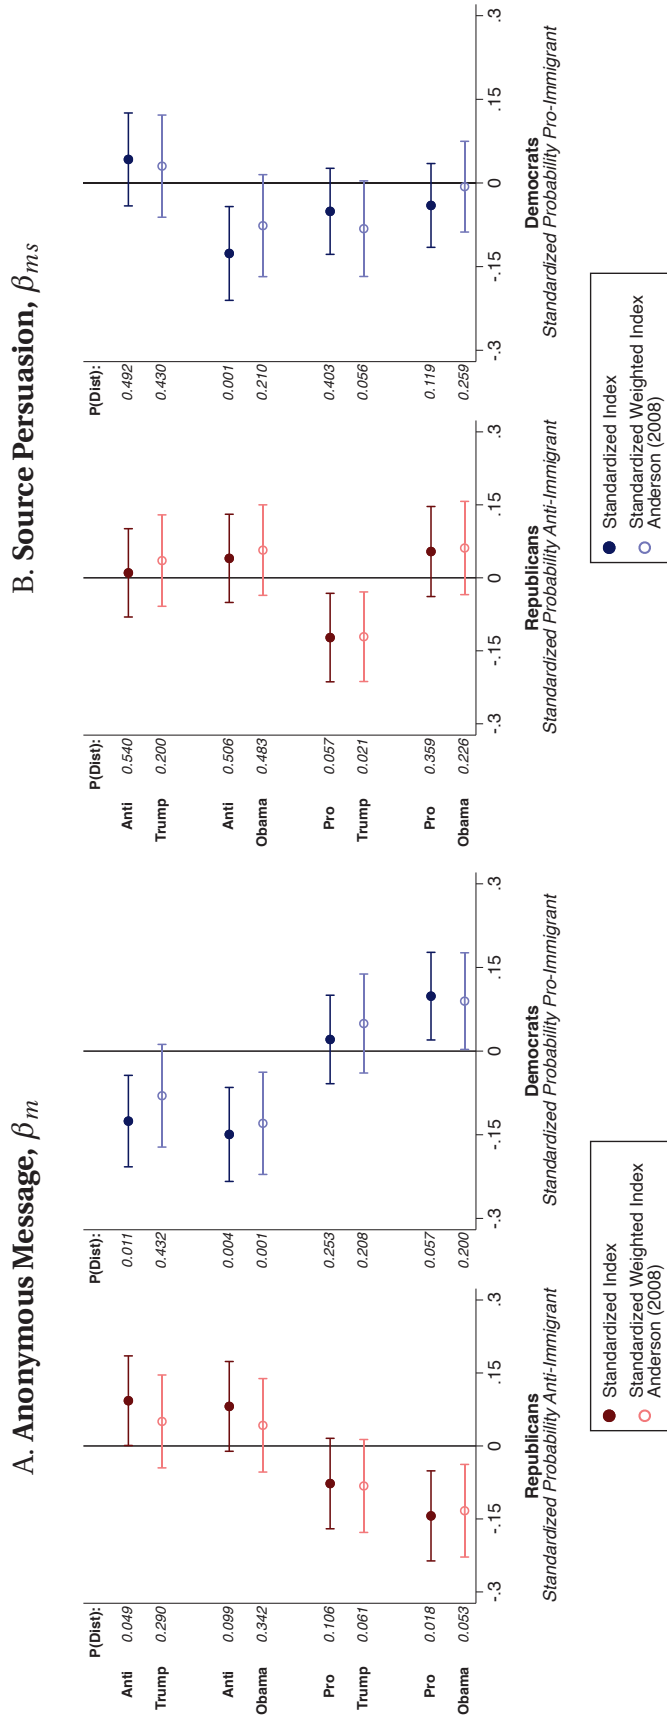
Notes: These figures replicate the baseline specification for Decomposition II for alternative sample restrictions and specifications. All outcomes are measured using $\ln(y + 1)$. “Without Covariates” is a specification that drops the covariate controls described in Appendix A3. “No Duplicate Responders” excludes any individuals who attempted the survey more than once; the baseline sample includes first attempts of the survey for these individuals. “5 Option Version” constructs the index using the full set of 5 options for questions that have 5 options (e.g. Strongly Agree, Agree, Neutral, Disagree, Strongly Agree), rather than the baseline version that collapses all answers to have 3 options to match 3 option questions (e.g. Agree, Neutral, Disagree). “14 Question Version” constructs the index to match the pre-registered version, which excludes questions on the societal contribution of legal immigrants or immigrants from English-speaking countries. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A6: Robustness of Log Specification



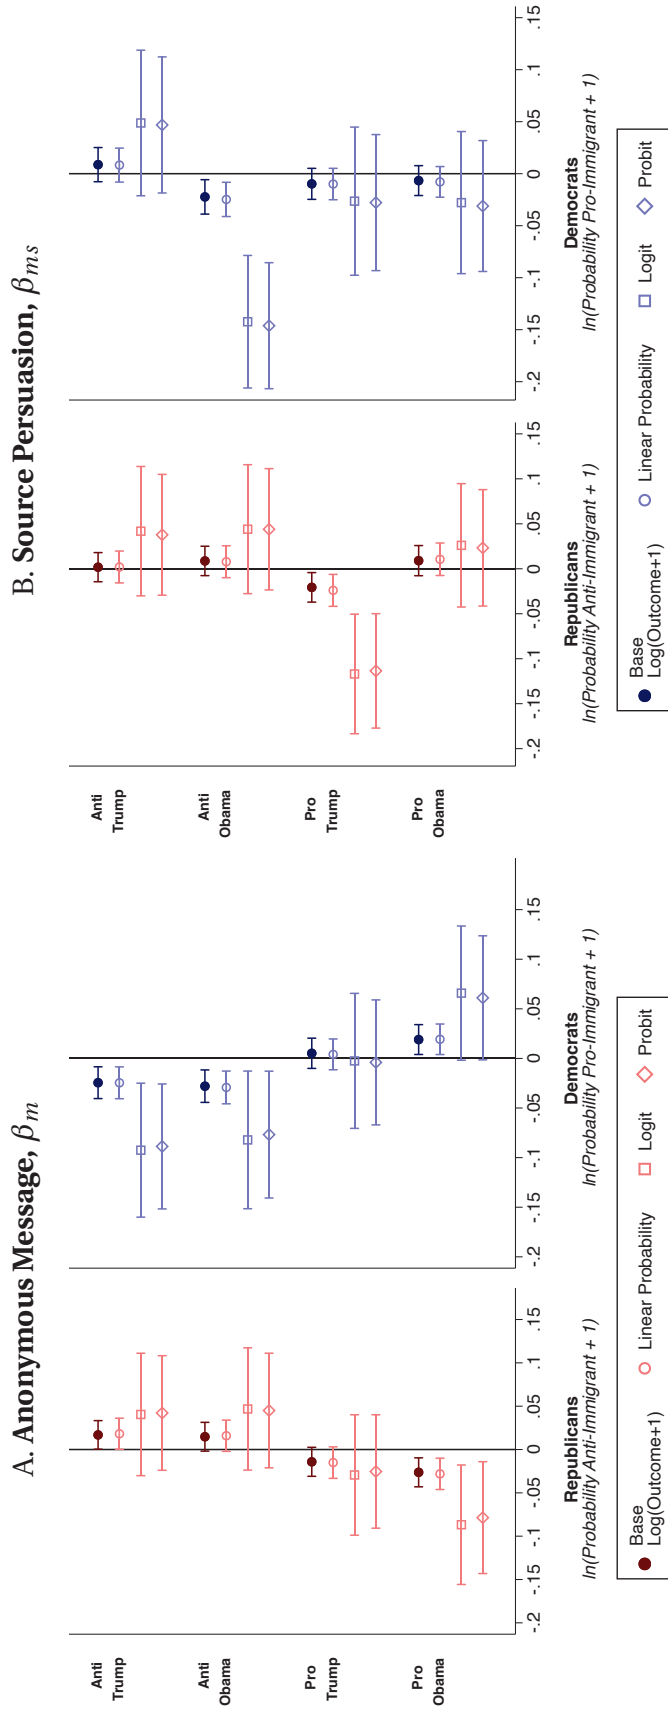
Notes: These figures replicate the baseline specification for alternative sample restrictions and specifications. The specifications compare the baseline transformation, $\ln(P(y) + 1)$, with the alternative transformation, $\ln(P(y))$, where any observations with $P(y) = 0$ are dropped. The figure also includes the unlogged outcome $P(y)$ as the dependent variable. The standardized version of the unlogged outcome is not included in this plot given the difference in scale of this outcome; however, the results using this transformation are identical in significance and direction as the unlogged outcome, given that standardization is a linear transformation. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A7: Standardized Index Outcomes



