

Persuasion in Parallel

Alexander Coppock

Draft, October 2021

© Copyright by Alexander Coppock 2021. All rights reserved.

Contents

1	Persuasion in Polarized America	1
1	The persuasion in parallel hypothesis	3
2	Example: Flat tax op-ed experiment	7
3	Parallel publics	10
4	What's at stake	16
5	Where we're headed	17
2	Reinterpreting a social psychology classic	18
1	The original Lord, Ross, and Lepper (1979) study	20
1.1	What went wrong?	21
2	The Guess and Coppock (2020) replication study	23
2.1	Replication study results	25
2.2	Biased assimilation	28
3	Summary	30
3	Definitions and distinctions	32
1	Treatment effect	32
2	Persuasive information	36
2.1	Group cues are not persuasive information	39
3	Policy attitudes and beliefs	44
3.1	Affective evaluations of messages and messengers	49
4	Positive	50
5	Small	50
6	Durable	51
7	Everyone	52
8	Summary	52
4	Research design	55
1	Model	57
2	Inquiry	60
2.1	What about mediated effects?	62
2.2	What about differences-in-CATEs?	63
2.3	What does parallel even mean?	64
2.4	What about hidden moderators?	67
3	Data strategy	69
3.1	Sampling of subjects	69
3.2	Assignment of information	72
3.3	Measurement of policy attitudes	75
4	Answer strategy	76

5	Summary	77
5	Persuasion experiments: originals, replications, and reanalyses	78
1	Three original persuasion experiments	79
1.1	Original study 1: Guess and Coppock (2020): Minimum wage	79
1.2	Original study 2: Guess and Coppock (2020): Gun control	80
1.3	Original study 3: Coppock, Ekins and Kirby (2018): Newspaper op-eds	82
2	Reanalysis and replication of five persuasion experiments	84
2.1	Reanalysis and replication 1: Chong and Druckman (2010): Patriot Act	86
2.2	Reanalysis and replication 2: Brader, Valentino and Suhay (2008): Immigration frames	87
2.3	Reanalysis and replication 3: Hiscox (2006): Free trade	88
2.4	Reanalysis and replication 4: Johnston and Ballard (2016): Expert economists	91
2.5	Reanalysis and replication 5: Hopkins and Mummolo (2017): Frame breadth	91
3	Reanalyses of five TESS studies	94
3.1	Reanalysis 1: Gash and Murakami (2009): Venue effects	96
3.2	Reanalysis 2: (Flavin, 2011): Political equality	98
3.3	Reanalysis 3: Kreps and Wallace (2016): Drone strikes	100
3.4	Reanalysis 4: (Mutz, 2017): Job loss	101
3.5	Reanalysis 5: Trump and White (2018): Income inequality	102
4	Meta-analysis	102
5	Two-sided messages	104
5.1	Two-sided messages 1: Guess and Coppock (2020) redesign of Lord, Ross and Lepper (1979): Capital punishment	107
5.2	Two-sided messages 2: Chong and Druckman (2010): Patriot Act	109
5.3	Two-sided messages 3: Hiscox (2006): Free trade	109
5.4	Two-sided messages summary	112
6	Summary	112
6	Persistence and decay	114
1	Extending the research design through time	116
1.1	Changes to the model	116
1.2	Changes to the inquiry	117
1.3	Changes to the data strategy	118
1.4	Changes to the answer strategy	120
2	Results	121
3	Summary	127
7	Models of information processing	130
1	Bayesian reasoning	132
1.1	Empirical assessments of Bayesian reasoning	133
2	Motivated reasoning	138
2.1	Bayesian interpretation of biased assimilation	140
2.2	Discussion and critique of Taber and Lodge (2006)	143
3	Summary	146
8	Persuasion is Possible	149

9	Appendix	152
1	Newspapers (Coppock, Ekins and Kirby, 2018)	153
2	Gun control (Guess and Coppock, 2020)	161
3	Minimum wage: (Guess and Coppock, 2020)	164
4	Replication and extension: (Lord, Ross and Lepper, 1979)	166
5	Reanalysis and replication 1: (Chong and Druckman, 2010)	170
6	Reanalysis and replication 2: (Brader, Valentino and Suhay, 2008)	172
7	Reanalysis and replication 3: (Hiscox, 2006)	176
8	Reanalysis and replication 4: (Johnston and Ballard, 2016)	178
9	Reanalysis and replication 5: (Hopkins and Mummolo, 2017)	180
10	Reanalysis 1: Gash and Murakami (2009)	183
11	Reanalysis 2: Flavin (2011)	185
12	Reanalysis 3: Kreps and Wallace (2016)	186
13	Reanalysis 4: Mutz (2017)	188
14	Reanalysis 5: Trump and White (2018)	190
10	Acknowledgements	191
	References	192

List of Tables

2.1	Reproduction of Lord, Ross, and Lepper (1979), Table 3	21
2.2	Reanalysis of Guess and Coppock (2020) à la Lord, Ross, and Lepper (1979)	26
2.3	Biased assimilation in the original and replication studies	29
3.1	Distinctions across three inquiries	35
3.2	Summary of predictions	54
4.1	Possible types when measurements are binary	65
4.2	Sample sources and sizes for 23 persuasion survey experiments	70
6.1	Persistence predictions for 12 panel survey experiments	120
6.2	Average persistence ratio estimates	124
6.3	Op-ed experiment persistence estimates	126
6.4	All persistence estimates	129
9.1	Coppock, Ekins and Kirby (2018): Treatment effect estimates	160
9.2	Guess and Coppock (2020): Gun control treatment effect estimates	163
9.3	Guess and Coppock (2020): Minimum wage treatment videos	164
9.4	Guess and Coppock (2020): Minimum wage treatment effect estimates	165
9.5	Guess and Coppock (2020): Capital punishment treatment effect estimates	169
9.6	Chong and Druckman (2010): Treatment effect estimates	171
9.7	Brader, Valentino and Suhay (2008): Treatment effect estimates	175
9.8	Hiscox (2006): Treatment effect estimates	177
9.9	Johnston and Ballard (2016): Treatment effect estimates	179
9.10	Hopkins and Mummolo (2017): Treatment effect estimates	182
9.11	Gash and Murakami (2009): Treatments and outcomes	183
9.12	Gash and Murakami (2009): Treatment effect estimates	184
9.13	Flavin (2011): Treatments and outcomes	185
9.14	Flavin (2011): Treatment effect estimates	185
9.15	Kreps and Wallace (2016): Treatments and outcomes	186
9.16	Kreps and Wallace (2016): Treatment effect estimates	187
9.17	Mutz (2017): Treatments and outcomes	188
9.18	Mutz (2017): Treatment effect estimates	189
9.19	Trump and White (2018): Treatment effect estimates	190

