



University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations


2013

Belief Echoes: The Persistent Effects of Corrected Misinformation

Emily Thorson

University of Pennsylvania, ethorson@asc.upenn.edu

Follow this and additional works at: <https://repository.upenn.edu/edissertations>

 Part of the [Communication Commons](#), and the [Political Science Commons](#)

Recommended Citation

Thorson, Emily, "Belief Echoes: The Persistent Effects of Corrected Misinformation" (2013). *Publicly Accessible Penn Dissertations*. 810.

<https://repository.upenn.edu/edissertations/810>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/edissertations/810>
For more information, please contact repository@pobox.upenn.edu.

Belief Echoes: The Persistent Effects of Corrected Misinformation

Abstract

The omnipresence of political misinformation in the today's media environment raises serious concerns about citizens' ability make fully informed decisions. In response to these concerns, the last few years have seen a renewed commitment to journalistic and institutional fact-checking. The assumption of these efforts is that successfully correcting misinformation will prevent it from affecting citizens' attitudes. However, through a series of experiments, I find that exposure to a piece of negative political information persists in shaping attitudes even after the information has been successfully discredited. A correction--even when it is fully believed--does not eliminate the effects of misinformation on attitudes. These lingering attitudinal effects, which I call "belief echoes," are created even when the misinformation is corrected immediately, arguably the gold standard of journalistic fact-checking.

Belief echoes can be affective or cognitive. Affective belief echoes are created through a largely unconscious process in which a piece of negative information has a stronger impact on evaluations than does its correction. Cognitive belief echoes, on the other hand, are created through a conscious cognitive process during which a person recognizes that a particular negative claim about a candidate is false, but reasons that its presence increases the likelihood of other negative information being true. Experimental results suggest that while affective belief echoes are created across party lines, cognitive belief echoes are more likely when a piece of misinformation reinforces a person's pre-existing political views.

The existence of belief echoes provide an enormous incentive for politicians to strategically spread false information with the goal of shaping public opinion on key issues. However, results from two more experiments show that politicians also suffer consequences for making false claims, an encouraging finding that has the potential to constrain the behavior of politicians presented with the opportunity to strategically create belief echoes. While the existence of belief echoes may also provide a disincentive for the media to engage in serious fact-checking, evidence also suggests that such efforts can also have positive consequences by increasing citizens' trust in media.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Communication

First Advisor

Diana Mutz

Keywords

belief echoes, misinformation

Subject Categories

Communication | Political Science

BELIEF ECHOES:
THE PERSISTENT EFFECTS OF CORRECTED MISINFORMATION

Emily A. Thorson

A DISSERTATION

In Communication and Political Science

Presented to the Faculties of the University of Pennsylvania

In Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2013

Supervisor of Dissertation

Diana Mutz, Ph.D.

Samuel A. Stouffer Professor of Political Science and Communication, Director,
Institute for the Study of Citizens and Politics

Graduate Group Chairperson

Joseph Turow, Ph.D.

Professor of Communication

Nancy Hirschmann, Ph.D.

Professor of Political Science

Dissertation Committee:

Joseph N. Cappella, Ph.D., Gerald R. Miller Professor of Communication

Michael Delli Carpini, Ph.D., Professor of Communication and Walter. H. Annenberg
Dean

Marc Meredith, Ph.D., Assistant Professor of Political Science

BELIEF ECHOES:
THE PERSISTENT EFFECTS OF CORRECTED MISINFORMATION

COPYRIGHT

Emily A. Thorson

2013

ACKNOWLEDGEMENTS

This dissertation would not have been possible without a number of people. First of all, my parents, Kristi Andersen and Stuart Thorson. My mother embodies what I hope to be as an academic: a creative thinker, careful researcher, and incredible mentor. Her advice has been invaluable. My father encouraged me to pursue this joint degree and has always inspired me to work harder and smarter. My sister Kjerstin Thorson has offered practical advice and has shown me that there is indeed a life after the dissertation. My colleagues at Annenberg and the political science department have kept me excited to go to work each day, especially Sarah Vaala, Tara Liss, and Chelsea Schafer. And finally, I cannot thank Stephan Stohler enough for everything he has done for this dissertation, from proofreading to substantive criticism to much-needed emotional support.