Notes: These figures replicate the baseline specification using an alternative formulation of the outcome. The first outcome takes the probability measure, $P(y)$, which is an average across all 16 questions and norms this outcome using values in the "No Audio" control group (separately for each party). The resulting standardized index has a mean of 0 and a standard deviation of 1 in the control group. The second outcome constructs a weighted standardized version of the index, where weights are determined by the inverse covariance matrix of the individual questions, following (Anderson, 2008). The second outcome is likewise normalized using values in the "No Audio" control group (separately for each party). Given that these outcomes cover the range of $\approx (-1, 1)$, the scale of effects differs from the primary figures. All regressions are adjusted for covariates in Appendix A3.

Figure A8: Change in Probability: Log, Linear Probability, Logit & Probit



Notes: These figures replicate the baseline specification for alternative sample restrictions and specifications. Unlike the preceding tables, all coefficients from specifications are converted to corresponding changes in underlying probability indexes using the delta method, or estimates of marginal effects (at means) from logit or probit models. The estimates are drawn from the baseline transformation, $\ln(P(y) + 1)$, a linear probability model with $P(y)$ as the dependent variable, a logit model, and a probit model. To estimate the probit and logit models, the continuous $P(y)$ measures are discretized at 0.5, or a “neutral” response probability. All regressions are adjusted for covariates in Appendix A3.

A2 Conceptual Framework: Additional Details

This section provides detailed derivations for results presented in Section 2.

Derivations for Equation (9). Adding 1 to both sides of Equation (8), we get

$$1 + P(\omega|s) = 1 + P(\omega) \times \Theta(s|\omega) = (1 + P(\omega)) \left(1 + \frac{P(\omega)}{1 + P(\omega)} (\Theta(m|\omega) - 1)\right) \quad (14)$$

Taking logs of both sides of the equation above gives us Equation (9).

Derivations for Remark 2. Using Equation (4) we have

$$\Theta(s|\omega, m) - 1 = \frac{P(s|\omega, m) - P(\omega|m)}{P(\omega|m)} \quad (15)$$

Now, note that

$$P(\omega|m) = P(\omega, s|m) + P(\omega, \neg s|m) = P(s|m)P(\omega|s, m) + P(\neg s|m)P(\omega|\neg s, m) \quad (16)$$

Finally, plugging this into the numerator in Equation (15) and using $P(\neg s|m) = 1 - P(s|m)$ gives us the expression of interest:

$$\Theta(s|\omega, m) - 1 = \frac{(1 - P(s|m))(P(\omega|s, m) - P(\omega|\neg s, m))}{P(\omega|m)} \quad (17)$$

Derivations for Equation (6). Adding 1 to both sides of Equation (4), we get

$$1 + P(\omega|s, m) = 1 + P(\omega|m) \times \Theta(s|\omega, m) = (1 + P(\omega|m)) \left(1 + \frac{P(\omega|m)}{1 + P(\omega|m)} (\Theta(s|\omega, m) - 1)\right) \quad (18)$$

Taking logs of both sides of the equation above gives us Equation (6).

A3 Control Variable Selection: Double Lasso

To select the control variables for our baseline specification we use the post double selection lasso (PDS) methodology of Belloni et al., (2012, 2014, 2015, 2016), using the STATA package by Ahrens et al., (2020). The procedure consists of first including all potential control variables from baseline and converting categorical variables into sets of indicators. Our set of potential controls is quite rich given the large number of survey questions we ask to participants prior to the treatment. These questions can be viewed in Supplemental Materials Appendix S1.

Next we keep the controls selected from the LASSO minimization of a regression which predicts treatment assignment. Then we repeat this exercise for a regression that predicts the outcome variable and take the union of controls selected from both procedures as our baseline controls. We do this procedure separately for Democrats and Republicans, given that our randomization was stratified by party and that the demographic characteristics of party groups differ in meaningful ways.

From the first step, none of the control variables are selected under the lasso procedure as important predictors of treatment status. From the second step, several variables are selected which are predictive of the outcome variable.

Republican Sample

The variables selected for Republicans are whether a participant is Hispanic, is in the age group 55 to 64, voted for Hillary Clinton in 2016, voted for a third party candidate in 2016, had the view that the level of immigration in the U.S. should decrease (pre-treatment), has the view that gun control regulations should be less strict, has the view that abortion should be illegal in all cases, believes taxes are too high, believes that there should be no government intervention in the healthcare system, is an occasional twitter user, is an occasional Buzzfeed reader, is a fan of Donald Trump, is a fan of LeBron James, is a fan of Taylor Swift, is a fan of Bill Gates, is a fan of Barack Obama, has no opinion towards Barack Obama, and views the most important policy issue to their vote as healthcare.

Democrat Sample

The variables selected for Democrats are whether a participant is Black, is Hispanic, is in the age group 35 to 44, is in the age group 45 to 54, has a high school degree as their highest level of education, voted for Hillary Clinton in 2016, had the view that the level of immigration in the U.S. should decrease (pre-treatment), had the view that the level of immigration in the U.S. should stay the same (pre-treatment), has the view that abortion should be illegal in all cases, has the view that abortion should be legal in some cases, believes taxes are too high, believes that there should be no government intervention in the healthcare system, feels neutral about government intervention in the healthcare system, is a daily, weekly, or occasionally reader of the New York Times (separate indicators), is a daily or occasional viewer of TV as a news source (separate indicators), is a daily newspaper reader, is a daily, weekly, or occasional viewer of Fox News (separate indicators), is a daily or weekly Breitbart News reader (separate indicators), is an occasional Buzzfeed reader, has no opinion of LeBron James, is a Barack Obama fan, is a Donald Trump fan, has no opinion of Donald Trump, lives in the Western U.S., region response west, and views the most important policy issue to their vote as taxes.

While these variables improve the precision of estimates, it is important to note that the models are robust to excluding all controls, as is shown in Appendix Figure A6.

A4 Sample Construction: Additional Details

This appendix includes detail about the construction of the analysis sample. Figure A9 displays these choices, separated by pre-treatment or baseline sample restrictions (in gray) and additional sample quality restrictions (in yellow).

Baseline Restrictions

The first set of treatment restrictions are determined pre-treatment, and we refer to these restrictions as baseline restrictions.

Several of these restrictions were pre-set within the survey environment by design, and failure to meet these qualifications resulted in the survey being automatically terminated for a respondent before treatment. These include whether a respondent: consented to the study, had audio capability to participate and was a U.S. citizen. Further, individuals were recruited as either Democrats or Republicans by our panel survey research partner. If they then entered the survey and answered that they were affiliated with the opposite party (e.g. recruited Democrats said they were Republicans within the survey), they were also exited from the survey environment. For this restriction, we did keep all individuals who answered that they were independents within the survey, although all participants were recruited as either Democrats or Republicans. Likewise, all survey respondents listen to a weather forecast audio clip prior to treatment and are asked questions about the content of the clip as an attention check. This attention check was included in our pre-analysis plan, and individuals who fail the attention check in the survey are automatically exited from the survey environment prior to treatment.