List of Figures

- 1.1 The persuasion in parallel hypothesis 4
- 1.2 Flat tax experimental results 9
- 1.3 Flat tax experimental results: 10-day follow-up 11
- 1.4 American attitudes toward gay marriage 2001-2019 13
- 1.5 American attitudes toward legal abortion by circumstance 1977-2018 15

- 2.1 Graphical treatments used in Guess and Coppock (2020) replication of Lord, Ross, and Lepper (1979) 25
- 2.2 Reanalysis of Guess and Coppock (2020): Self-assessed outcome measure 27
- 2.3 Note: Data from 682 Mechanical Turk respondents. Proponents’ responses are plotted with the letter P and opponents’ responses with the letter O. 27
- 2.4 Reanalysis of Guess and Coppock (2020): Direct outcome measure 28

- 3.1 Effects of persuasive information on target and non-target attitudes 38
- 3.2 Reanalysis of 1982 California poll party cue experiment 41
- 3.3 Reanalysis and replication of Nicholson (2012) 42
- 3.4 Illustration of the attitudes model 48

- 4.1 Model of persuasive information and policy attitudes in the panel survey experiment design 58
- 4.2 CATEs on measured outcomes can be different even if CATEs on latent outcomes are equal 66

- 5.1 Reanalysis of Guess and Coppock (2020): Minimum wage 81
- 5.2 Reanalysis of Guess and Coppock (2020): Gun control 83
- 5.3 Reanalysis of (Coppock, Ekins and Kirby, 2018): Newspaper op-eds 85
- 5.4 Reanalysis and replication of Chong and Druckman (2010): Patriot Act 87
- 5.5 Reanalysis and replication of Brader, Valentino and Suhay (2008): Immigration frames 89
- 5.6 Reanalysis and replication of Hiscox (2006): Expert treatment 90
- 5.7 Reanalysis and replication of Hiscox (2006): Valence treatments 92
- 5.8 Reanalysis and replication of Johnston and Ballard (2016): Expert economists 93
- 5.9 Reanalysis and Replications of Hopkins and Mummolo (2017): Frame breadth 95
- 5.10 Reanalysis of Gash and Murakami (2009): Venue effects 97
- 5.11 Reanalysis of (Flavin, 2011): Political equality 99
- 5.12 Reanalysis of Kreps and Wallace (2016): Drone strikes 100
- 5.13 Reanalysis of (Mutz, 2017): Job loss 101
- 5.14 Treatments in Trump and White (2018) 102
- 5.15 Reanalysis of Trump and White (2018): Income inequality 103
- 5.16 Correlation of CATEs across demographic subgroups 105

5.17	Two-sided messages in Guess and Coppock (2020) redesign of Lord, Ross and Lepper (1979): Capital punishment	108
5.18	Two-sided messages 2: Chong and Druckman (2010): Patriot Act	110
5.19	Two-sided messages 3: Hiscox (2006): Free trade	111
6.1	Overtime ATE estimates	123
6.2	Persistence estimates	124
6.3	Newspaper op-eds overtime treatment effect estimates	125
7.1	Treatment effects of information depending on priors and likelihood functions . . .	137
9.1	Minimum wage treatment videos	164
9.4	Trump and White (2018): Treatments	190

CHAPTER 1

Persuasion in Polarized America

In dismay over the deep disagreements that divide Americans by party, by race, by class, by generation, by geography, we sometimes give in to the idea that at least the other side is stupid. At least the other side is full of irrational idiots who fail to see reason. They never learn from new information, never change their minds, never admit when they're wrong. If anything, trying to convince them on the basis of evidence and logic makes things worse and we become even more polarized. Each successive day seems to widen the gulf between Democrats and Republicans, between progressives and conservatives, between the haves and have nots, between the white supremacists and the rest of us, between we who are reasonable and they who are unreasonable.

It's a tempting feeling. When we see the unbelievable political positions some people support, the notion comes naturally that their ability to reason must be broken. When others come to different conclusions from us, despite having access to the same information we do, we think there must be something wrong with their information processing, or how they learn from the world to inform their policy attitudes.

In the academic literature, this idea goes by the name "motivated reasoning." A number of variants of the theory exist, but most posit that individuals are motivated by both accuracy goals (wanting to get it right) and directional goals (wanting to arrive at a particular conclusion). When directional goals dominate, people distort their interpretation of new information, always learning they were more right than they knew, regardless of what the new information shows. When people are motivated to reason to arrive at preferred conclusions, the effect of evidence on attitudes is "heads I'm right, tails you're wrong," because no matter how the evidentiary coin comes up, it is twisted toward directional goals. Under this theory, our political opponents' capacities for reason actually are broken, since everyone's are.

A key implication of motivated reasoning theory is backlash: exposure to counterattitu-

dinal evidence will cause people hold more strongly to their pre-existing positions. Festinger, Riecken and Schachter (1956) give a forceful articulation of this idea in the introduction to *When Prophecy Fails*, their study of cultists after a predicted doomsday fails to materialize:

“Suppose an individual believes something with his whole heart; suppose further that he has a commitment to this belief, that he has taken irrevocable actions because of it; finally, suppose that he is presented with evidence, unequivocal and undeniable evidence, that his belief is wrong: what will happen? The individual will frequently emerge, not only unshaken, but even more convinced of the truth of his beliefs than ever before. (p. 3)”

Two decades later, Lord, Ross and Lepper (1979) claim to have demonstrated this sort of “attitude polarization” in the laboratory. They conclude:

“If our study demonstrates anything, it surely demonstrates that social scientists can not expect rationality, enlightenment, and consensus about policy to emerge from their attempts to furnish “objective” data about burning social issues. If people of opposing views can each find support for those views in the same body of evidence, it is small wonder that social science research, dealing with complex and emotional social issues and forced to rely upon inconclusive designs, measures, and modes of analysis, will frequently fuel rather than calm the fires of debate. (p. 2108)”

Motivated reasoning theory was introduced into political science by Taber and Lodge (2006), who also purport to show that people respond to information by updating their views in opposite directions. They write:

“Our studies show people are often unable to escape the pull of their prior attitudes and beliefs, which guide the processing of new information in predictable and sometimes insidious ways. (p. 767)”

With this book, I hope to convince you that this idea from motivated reasoning theory simply does not describe how people respond when presented with persuasive information. On the contrary, people update in the direction of information, by a small amount. I will show how the evidence in favor of the most dire prediction of motivated reasoning – backlash – rests on weak research designs. When these designs are strengthened, the conclusions flip. The idea that trying to persuade the other side is counterproductive may *feel right*, but the goal of this book is to demonstrate that it *is wrong*.