I had the best dissertation committee possible. From when I first started at Annenberg, Joe Cappella and Michael Delli Carpini have offered me substantive criticism on my projects. And pursuing the dual degree was worth it if only because it provided the opportunity to work with Marc Meredith, who despite the busy schedule of an assistant professor always took time to help me think through my experiments and offer career advice.

Finally, this dissertation would not exist without the guidance of Diana Mutz. I can honestly say that she has taught me how to be a social scientist. Her commitment to asking normatively important questions and careful research design is extraordinary. I feel very lucky to have had the opportunity to work with her.

ABSTRACT

BELIEF ECHOES:
THE PERSISTENT EFFECTS OF CORRECTED MISINFORMATION

Emily A. Thorson

Diana Mutz

The omnipresence of political misinformation in the today's media environment raises serious concerns about citizens' ability make fully informed decisions. In response to these concerns, the last few years have seen a renewed commitment to journalistic and institutional fact-checking. The assumption of these efforts is that successfully correcting misinformation will prevent it from affecting citizens' attitudes. However, through a series of experiments, I find that exposure to a piece of negative political information persists in shaping attitudes even after the information has been successfully discredited. A correction—even when it is fully believed—does not eliminate the effects of misinformation on attitudes. These lingering attitudinal effects, which I call “belief echoes,” are created even when the misinformation is corrected immediately, arguably the gold standard of journalistic fact-checking.

Belief echoes can be affective or cognitive. Affective belief echoes are created through a largely unconscious process in which a piece of negative information has a stronger impact on evaluations than does its correction. Cognitive belief echoes, on the other hand, are created through a conscious cognitive process during which a person recognizes that a particular negative claim about a candidate is false, but

reasons that its presence increases the likelihood of other negative information being true. Experimental results suggest that while affective belief echoes are created across party lines, cognitive belief echoes are more likely when a piece of misinformation reinforces a person's pre-existing political views.

The existence of belief echoes provide an enormous incentive for politicians to strategically spread false information with the goal of shaping public opinion on key issues. However, results from two more experiments show that politicians also suffer consequences for making false claims, an encouraging finding that has the potential to constrain the behavior of politicians presented with the opportunity to strategically create belief echoes. While the existence of belief echoes may also provide a disincentive for the media to engage in serious fact-checking, evidence also suggests that such efforts can also have positive consequences by increasing citizens' trust in media.

Table of Contents

Table of Contents	vi
List of Figures	viii
List of Tables	ix
Chapter 1 – The rising concern over political misinformation.....	1
Misinformation and corrections in today’s media environment	2
Why misinformation is normatively troubling.....	3
Overview of dissertation	5
Chapter 2 – How misinformation emerges and is corrected	8
Defining misinformation	8
The process of misinformation and correction	10
Chapter 3 – Belief echoes and the role of delayed versus immediate	
corrections	20
Misinformation and the modern marketplace of ideas.....	20
Overview of experimental procedure	26
Experiment 1: The existence of belief echoes	34
Experiment 2: Belief time and motivated reasoning.....	42
Generalizability.....	51
Potential consequences of belief echoes	53

Chapter 4 - Affective and cognitive processes of belief echo creation	55
The affective process of belief echo creation.....	56
The cognitive process of belief echo creation	61
Using partisanship to distinguish between affective and cognitive processes.....	64
Experiment 3: Creating cognitive belief echoes through elaboration	68
Experiment 4: Distinguishing between affective and cognitive belief echoes.....	76
Implications of affective and cognitive belief echoes.....	84
Chapter 5 - Changing incentives for politicians and journalists.....	86
Three major approaches to fact-checking	87
Experiment 5: Constraints on politicians and media.....	90
Preventing belief echoes: incentives and disincentives.....	105
Chapter 6 - Conclusion.....	108
Limitations.....	109
Belief echoes in the real world.....	112
Practical lessons for the media	119
Belief echoes: part of a larger misinformation and correction process.....	123
References	125
Appendix	140
Appendix A. Article Manipulations	140
Appendix B. Question Wording	146

Appendix C. Distractor Task	152
Appendix D. Demographics	155
Appendix E. Other experiments conducted as part of research for this dissertation	156