Our baseline sample restrictions also include dropping individuals who did not complete the survey with a valid recruitment survey ID (or were not recruited through our partner), and keeping only the first attempts of the survey for people who took the survey multiple times. These restrictions were coordinated with our survey panel partner, who replaced observations when these conditions failed.

As shown in Figure A9, the recruitment screen conditions (all except the attention check) yield an initial starting sample of 18,199 persons. After excluding individuals who fail the attention check, the sample is reduced to 14,356 respondents.

Additional Quality Restrictions

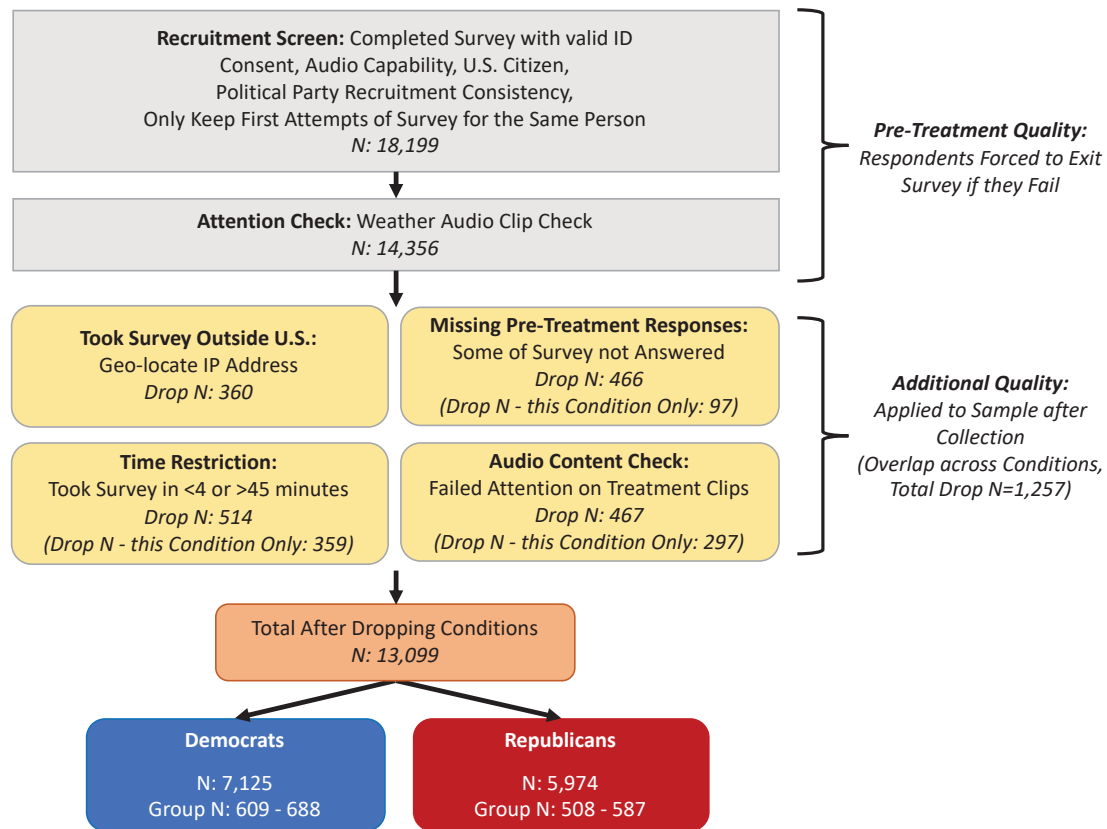
Next, we apply four additional restrictions to the sample in order to increase the quality of the responses in our final data set (and the precision of our estimates). Collectively, these restrictions moderately reduce the sample size by 1,257 individuals, or reducing the sample by 8.75% relative to the sample with only the baseline set of restrictions.

The first restriction is to exclude individuals who took the survey outside the U.S., based on the geolocation of I.P. addresses provided by Qualtrics (360 individuals). We targeted recruitment of U.S. citizens who are eligible to vote for this study, and we were concerned that responses entered outside the country may have been less likely to fit this qualification.

The second restriction is to exclude individuals who did not answer all of the pre-treatment questions (466 individuals). This restriction allows us to apply the Lasso procedure and control for select pre-treatment covariates in our regressions, by assuring that all respondents have complete pre-treatment data.

Third, we excluded respondents who took the survey in an exceptionally short or long time period (< 4 or > 45 minutes) (514 individuals). Time of survey is calculated net of the time of the

Figure A9: Sample Restrictions



Notes: This figure displays the sample restrictions used to construct the study data set. The restrictions in gray are baseline restrictions. Individuals who do not consent to the survey, indicate that they are not U.S. citizens, or do not have audio capability to complete the survey are automatically exited from the survey following these responses. Likewise, individuals who fail the attention check in the survey are automatically exited from the survey environment. Further, individuals who are recruited with a particular party affiliation (Democrat or Republican) and then answer with an inconsistent party affiliation within the survey are automatically exited. The baseline restrictions also include a requirement that we only keep individuals with a valid ID and first attempts for individuals who completed the survey multiple times. The second set of restrictions, in yellow, are applied for sample quality. These include dropping individuals who did not answer all pre-treatment questions, dropping individuals who took the survey outside of the U.S., and dropping individuals who took the survey in a very short or very long time period (< 4 or > 45 minutes). We also exclude individuals who were assigned to an audio clip in the survey and failed a content check question after listening to the clip. Individuals can fall into multiple categories in the second set of restrictions, and the overall number of individuals dropped from these restrictions is 1,257 respondents.

different audio clips that individuals are assigned to according to their treatment group. After applying the baseline sample restrictions, the 5th percentile of survey time is 4.9 minutes and the 95th percentile is 25.75 minutes; thus our restriction is only dropping respondents at the far tails of this distribution. Our data quality concern here is that individuals may be racing through the survey and not paying attention, or that they take a very long time completing the survey because they are not fully engaged and may be occupied with other tasks.