1. *The persuasion in parallel hypothesis*

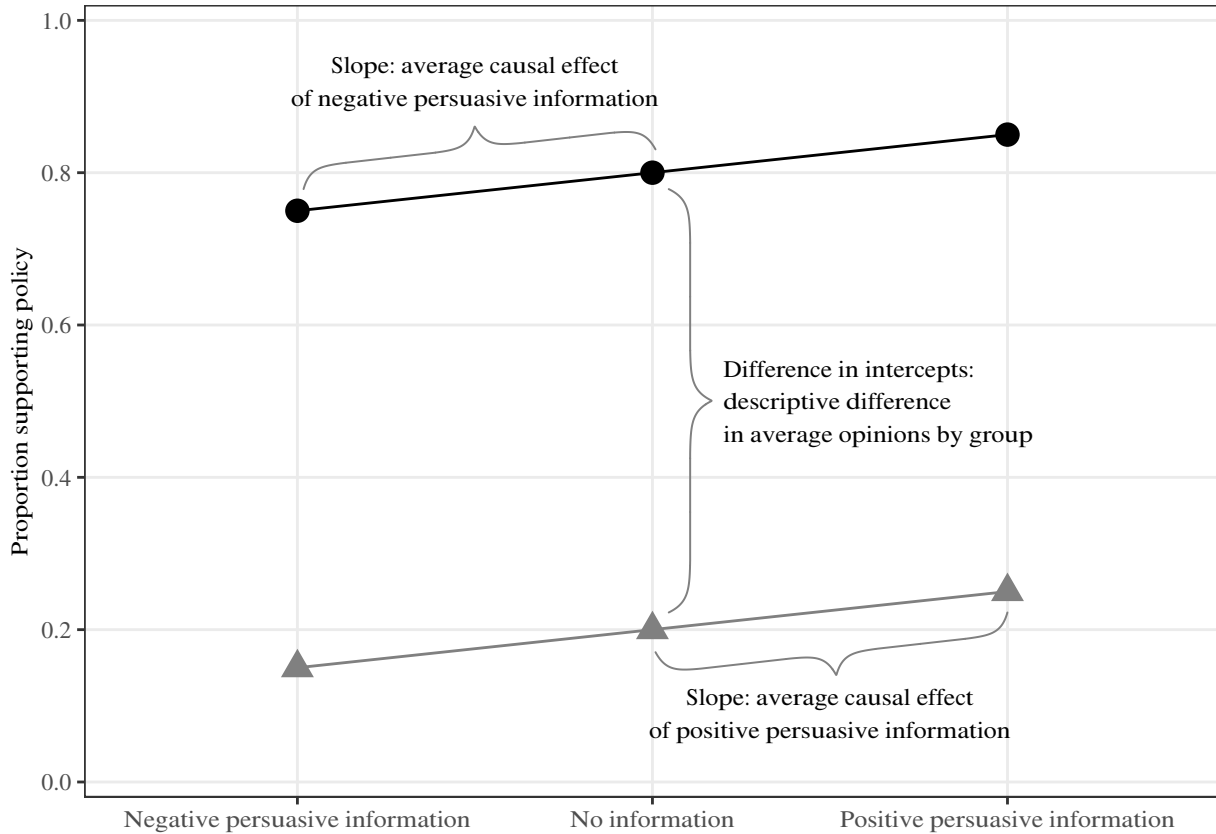
This book makes a single argument, over and over: Persuasion occurs in parallel. Using evidence gleaned from many randomized experiments, I will show that when people encounter new information, they don't distort it to further entrench their pre-existing views – instead, they are persuaded in the direction of that information. If persuasion occurs in parallel, people from different groups respond to persuasive information in the same direction and by about the same amount. While baseline political views are very different from group to group, responses to information are quite similar. For example, when people encounter pro-immigration arguments, regardless of whether they are immigration opponents or proponents, they increase their support for immigration a little bit. The converse is also true. Anti-immigration arguments will decrease support among both immigration proponents and opponents just the same. We see this common pattern across dozens of policy issues and across many different subdivisions of Americans.

Throughout the book, I'll rely on the visual metaphor of parallel lines to describe this idea. Parallel lines have the same slope (changes with respect to the horizontal axis) but may have different intercepts (positions on the vertical axis). Figure 1.1 is a schematic representation of the persuasion in parallel hypothesis. The figure shows two parallel lines, one for the circle and one for the triangles group. The separation of people into circles and triangles groups stands in for any division of Americans into opposing subgroups: proponents and opponents of a policy, Republicans and Democrats, or young and old, to name a few. The difference in intercepts reflects the descriptive difference in the average opinions of the circle and triangle groups. This difference can be large, because groups in our polarized society can be bitterly divided over political issues, but this difference does *not* refer to supposedly motivated responses to information.¹ When I say that the circle and triangle groups are persuaded in parallel, I mean that the average causal effects of persuasive information are the same for both groups, *even though* the groups still disagree when exposed to positive or negative information.

The claim that persuasion always occurs in parallel for everyone, regardless of the content or provenance of the persuasive information, is obviously far too broad. I promise that impor-

¹The descriptive difference in average opinions likely has many causes, but the main contenders in my mind are the effects of group cues and differential exposure to positive and negative persuasive information. The average causal effects of persuasive information themselves are represented by the slopes of the lines.

Figure 1.1: The persuasion in parallel hypothesis



Note: This figure is a schematic representation of the persuasion in parallel hypothesis. The vertical axis represents support for a hypothetical policy and the horizontal axis represents exposure to different levels of a hypothetical treatment. The triangles and circles stand in for any politically salient subdivision of Americans into groups that disagree over a public policy. The descriptive difference between the groups is large but the average causal effects of information are the same for both groups.

tant caveats and conditions are coming for readers who forge on. In the meantime, I want to emphasize that although there may be exceptions, persuasion in parallel is the norm. It happens many millions of times a day as people scroll through social media feeds and talk with friends and coworkers about the news. Politicians, journalists, pundits, academics, and advertisers are constantly barraging people with persuasive attempts to vote for a candidate, to believe a fact, to behave a certain way, or to buy a product. These attempts are probably a little bit effective for most everyone who hears them.

Persuasive information encompasses a wide set of political communications. It refers to facts and arguments designed to move target attitudes in a particular direction. Claims about why we need particular policies, what they would do, and how much they would cost are all persuasive information. The persuasion in parallel hypothesis is that about persuasive information only, not other kinds of communication that lack a target or a direction.

Probably the largest class of political communication that does not count as persuasive information is group cues. Group cues are messages about which groups support which positions. Party cues – information that indicates where a political party stands on the issues – are prominent examples of group cues. Group cues can exert powerful influence over policy attitudes and whether the effect is positive or negative can depend on group membership. For example, learning that Republicans in Congress support a particular bill usually increases support among Republican survey respondents but decreases support among Democratic respondents. The distinction between group cues and persuasive information can sometimes be blurry because messages contain elements of both, but we will nevertheless keep these two kinds of communication theoretically distinct.

The amount of opinion change in response to persuasive information is usually small. Small means something like five percentage points or a tenth of a standard deviation in response to a treatment like an op-ed, a video advertisement, or a précis of a scientific finding. Small changes make sense. If persuasive effects were much bigger, wild swings in attitudes would be commonplace and people would be continually changing their minds depending on the latest advertisement they saw.

Persuasive effects decay. Ten days after people encounter persuasive information, average effects are about one-third to one-half their original magnitude. After ten days, we have only

limited evidence about whether they persist or fade. In one study only, I measured persistence after 30 days, finding a “hockey stick” pattern in which treatment effects had decayed somewhat by day 10 but no further by day 30.

The strongest evidence for the claim that people are persuaded in parallel derives from randomized experiments in which some people are exposed to information (the treatment group) while others are not (the control group). These experiments show over and over that the average treatment effects of information are positive. By positive, I mean that they are in the direction of information. These average causal effects also hold for subgroups: young and old, better and less-well educated, Republican and Democrat, black and white, women and men, we all respond in the direction of information by about the same amount.

The strong form of the persuasion in parallel hypothesis is that information has the exact same effect for everyone. Falsifying this hypothesis is trivially easy. All we would need is one statistical test that shows that effects are stronger for one group than another. The experiments described in this book offer occasional examples of such tests. When the average effect for Democrats is 3 percentage points and the average effect for Republicans is 5 percentage points, a sufficiently large experiment would declare these two average effects “statistically significantly different.” But these sometimes statistically significant differences are rarely politically significant.² Qualitatively speaking, the effects of treatment on policy attitudes are quite similar even across wildly diverse groups of people.

The weak form of the hypothesis is that *backlash* doesn’t occur. Backlash (or backfire – I don’t draw any distinction between the two terms) would occur if information had positive effects for some but negative effects for others. Falsifying this weaker hypothesis would also be easy. All we would need is one statistical test that finds evidence of a positive effect for one group but a negative effect for a different group. None of the many experiments reported in

²Understanding statistical significance involves working through a bizarre thought experiment. Suppose we stipulate a model of the world (called a null model) that the true difference is zero but that estimates bounce around this true zero because of sampling variability. Under that null model and using our current empirical strategy, the thought experiment asks, what’s the probability would we obtain an estimate as large as or larger than the one we did obtain? If the answer is “not very often, less than 5% of the time” the estimate is deemed statistically significant. But that probability depends on the details of the empirical strategy, especially sample size. Unless the null model happens to be exactly correct (which it almost certainly isn’t), then as sample size gets bigger, we’re almost guaranteed to find statistical significance. Annoyingly, with very large experiments, even substantively meaningless differences can be declared “significant.” The cure for this problem is to focus more on effect estimates and their precision and less on statistical significance.

this book measure any instances of backlash, but other authors have claimed to find them (e.g., Lazarsfeld, Berelson and Gaudet, 1944; Nyhan and Reifler, 2010; Schaffner and Roche, 2016; Zhou, 2016). Backlash has probably occurred, but it is definitely not the norm (Wood and Porter, 2018; Nyhan et al., 2019; Swire-Thompson, DeGutis and Lazer, 2020). In their book *False Alarm*, Porter and Wood (2020) randomize more than fifty factual corrections of misinformation and uncover exactly zero instances of backlash.

Even so, many theories accommodate and predict backlash, including the Receive-Accept-Sample model (Zaller, 1992), the John Q. Public model (Lodge and Taber, 2013), and the Cultural Cognition model (Kahan, 2012). The Receive-Accept-Sample model allows backlash through “countervailing resistance” to uncongenial communication flows (Zaller, 1992; , p. 122). The John Q. Public model includes the “affect transfer” and “affect contagion” postulates through which the negative affective evaluation of a communication will cause backlash effects on the attitudes the persuasive attempt was meant to change (Lodge and Taber, 2013; pp. 56-58). Under the Cultural Cognition model, people engage in “protective cognition” to defend their variously individualistic or communitarian cultural values against challenges from scientific information, thereby “reinforc[ing] their predispositions (Kahan, 2010; , p. 296).” These models may be useful for explaining other phenomena, but at a minimum, their predictions of backlash are (in my view) incorrect. My critique of these theories will be mostly empirical – the predicted backlash doesn’t materialize across many dozens of tests for it – but I will draw out a theoretical critique as well. In brief, my view is that these theories mistake affective evaluations of messages and messengers for the persuasive effects of those messages. Since many people don’t *like* counter-attitudinal messages, we might mistakenly think they “reject” them. But as we will see, people update their policy attitudes in the direction of information they like and information they don’t like just the same.