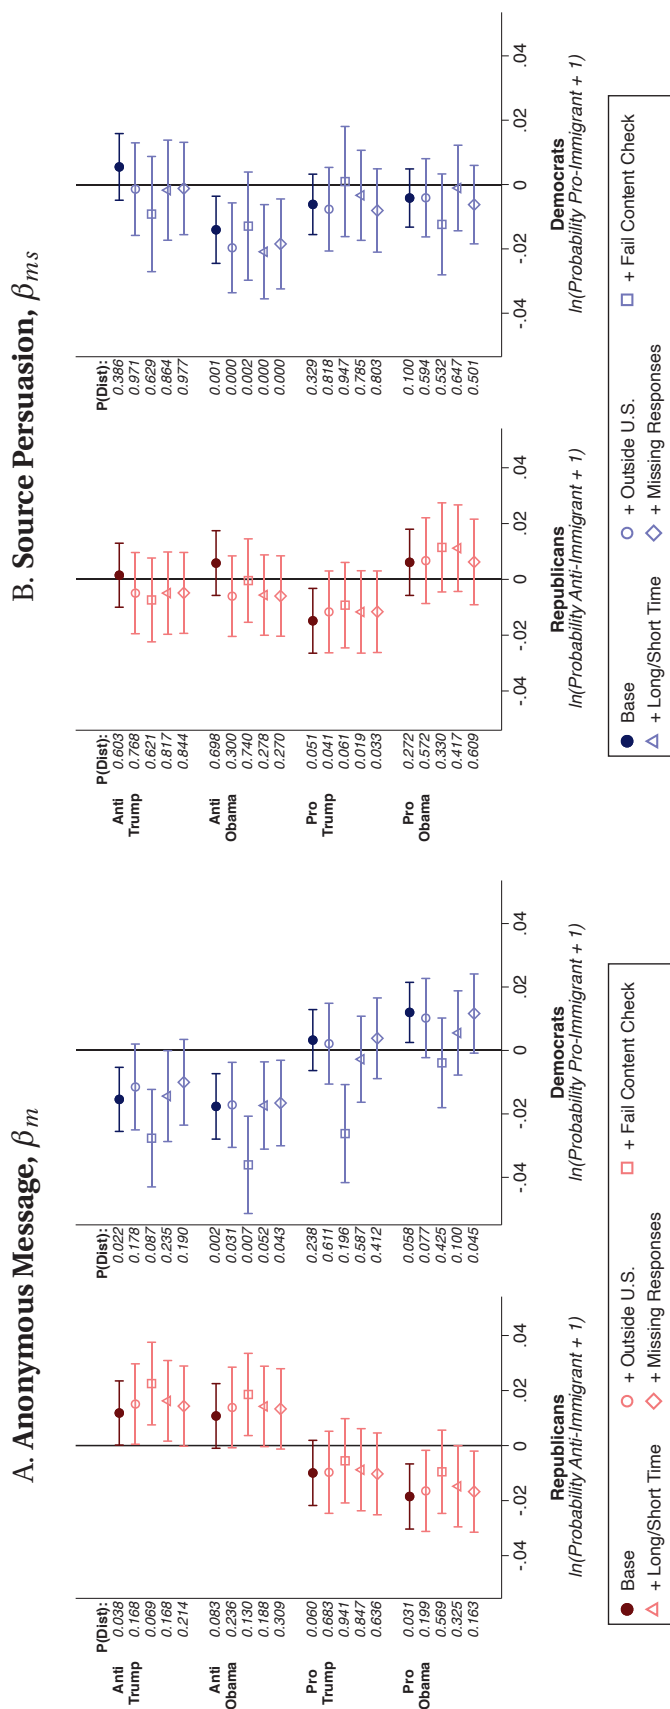
These first three additional quality restrictions are applied to individuals regardless of treatment assignment. The last restriction is an attention check on audio clips that individuals were assigned to in the survey (either immigration audio clips from the actor or president, or the

president turkey pardon clips). Here, we asked participants what topic was covered by the audio clip, and we remove them if they did not answer correctly. As such, this restriction cannot be applied to the “no audio” control group. Despite this potential imbalance, we chose to include this restriction because individuals who could not answer the minimal attention check that an audio clip was about “immigration” (vs. a different policy area) were not actually effective participants in the study. This restriction applies to 467 individuals. However, it only applies to 297 individuals (2% of the sample after baseline restrictions) who would not otherwise be dropped by one of the other three sample restrictions described above (as there is overlap across these quality conditions).

Figure A10 displays the results when we sequentially add back individuals from these additional data quality restrictions, as compared to results in our final sample. For the last quality restriction of the attention check on treatment audio clips (“+Fail Content Check”), we see a movement towards an increase in anti-immigrant beliefs for the anonymous message effect on the Pro, Trump (Actor) clip for Democrats, in contrast to a null effect in the analysis sample. This effect is not significant as a distributional difference but is significant as a movement of the average (coefficient). We also see moderately larger movements in β_m towards anti-immigrant positions for anti-immigrant messages for Democrats in this group, though these accord with the direction and significance of the analysis sample. It is not surprising that we might get some differences for this group, as these respondents did not know that immigration was the topic of the audio treatment clip. However, we find it reassuring that this sample restriction does not change the pattern of results except in a single case (Anonymous Message; Democrats; Pro, Trump (Actor)).

For all other sample restrictions, the point estimates that add back the dropped data are quite close to our baseline estimates, though the precision is slightly lower when the restrictions are relaxed. This overall stability provides reassurance that our choice of restrictions is not driving the patterns we observe in the data.

Figure A10: Impact of Sample Construction Choices



Notes: This figure plots the central results of the paper, adjusting the sample to “add back” respondents that were dropped due to data quality concerns. The sample restrictions included correspond to the restrictions in yellow in Figure A9. “+Outside U.S.” adds back individuals who completed the survey outside of the U.S., based on IP address location. “+Fail Content Check” adds back individuals who listen to an audio clip within the survey and then do not correctly answer a question about the content of the audio clip afterwards. “+Long/Short Time” adds back individuals who take the survey very quickly or over a very long time period (< 4 or > 45 minutes). “+Missing Responses” adds back individuals who did not answer all of the pre-treatment questions in the survey.

A5 Heterogeneity

We designed our experiment to explicitly address differences in treatment effects along political party lines. Our primary results show starkly mirrored effects across political party affiliation. As such party affiliation appears is a key characteristic that predicts both prior beliefs about immigration as well as responses to different leaders.

In this appendix, we discuss analysis of the results along several additional dimensions of heterogeneity that are worth exploring within party group, and we examine these dimensions to the extent possible given our sample size and statistical power. Examining sub-groups of the data also permit insights about some of the decision models hypothesized to underly political belief formation that have been referenced in earlier work; namely, (1) motivated reasoning models, which predict that responses will vary by the strength of an individual's prior belief about immigration (or their party allegiance), and (2) dual-processing models, which predict that responses will vary by the attention or focus of the participant. Again, we focus on the results from Decomposition II for these sub-analyses.

First, we test the importance of the strength of party affiliation as well as the strength of prior beliefs about immigration in Figure A11. We do this by including results for alternate samples that exclude individuals who voted for the opposite party candidate in 2016 ("Drop Flip Voters"), exclude individuals who answered the pre-treatment immigration question by stating that immigration should be increased (decreased) who were Republicans (Democrats) ("Drop Anti-Party Views") and additionally exclude individuals who stated that immigration should remain the same ("Only Party Views"), exclude individuals who identified as independents within the survey ("Drop Independents"), exclude individuals who stated that they were not a fan of their own party president ("Drop Non-Fan President"), and lastly separately individuals who were predicted to have far right or left views on immigration using information on all pre-treatment questions in the survey ("Drop Polarized" and "Only Polarized").³² In a motivated reasoning model, we might expect movement to the opposing position to be more likely among individuals with lower levels of attachment to their party prior. However, across each of these sub-groups of the data, the estimates are nearly identical to the baseline sample that includes all participants. An exception is that the message effects tend to move participants leftward for polarized Republicans, a feature that could be a result of these individuals having extreme right priors before treatment that do not permit them to move rightward to the same degree as other groups. Notably, this outlier result contrasts with the predictions of a motivated reasoning model which would anticipate larger leftward movements for moderate Republicans. Overall, the stability across groups suggests that neither moderates nor extreme members of each party are driving the results. Rather, consistent with the plots in Figure A4, it appears that there is a full shift of the distribution as a result of treatment that is not stronger at either tail.

Next, we test whether differences in political engagement or news consumption matter in Figure A12. In a dual processing model, we might expect participants who are less engaged or attentive to be more likely to update their views in a direction that opposes their priors. Here, we re-estimate the models excluding individuals who consume news from at least two platforms