2. Example: Flat tax op-ed experiment

The evidence in favor of the persuasion in parallel hypothesis will come much later in the book, in Chapter 5. The reasons to trust in that evidence – the experimental designs – will be described in Chapter 4. And the details of the theoretical structure those designs depend on will be laid out in Chapter 3. Since that’s a lot to work through before getting to the punchline, this

section will describe one example of persuasion in parallel to tide us over.

Emily Ekins, David Kirby, and I ran an experiment to measure how much, if at all, an opinion piece in the *Wall Street Journal* changes minds on tax policy. The randomly-assigned treatment group read an op-ed by Senator and then-(at the time of the experiment)-presidential hopeful Rand Paul of Kentucky. In “Blow Up the Tax Code and Start Over” (Paul, 2015), Senator Paul proposed a 14.5% flat tax that he predicted would cause the economy to “roar.” Paul primarily argues for his flat tax proposal on fairness grounds. He anticipated the objections that the proposal is a giveaway to the rich (he’d close loopholes) and that the proposal would induce massive deficits (he’d balance the budget). He called the IRS a “rogue agency” and blamed Washington corruption for the convolutions of the tax code. The op-ed is 1,000 words long, makes a complicated argument, and demonizes relatively obscure bureaucrats most Americans wouldn’t have heard of. It’s insidery, a little punchy, and lacks strong evidence for its claims, which is to say, it’s a good example of contemporary political communication.

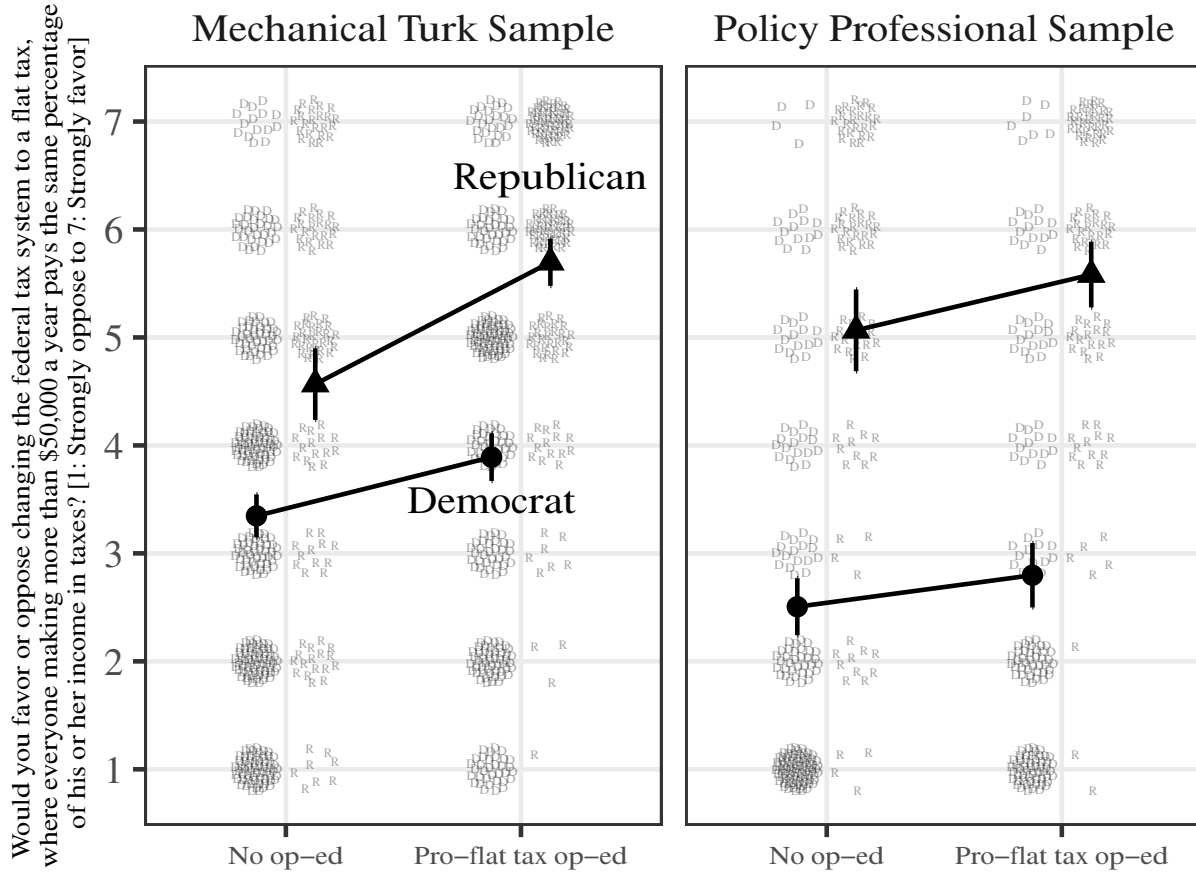
Post-treatment, we asked both the treatment group and a control group that did not read any op-ed “Would you favor or oppose changing the federal tax system to a flat tax, where everyone making more than \$50,000 a year pays the same percentage of his or her income in taxes?” While the op-ed does identify the author as a “Republican from Kentucky,” and so includes a partisan cue, I argue that the bulk of the treatment operates via persuasive information, largely because the goal Senator Paul’s op-ed was to contrast his tax policy with those of his opponents in the Republican presidential primary.

We ran this experiment twice. We conducted the first version once with a convenience sample obtained via Amazon’s Mechanical Turk (MTurk) service. The people on Mechanical Turk aren’t representative of all Americans, but they are nevertheless Americans.³ We also ran the experiment on a sample of policy professionals. These people are also not representative of all Americans – they are DC staffers, journalists, lawyers, and other professionals with some degree of connection to policymaking. For the moment, please don’t let the sample’s lack of representativeness bother you. We’ll return to the questions of generalizability and external validity in

³“Convenience sample” is a term of art that means the sample is made up of easy-to-interview people rather than randomly selected people from a well-defined population. MTurk is an online labor market where people get paid to perform small tasks, like tagging photos, transcribing videos – or answering academic surveys. MTurk makes it easy to interview large, diverse, and indeed, convenient samples.

Chapter 4.

Figure 1.2: Flat tax experimental results



Note: Survey experimental data from 860 MTurk respondents and 518 policy professionals who provided immediate and 10-day follow-up responses (Coppock, Ekins and Kirby, 2018). Democratic responses are plotted with the letter D and Republican responses with the letter R. The slopes of the lines connecting the average outcomes by condition represent average causal effects. On MTurk, the average effects are 1.13 (robust standard error: 0.20) for Republicans and 0.54 points (0.15) for Democrats. Among policy professionals, these values were 0.52 (0.25) for Republicans and 0.29 (0.20) for Democrats.

Figure 1.2 shows the results of both versions of the experiment. The MTurk experiment is on the left and the policy professionals experiment is on the right. I've overlaid the group averages by sample (MTurk or Elite), partisanship (Republican or Democrat), and treatment condition (treatment or control) on top of the raw data. Democrats and Republicans clearly differ with respect to the flat tax. On MTurk, partisans in the control group differ on average by over a full point on the 1 to 7 scale. Among the policy professionals, the gap is closer to 2.5 points.

Despite these baseline differences, both Republicans and Democrats on MTurk change their minds in response to the op-ed by similar amounts: 1.13 (robust standard error: 0.20) for Republicans and 0.54 points (0.15) for Democrats. A similar pattern holds for the policy professionals: 0.52 (0.25) for Republicans and 0.29 (0.20) for Democrats. These data offer clear evidence of persuasion in parallel. Despite some mild differences in the magnitude of the effects, they are qualitatively similar and are plainly all in the same direction. If there were backlash along partisan lines, the slopes for Democrats and Republicans would be oppositely signed. Instead of parallel motion, we would have contrary motion. That’s not what we find here – nor is it what we find in any of the persuasive information experiments to come.


We also recontacted our experimental subjects 10 days after they did or did not read Senator Paul’s opinion piece and we again asked subjects for their opinions about the flat tax. Figure 1.3 shows the results. Among the elite sample, the effects are less than half as large after 10 days as they were immediately post-treatment. Among the MTurk subjects, effects persist at 80% of the original magnitude for Republicans and 58% for Democrats. While it’s clear that the effects of persuasive information do dissipate over time, we nevertheless observe persistent treatment versus control differences even after 10 days.

The rest of the book contains many figures that look just like Figures 1.2 and 1.3, with various elements swapped in and out. I will report experiments that I’ve conducted myself with collaborators, experiments by others that I have replicated on new samples, and experiments by others that I have reanalyzed using my preferred set of tools. While the specifics of the randomly assigned persuasive information, the survey items used to measure policy opinions, and the subgroup divisions will vary, the pictures tell very similar stories: small effects in the direction of information.

3. *Parallel publics*

The idea that public opinions move together in parallel has a long history in American politics scholarship. In *The Rational Public*, Page and Shapiro (1992) used surveys that ask the same questions over many years to repeated cross-sections of the U.S. population to argue for the existence of “parallel publics.” They found that for most issues, opinion changes trend in the same direction for many segments of society. According to Page and Shapiro, opinion doesn’t

Figure 1.3: Flat tax experimental results: 10-day follow-up



Note: Survey experimental data from 860 MTurk respondents and 518 policy professionals who provided immediate and 10-day follow-up responses (Coppock, Ekins and Kirby, 2018). On MTurk, the average effects after 10 days were are 0.908 (robust standard error: 0.22) for Republicans and 0.32 points (0.15) for Democrats. Among policy professionals, these values were 0.18 (0.25) for Republicans and 0.13 (0.20) for Democrats.

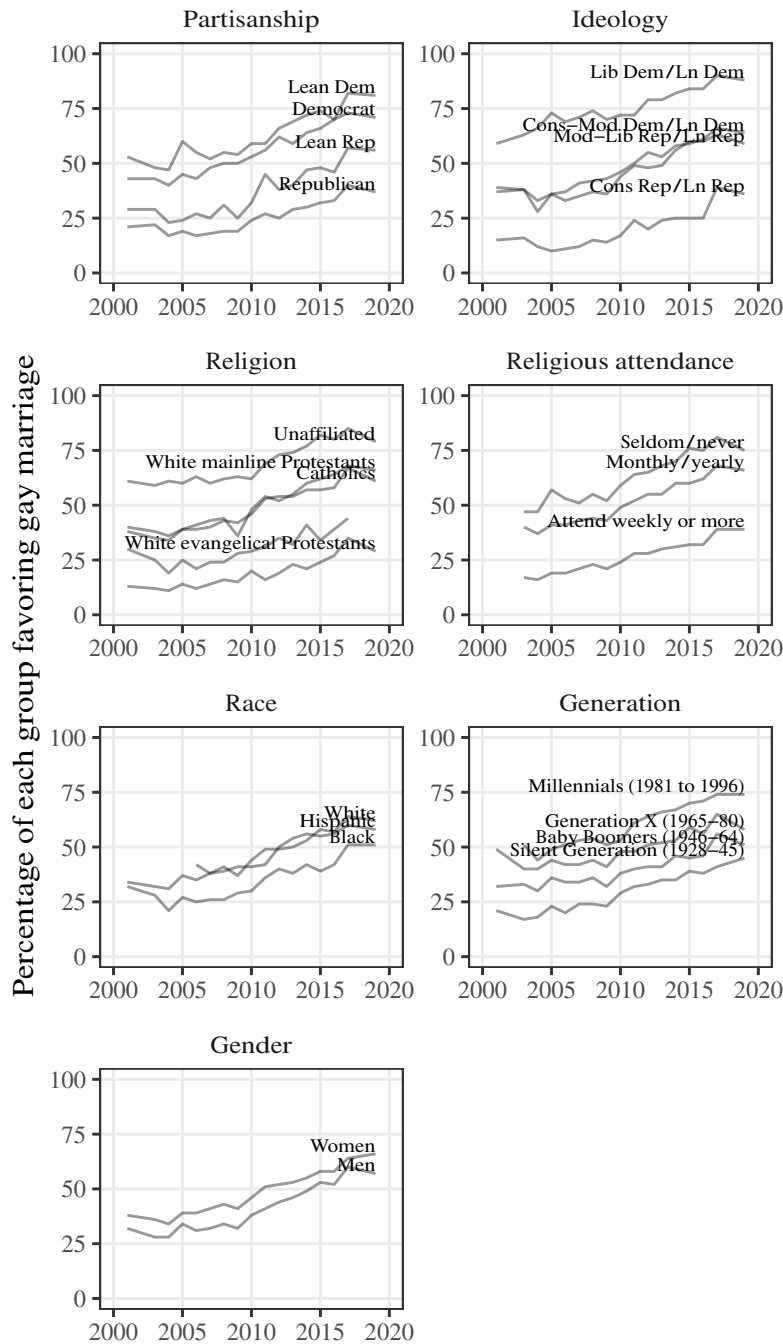
tend to polarize in the sense that as Democrats become more supportive of an issue, Republicans become less supportive of it. On the contrary, if Democrats warm to an issue, so too do Republicans. Since *The Rational Public*, political scientists have amassed evidence in favor of the parallel publics thesis in a huge number of domains: defense spending, redistribution, presidential approval, crime, even healthcare. To name a few of the dozens of articles and books that echo this finding, see Huxster, Carmichael and Brulle (2015) on climate change, Eichenberg and Stoll (2012) on defense spending, Enns (2007) on welfare spending, Kellstedt (2003) on bus-ing, or Porter (2020) on government waste. Green, Palmquist and Schickler (2002; p.135) and Hochschild and Einstein (2015; p.54) both document clearly parallel trends on attitudes about the Lewinsky scandal by partisan group. In April of 1998, 32% Democrats and 70% of Republicans believed Clinton had the affair; those figures increased by 14 points among both groups to 46% among Democrats and 84% among Republicans by the end of July.