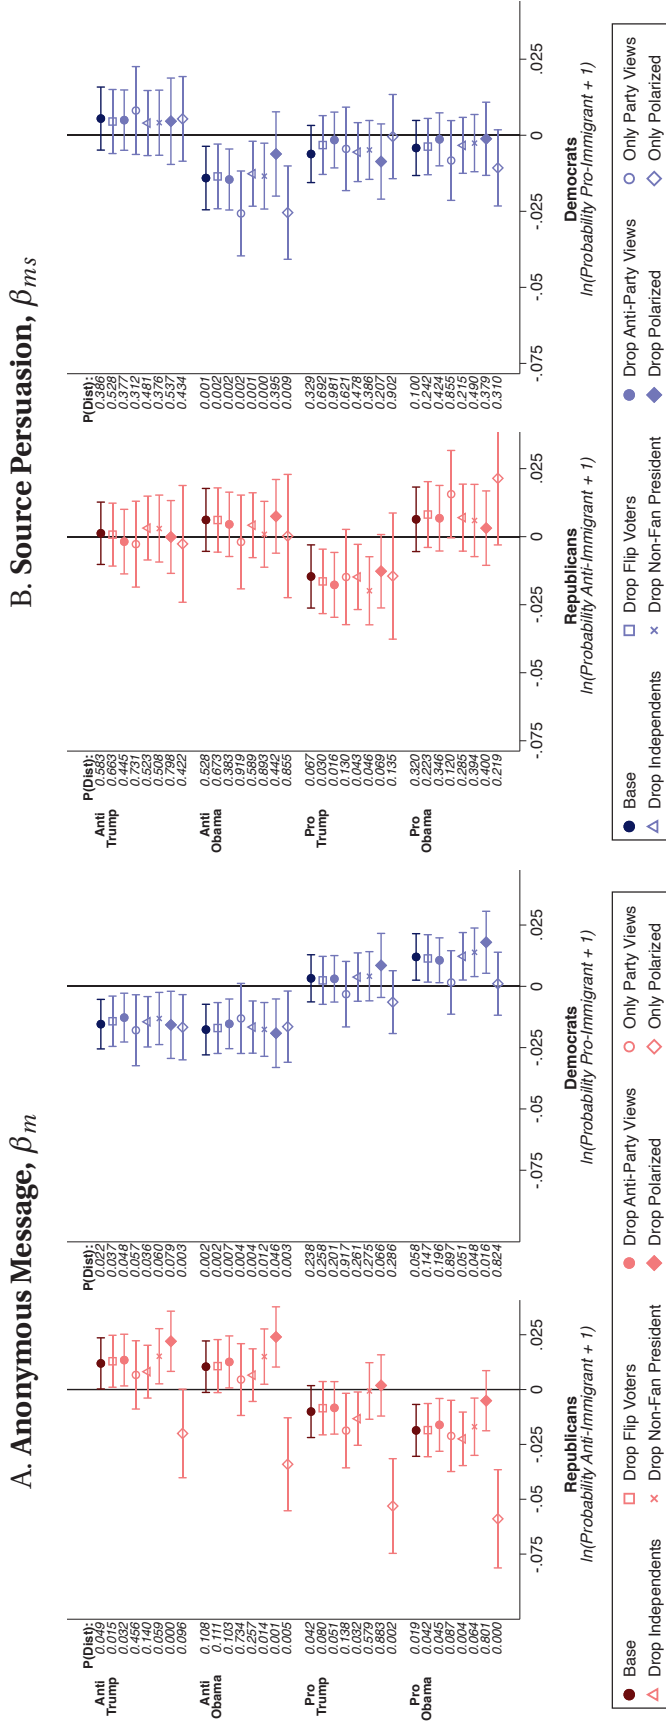
³²"Polarized" is an estimated characteristic assigned to participants in the sample. We do this by predicting whether an individual has views in the top quartile of the party position (pro-immigrant for Democrats, anti-immigrant for Republicans) using only the control group and pre-treatment question responses. We use the coefficients from the estimation to extrapolate who would be polarized (without treatment) for the whole sample.

(Newspaper, TV, Twitter, Facebook) and at least daily (“Multiple News Types”), individuals who consume news from both a left-leaning and a right-leaning source at least weekly (“Bipartisan News”), and individuals who did not vote in the 2016 election (“Non-Voters (2016)”). Again, the estimates are remarkably stable across groups, suggesting that neither highly informed/engaged nor uninformed/unengaged participants are driving the findings of the study.

Lastly, we explore a number of demographic dimensions of heterogeneity in Figures A13 and A14. Overall, we find very few notable differences in the estimates according to gender, race, employment status, or educational attainment. The confidence intervals of nearly all of the comparison estimates overlap; and we are hesitant to draw meaningful conclusions from suggestive differences in point estimates given limitations due to sample size.

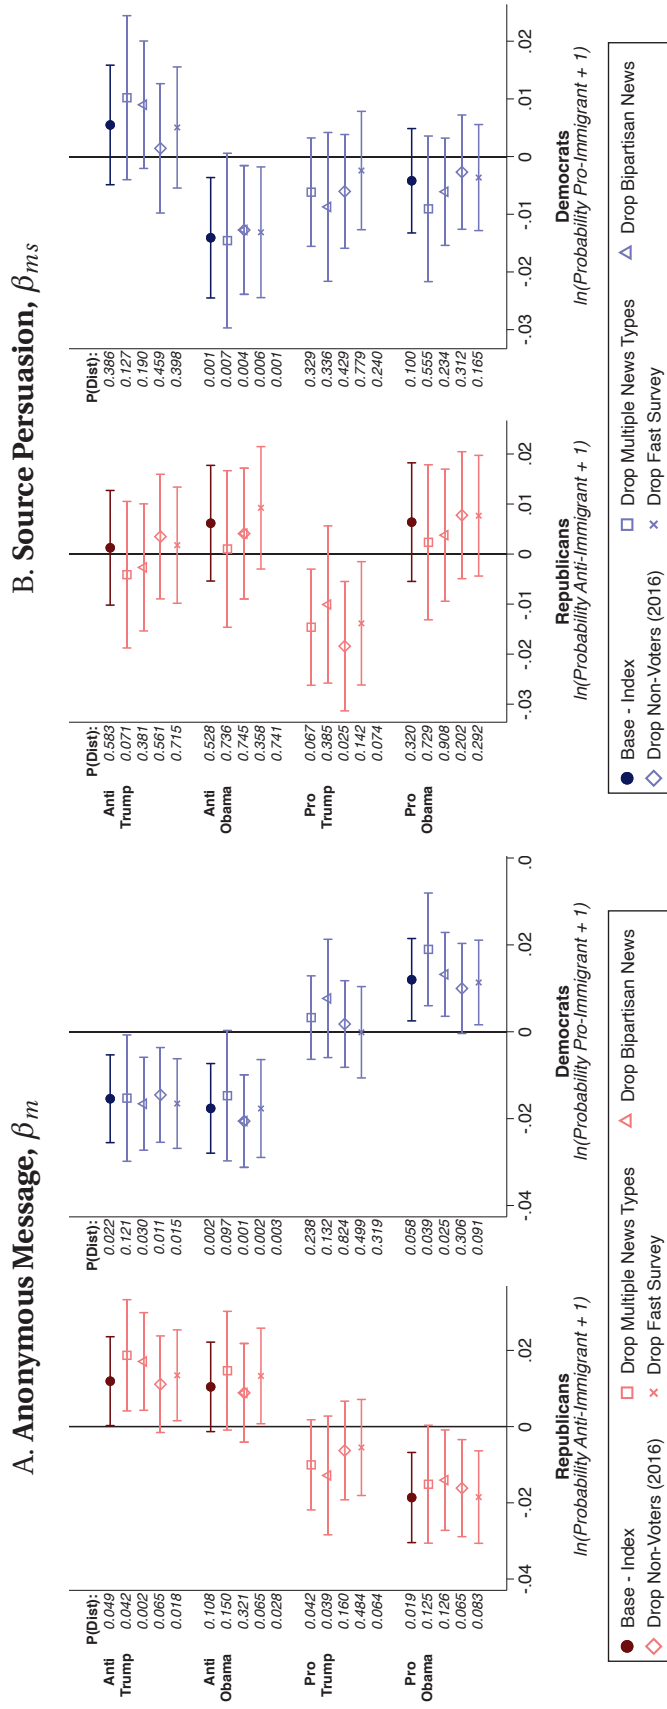
Across each of these characteristics and dimensions, we find limited variation in our estimated effects. This result could be attributable to limited power to identify these subgroup differences, as our experiment was not powered to detect subgroup effects. However, we note that while our estimates for some subgroups are less precise, the point estimates for these effects are quite consistent with the full sample estimate. This consistency differs from prior work in political science using motivated reasoning models, which posits that individuals with stronger party affiliation will have stronger effects, or dual processing models, which posits that individuals who are less informed about policy issues will have stronger effects (Bullock, 2020). Instead, it is the shifts in beliefs that we observe are present for all segments of the distribution within party. Our results support the hypothesis that *party affiliation* is the most important factor that determines responsiveness to treatment.