To see some evidence of persuasion in parallel from the sort of nonexperimental data used in *The Rational Public*, consider Figure 1.4, which shows how attitudes toward same-sex marriage have evolved over time. The data come from repeated cross-sectional polls of Americans conducted by Pew Research Center between 2001 and 2019 (Fact Sheet: Changing Attitudes on Gay Marriage, 2019). Pew researchers estimated support for same-sex marriage in 25 separate demographic subgroups based on partisanship, ideology, religion, religious attendance, race, generation, and gender. In all 25, the proportion favoring gay marriage was higher in 2019 than in it was 2001, the beginning of data collection. Fitting straight lines to each series, we can estimate the average amount each group changed its position over time. Overall, the average change (or slope with respect to time) is about 1.7 percentage points per year. The slopes for some groups are slightly larger (Democrats: 2.1 points per year, White mainline Protestants: 2.1 points per year) than for others (Republicans: 1.2 points per year, White evangelical Protestants: 1.2 points per year) but the overall pattern across subgroups is very similar from one to the next.

As we'll delve into in Chapter 4, comparing magnitudes of change is not straightforward because the comparisons are sensitive to scaling. Whereas Democrats saw a greater percentage *point* change (28 points) than Republicans (16 points), Republicans experienced a larger *percent* change (76% increase) than Democrats (65% increase). It is not at all obvious which increase is "bigger." Suffice it to say that the magnitudes of change are similar but not the same and they

are difficult to rank.

Figure 1.4: American attitudes toward gay marriage 2001-2019



Note: Data from the Pew Research Center (2019). Despite clear differences in intercepts (position on the vertical axis), the slopes (change with respect to the horizontal axis) are all similar, with an average slope of a 1.7 percentage point increase in support for gay marriage per year.

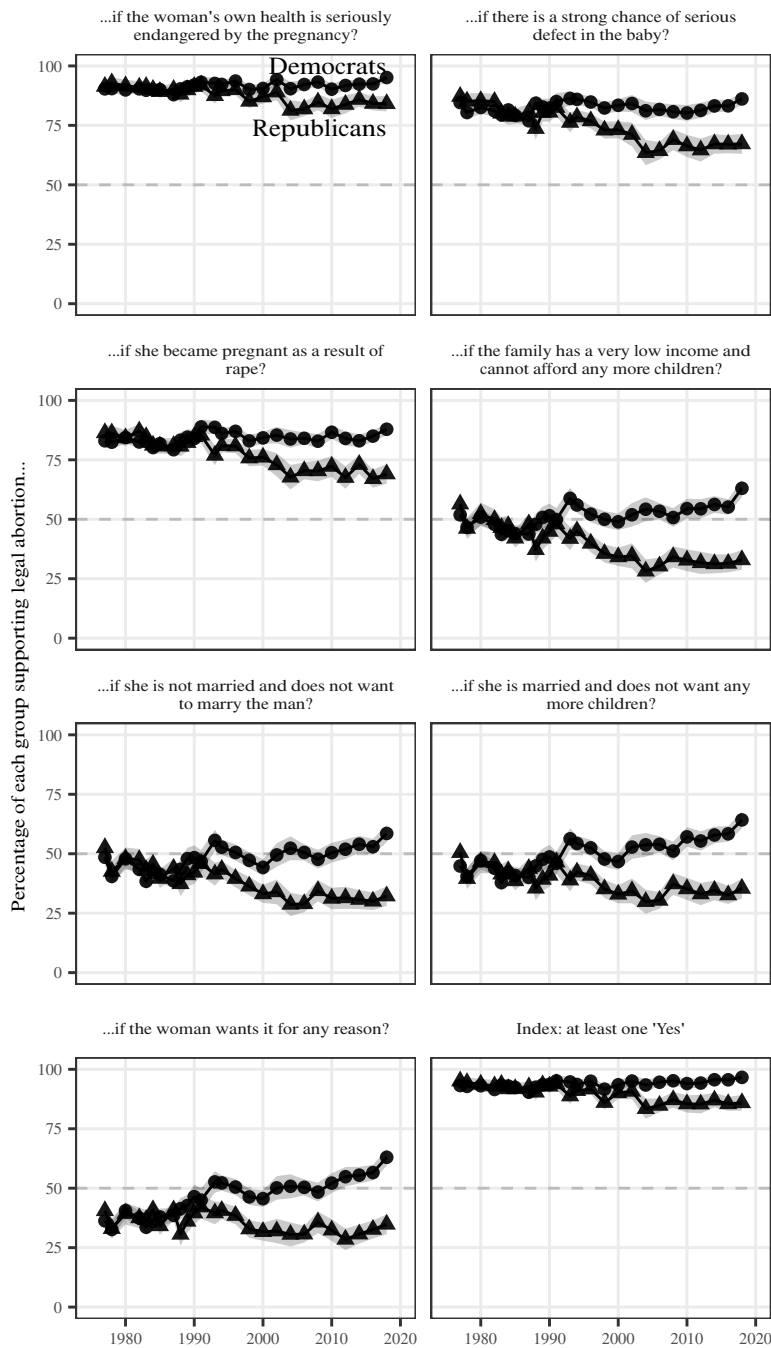
The parallel publics pattern is quite widespread but there are, however, some clear-cut ex-