Figure A11: Heterogeneity, Moderates and Partisan Respondents




Notes: This plot shows the estimates for the baseline outcomes as compared to samples that remove either moderates or extremists. The outcome is $\ln(P(\text{Outcome} + 1))$. “Drop Flip Voters” removes Republicans who voted for Hillary Clinton in 2016 and Democrats who voted for Donald Trump in 2016. “Drop Anti-Party Views” removes Republicans (Democrats) who answered a pre-treatment question saying that they think immigration should be expanded (restricted). “Drop Independents” drops individuals who were recruited as either Democrats or Republicans and then later stated that they were Independents within the survey. “Drop Non-Fan President” excludes individuals who are not fans of the president from their party. “Drop Polarized” excludes those who are predicted to have extreme immigration views using a probit predictive model with pre-treatment characteristics estimated on the control group and extrapolated to the rest of the sample; “Only Polarized” isolates this group. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A12: Heterogeneity, Informed and Engaged Respondents




Notes: This plot shows the estimates for the baseline outcomes as compared to samples that remove engaged/informed or non-engaged/non-informed participants. The outcome is $\ln(P(Outcome + 1))$. “Drop Multiple News Types” removes participants who consume news at least daily through more than one mode: Newspaper, TV, Twitter and/or Facebook. “Drop Bi-Partisan News” removes participants who consume news from both a right-leaning and left-leaning news outlet at least weekly. “Drop Non-Voters (2016)” removes individuals who did not vote in the 2016 presidential election. “Drop Fast Survey” removes individuals who took the survey in 5 minutes or less, adjusted for the time span of audio treatments within the survey. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A13: Heterogeneity, Demographics



Notes: This plot shows the estimates for demographic sub-groups of the party samples. Note that the demographics of the party groups differs by design, to approximate the demographics of each political party. The outcome is $\ln(P(\text{Outcome} + 1))$. All regressions are adjusted for covariates in Appendix A3. "P(Dist)" is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A14: Heterogeneity, Demographics (Continued)



Notes: This plot shows the estimates for demographic sub-groups of the party samples. Note that the demographics of the party groups differs by design, to approximate the demographics of each political party. The outcome is $\ln(P(\text{Outcome} + 1))$. All regressions are adjusted for covariates in Appendix A3. "P(Dist)" is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

A6 Results By Component Questions

Our central results rely on an index composed of 16 different questions to incorporate multiple dimensions of immigration attitudes. Here, we test the robustness of the index to its components, again focusing on the results from the central Persuasion. First, we consider the sensitivity of the index by re-calculating the estimates leaving out one question at a time in Figure A15 and by estimating the models separately for each question in Figure A16. While the results are more disperse when separately estimated by question, across both sets of tests the findings are quite consistent with the baseline index outcome.

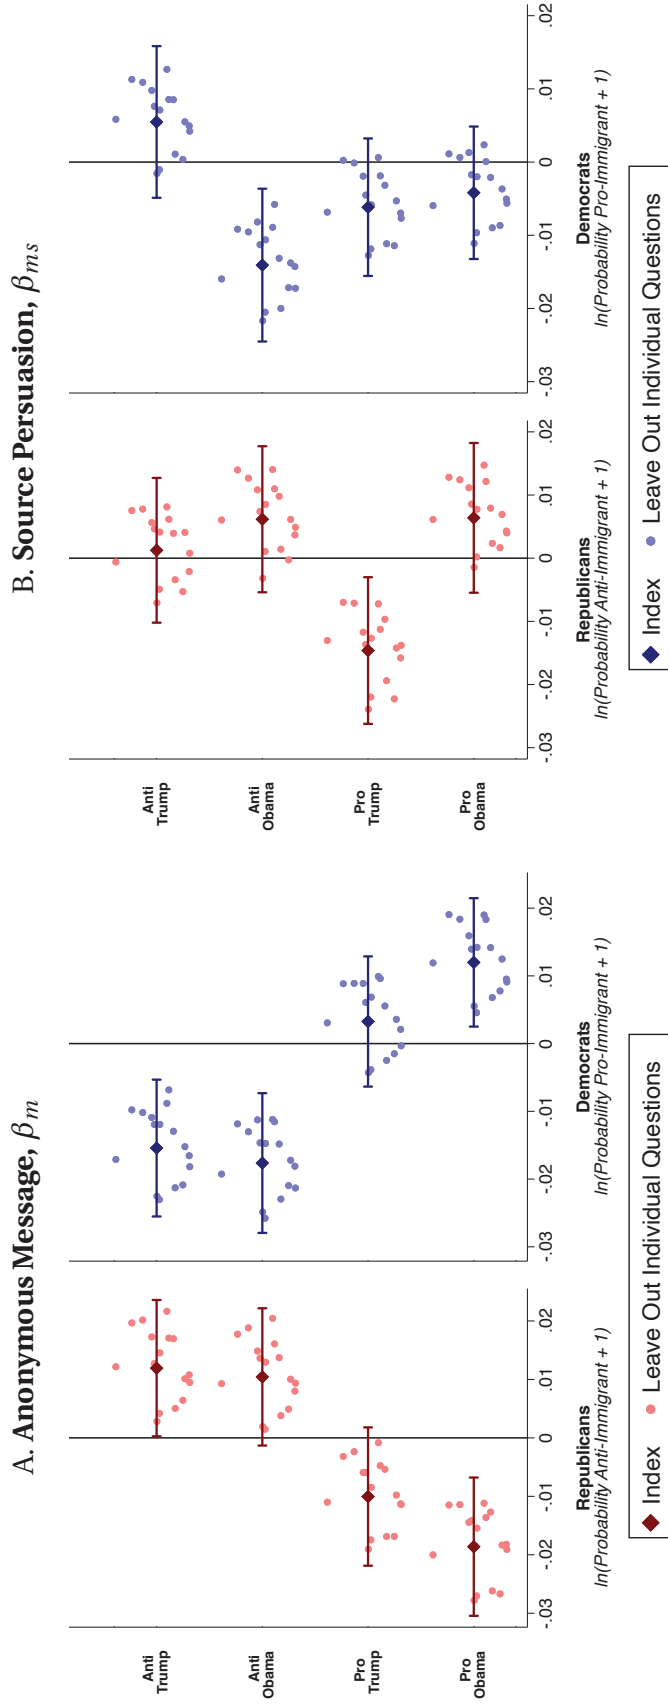
Figures A17 and A18 display the results for each individual question. The results are noisier when estimated at the individual question level, but some suggestive patterns emerge. Focusing on the source persuasion effects, Republicans become more pro-immigrant (for the Trump pro-immigrant treatment) on questions related to their “overall view” of immigrants, providing a path to citizenship for immigrants, especially for immigrants who came to the U.S. illegally as children (“Dreamers” or DACA recipients), and their views on whether all immigrants residing in the U.S. illegally should be deported. Arguably, these questions capture aspects of the immigration debate that are the most modest, or most open to support from right-leaning individuals.

For Democrats, we observe a pattern in the source persuasion effect whereby partisans are more likely to become anti-immigrant on issues that penalize immigrants who may be viewed as the most egregious rule-breakers. Here, we observe anti-immigrant movement on questions related to the normative societal contribution of immigrants living in the U.S. illegally, the economic contribution of immigrants, interest in deporting immigrants living in the U.S. illegally, especially those with a criminal record, the expansion of immigration background checks for workers, and increases in funding for U.S. border patrol agents. This pattern of findings suggests that Democrats may be prone to adopting anti-immigrant attitudes towards immigrants that may be more likely to be viewed as the most severe “bad actors.”