ceptions. For example, Figure 1.5 shows how the last 30 years have seen a dramatic partisan *divergence* in abortion attitudes. Given the sharp contemporary divisions by party, it may be surprising to learn that Republicans and Democrats in the '70s and '80s held almost identical average opinions about the circumstances under which abortion should be allowed. The parties have moved in opposite directions on this issue ever since. Here we have obvious evidence of contrary, rather than parallel motion. In passing, I will note that the parties appear to mostly agree on the ranking of the “reasons,” and very large majorities of both parties support legal abortion in at least one case.

The evidence from repeated cross-sectional polls like those shown in Figures 1.4 and 1.5 can only take us so far. The descriptive patterns of how the average opinions of various subgroups of society do or do not move together are interesting, but important methodological issues arise when we want to use these data to draw causal inferences. These issues fall into three main categories.

First, we *want* to think of overtime change in opinion as the result of a causal process – but what treatment are these changes in response to? The variable on the horizontal axes of all these graphs is time, not some particular set of persuasive messages. It is difficult to conceptualize what the treatment is, since the media transmit a complex mix of persuasive messages on the one hand and group cues on the other. One “easy” explanation for the discrepancy between Figures 1.4 and 1.5 is that media environment may have included more positive persuasive messages for gay marriage over time but more polarizing group cues on abortion over time. Second, even if we could reasonably claim that the treatment is something like the balance of persuasive messages transmitted via mass media, we can't be sure that different segments of the population are exposed to the same set of messages. Differential exposure to pro- and counterattitudinal messages, even if mild, could seriously confound our inferences. Third, these repeated cross-sectional polls aren't panel studies. In a panel study, the same people are reinterviewed at multiple points in time but in a repeated cross-sectional design, the people who respond to the survey are different each year. Crucially, the cross-sectional design means that the composition of the groups could change over time. The kinds of people who call themselves Republicans in 1980 may be different from the kinds who call themselves Republicans today. Stated differently, the partisan divergence in abortion attitudes might not be the result of Republicans and Democrats being

Figure 1.5: American attitudes toward legal abortion by circumstance 1977-2018



Note: Data from the General Social Survey (1977-2018). The question asked, “Please tell me whether or not you think it should be possible for a pregnant woman to obtain a legal abortion...” in each of these seven circumstances. Despite clear partisan agreement on abortion through 1990, the gap between Republicans and Democrats has grown ever since, with more Democrats but fewer Republicans supporting legal abortion in each circumstance.

persuaded in opposite directions, but rather the result of pro- and anti-choice people sorting themselves into the parties differently over time (Killian and Wilcox, 2008; Levendusky, 2009).

Because of the inferential difficulties associated with repeated cross-sectional data, the main source of evidence in this book will come from randomized experiments. These studies have their own weaknesses and infirmities too, and I'll try to be as clear and forthcoming about those as possible throughout. But the main reason to turn to experiments to study persuasion is that we can be in control of the main causal agent we want to study. We're studying persuasive information, so that's what we'll randomize.

4. *What's at stake*

By the end of the book, I hope to have convinced you that the persuasion in parallel hypothesis is at least approximately correct. This single finding – approximately parallel changes in attitudes in response to persuasive information – has major implications for how we think about people who hold different views from us and their capacity for change.

First and foremost: *the other side is persuadable*. Even though polarization may have taken hold in many quarters, it is nevertheless worth our time to make arguments in favor of our preferred policies because we end up changing minds, even if just a little. That said, having conversations about politics is often painful. It's painful at family gatherings, it's painful on social media, it's painful among friends and coworkers. We pay a social cost when we disagree with others, but that doesn't mean the attempt has no effect at all on others' attitudes.

Second, political misinformation is dangerous, because people are persuaded by false information and true information alike. We have to hold those who control media platforms of every stripe – print, broadcast, or social – accountable for the spread of lies, conspiracy theories, and propaganda. Corrections to misinformation are effective since they too are a kind of persuasive information, but we would obviously be far better off if misinformation were not spread in the first place.⁴

⁴Thorson (2016) uses an elegant three-group candidate evaluation experiment to demonstrate this point. In her experiments, subjects could be assigned to a control condition, a negative misinformation condition, or a misinformation plus correction condition. Candidate evaluations in the misinformation group were lower than in control; they were higher in the misinformation plus correction condition than in the misinformation only condition. But in a pattern Thorson calls “belief echoes,” evaluations in the misinformation plus correction condition were still lower than in control – the corrections were unable to fully undo the damage done by the false information.

Third, we must recognize that we are ourselves persuadable as well. Being open to arguments from our opponents doesn't make us hypocrites, it just means we are like everyone else: a little bit persuadable.

5. *Where we're headed*

This book is aimed at Chapter 5, which will lay out the evidence from a large number of survey experiments that persuasion occurs in parallel. To get there, we're first going to correct the record on perhaps the most widely-cited and influential study that claimed the opposite. Lord, Ross, and Lepper (1979) purports to demonstrate that information causes "attitude polarization," which is equivalent to backlash as we've defined it. Chapter 2 will show that that claim is not correct, and unpacking the study's research design will explain why. In Chapter 3, I will provide definitions and distinctions and in Chapter 4, I will explain how my research design (the panel survey experiment) allows us to evaluate the persuasion in parallel hypothesis. Chapter 5 presents the evidence from those experiments. Chapter 6 demonstrates the over-time durability of these persuasive effects. The first six chapters of this book will be light on theory, but Chapter 7 will show what the evidence from these persuasion experiments does and does not mean for two of the most important theories of information processing, Bayesian learning and motivated reasoning. Chapter 8 will conclude by offering an explanation for why it *feels* like backlash is common though it is not and why we should nevertheless persist in trying to persuade the other side.