Alternatively, in Figure A19, we split the questions into two sub-indices focused on questions relating to immigration policies vs. beliefs about immigrants. The policy index includes 7 questions on the following topics (using shorthand from prior figure): Dreamer Path Citizen, Path Citizen, Deport Crime, Deport All, Worker Status, Border Patrol, Border Wall. The belief index includes the 9 remaining questions. Here, the questions are not split by modesty of proposed policies or types of immigrants who may be viewed as “rule-breakers” or “bad actors”, and consequently, we find little difference in results across these two sub-indices.

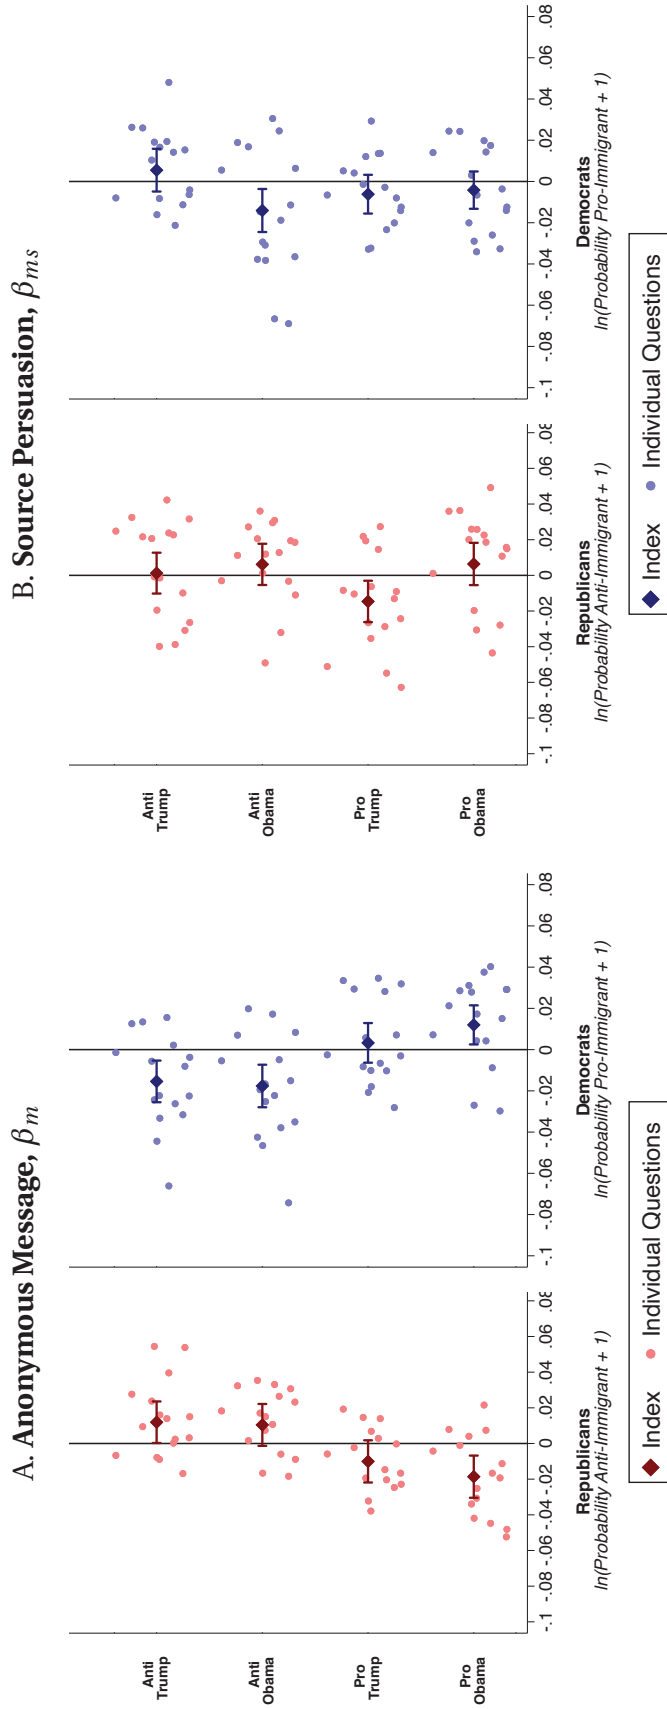
The “by question” results also highlight a strength of the study design: the fact that we are able to combine responses of multiple dimensions of immigration attitudes into a summary index. Immigration is a complex policy issue, and partisan voters may change their beliefs on some aspects of immigration and not others in response to political signals. Using information from multiple varied questions increases the total precision of our outcome. If we had alternatively only considered 1-2 outcome questions with a general focus (the approach of earlier work), we may not have found any evidence for the robust and symmetric effects that we observe for both parties in our study.

Figure A15: Leave Out Question Distribution




Notes: This figure overlays the baseline estimates of each test using the full index on top of 16 separate versions of each outcome that each leave out a single question from the index. All outcomes are measured using $\ln(P(y) + 1)$ and coefficients represent changes in this transformed outcome. All regressions are adjusted for covariates in Appendix A3. “p(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A16: By Question Distribution




Notes: This figure overlays the baseline estimates of each test using the full index on top of the outcome for each of the 16 questions that comprises the index. All outcomes are measured using $\ln(P(y) + 1)$ and coefficients represent changes in this transformed outcome. Each individual question has the same scale as the index and can assume values of $P(y) \in [0, 0.5, 1]$. All regressions are adjusted for covariates in Appendix A3. "p(Dist)" is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A17: Results by Question



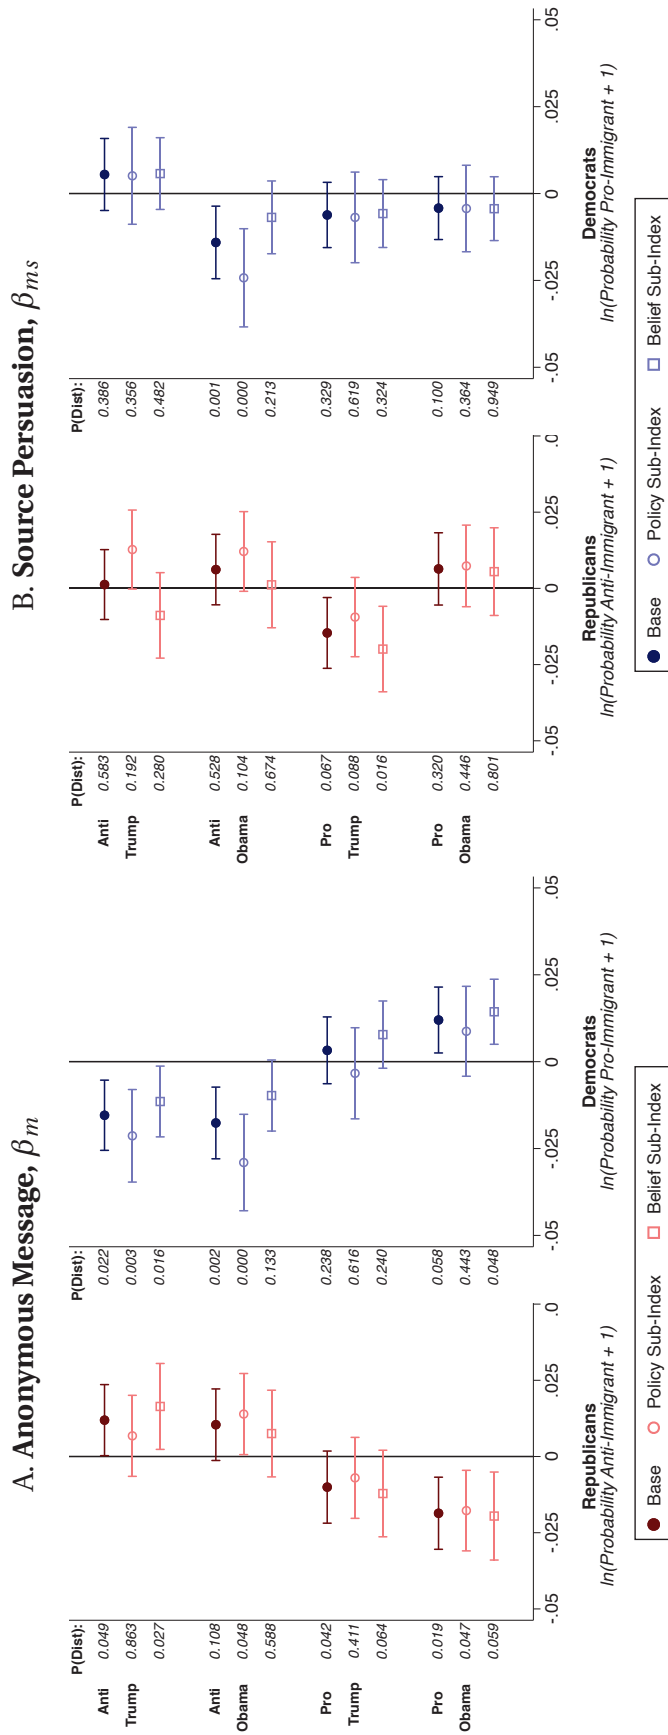
Notes: This figure displays the baseline estimates of each test using the full index along with results for each of the 16 questions that comprises the index. All outcomes are measured using $\ln(P(y) + 1)$ and coefficients represent changes in this transformed outcome. Each individual question has the same scale as the index and can assume values of $P(y) \in \{0, 0.5, 1\}$. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A18: Results by Question (Continued)



Notes: This figure displays the baseline estimates of each test using the full index along with results for each of the 16 questions that comprises the index. All outcomes are measured using $\ln(P(y) + 1)$ and coefficients represent changes in this transformed outcome. Each individual question has the same scale as the index and can assume values of $P(y) \in \{0, 0.5, 1\}$. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

Figure A19: Sub-Indices for Questions about Immigration Policy vs. Beliefs



Notes: This figure displays the baseline estimates of each test using the full index as well as alternative indices that split questions based on whether they refer to immigration policy positions or immigration beliefs. The policy index includes 7 questions on the following topics (using shorthand from prior figure): Dreamer Path Citizen, Path Citizen, Deport Crime, Deport All, Worker Status, Border Patrol, Border Wall. The belief index includes the 9 remaining questions. All outcomes are measured using $\ln(P(y) + 1)$ and coefficients represent changes in this transformed outcome. Each individual question has the same scale as the index and can assume values of $P(y) \in \{0, 0.5, 1\}$. All regressions are adjusted for covariates in Appendix A3. “P(Dist)” is the p-value from a Kolmogorov-Smirnov test of equality of distributions between the treatment and comparison groups for each estimate.

A7 Out of Sample Outcomes Measured in Survey

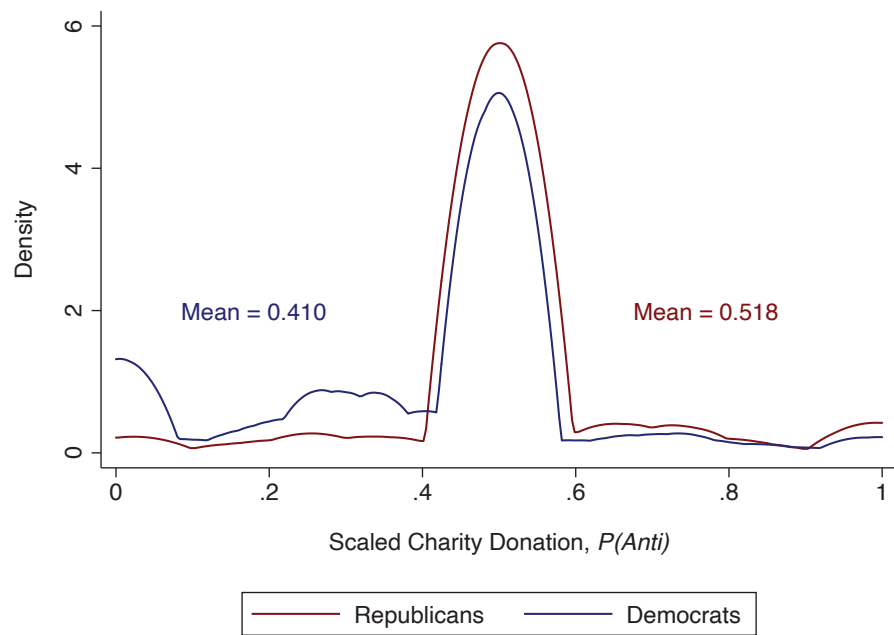
In addition to collecting information on immigrant views, we also attempted to measure two out-of-sample outcomes as a way to capture how actions might respond to a change in immigration beliefs. First, we asked participants about their intended vote in the 2020 presidential election, where we interpret a vote for Donald Trump as an anti-immigrant outcome and a vote for Joseph Biden as a pro-immigrant outcome. Second, we asked whether the participant would be interested in donating money to an anti-immigrant (Federation for American Immigration Reform (FAIR)) or a pro-immigrant charity (American Civil Liberties Union (ACLU) Immigrants' Rights Project), out of potential lottery earnings, if the participant wins a lottery that was conducted as part of the survey. The lottery offered a potential prize of \$25 payable directly to the winner. If the winner chooses to donate to one of the two charities, she could choose the amount to be donated of the earnings. Any donation would be doubled before donation, such that the maximum donation amount was \$50.

Figure A21 shows the results for these two outcomes as compared to the baseline immigration index outcomes. Unfortunately, the results here are noisy, imprecise and cannot reject the null hypothesis of no effect for either outcome. In cases where there is a significant change in the coefficient (or mean difference across groups), there is no corresponding significant difference in the distributions using the Kolmogorov-Smirnov test, and vice versa.

Multiple factors may contribute to the lack of effects we find for these outcomes. First, we may fail to find effects for these outcomes could be due to lack of precision and limited power. In addition, our ability to measure the voting outcome is limited by the fact that many participants may have voted prior to taking the survey, through early or on-time voting at polling places. While we intended to complete the study prior to the election on November 3, 2020, our survey recruitment period took longer than we had anticipated and continued past the presidential election to November 10th. We need to exclude the data collected on or after the election for the voting outcome, which also limits the sample size for this test.

For the charity donation outcome, our total variation in the outcome is quite limited. Figure A20 plots the charity donation outcome in the control group of the study, where this outcome is normalized to be between 0 and 1. Here, 0 corresponds to a choice to donate the entire potential prize to the ACLU (pro-immigrant charity), 1 corresponds to the choice to donate the entire potential prize to FAIR (anti-immigrant charity), and 0.5 corresponds to the choice not to donate to either charity. Nearly all survey participants chose to award themselves the money (if they were to win). One reason that we may not find frequent responses to donate to one of the charities is that the prize may not be big enough, a constraint that we faced given our research budget. Effectively, the lack of variation we see in this outcome contributes to our limited ability to observe meaningful or significant changes in the charity donation choice after different treatments. It is not surprising that we find noisy and imprecise effects for this outcome given that nearly all participants chose to give themselves the prize, regardless of treatment status.

Figure A20: Charity Donation Outcome in No Audio Control Group



Notes: This plot shows the baseline distribution of the charity donation outcome in the no audio control group. The charity outcome is measured from a question where participants are told that they have been automatically entered into a lottery for \$25, of which they can donate a portion to either an anti-immigrant charity, the Federation for American Immigration Reform (FAIR), or a pro-immigrant charity, the American Civil Liberties Union (ACLU) Immigrants' Rights Project, at a rate that will double any dollar that is donated. This outcome re-scales the donation choice based on these dollar donations to have a scale of [0, 1]. A value of 0.5 corresponds to the participant electing not to make any donation.

Figure A21: Out of Sample Outcomes