List of Figures

Figure 3.1. Experiment 1: Effectiveness of the correction	39
Figure 3.2. Experiment 1: Effect of misinformation on attitudes.....	40
Figure 3.3. Experiment 1: Effect of misinformation on electability perceptions	41
Figure 3.4. Experiment 2: Design	46
Figure 3.5. Experiment 2: Effectiveness of correction	49
Figure 3.6. Experiment 2: Effect of misinformation on attitudes.....	50
Figure 4.1. Experiment 3: Design	69
Figure 4.2. Experiment 3: Example Article Format.....	70
Figure 4.3. Experiment 3: Evaluations of candidate by same-party condition	74
Figure 4.4. Experiment 3: Evaluations of candidate opposing-party condition	75
Figure 4.5. Experiment 4: Belief echoes in post-attitude recall task group.....	80
Figure 4.6. Experiment 4: Belief echoes in pre-attitude recall task group.....	81
Figure 5.1 Experimental Design.....	96
Figure 5.2 Experiment 5: Evaluations of accuser.....	100
Figure 5.3 Experiment 5: Perceived bias of <i>Iowa Ledger</i>	102
Figure 5.4 Experiment 5: Media Evaluations	105

List of Tables

Table 2.1 The prerequisites of belief echo creation.....	10
Table 3.1. Experimental Procedure.....	29
Table 4.1. Experiment 3: Misinformation condition headlines.....	73
Table 4.2. Experiment 4: Open-ended response distribution.....	82

Chapter 1 – The rising concern over political misinformation

In 2011, *New York Times* columnist Paul Krugman predicted that the 2012 Presidential campaign would mark the nation's entrance into what he called "post-truth politics" (Krugman 2011). In a world of post-truth politics, Krugman argued, politicians would strategically manipulate the truth for political advantage and the media would fail to penalize them in any meaningful way. Krugman published his piece in the midst of the 2011 Republican primaries, and over the next year, the unfolding Presidential campaigns seemed to at least partially confirm his forecast.

For example, the October 11 debate between Vice-President Joseph Biden and Republican vice-presidential candidate Paul Ryan was littered with dubious factual claims, exaggerations, and misleading statements (Gerstein & Samuelsohn 2012). However, media coverage of the debate disproved the second half of Krugman's prediction. Rather than ignoring Biden and Ryan's falsehoods, the media enthusiastically engaged in fact-checking the candidates' statements. The *Washington Post's* headline read "Fact-Checking the Vice-Presidential Debate," while the more dramatic headline from the *Daily Beast* blog asked "Who Lied?" These headlines were not an anomaly: they exemplify an increasing media commitment to correcting false claims. Whether or not the absolute level of misinformation in politics has changed, it is certainly the case that conflict over what constitutes the truth have emerged as a major narrative in American politics.

Misinformation and corrections in today's media environment

As the headlines that followed the Ryan-Biden debate demonstrate, corrections of false claims can emerge quickly in today's media environment, and they are often the focus of political news coverage. This means that while misinformation may be common, misleading factual claims are frequently contested, as opposing political forces and the media themselves expend resources to debunk false claims.

A 2010 incident aptly illustrates this rapid process of misinformation and correction. On November 2, the *Press Trust of India* published an online piece about President Obama's upcoming trip to India. The article claimed that three thousand people would accompany the President, costing the U.S. government \$200 million per day. Within hours several right-wing blogs republished the claim, and within a day Fox News broadcast the story. But just as quickly, the misinformation was corrected. Just 48 hours after the initial claim was made, FackCheck.org declared it incorrect, and soon after, a *Wall Street Journal* editorial called the claim "demonstrably false" (Weisman 2010). The entire process of misinformation and correction took fewer than three days.

Patterns of misinformation and correction have occurred around topics ranging from candidates' biographies to federal funding allocations. In many of these cases, the correction has become a bigger story than the misinformation itself. For example, the first and only article published by the *L.A. Times* about the cost of Obama's India trip was entitled "Purported cost of Asia trip 'wildly inflated'" (Times 2010). The concern with correcting false claims has led to the rapid growth fact-

checking organizations as well as to a recent re-commitment by journalists to “shucking the old he-said-she-said formulation and directly declaring that certain claims are false” (Graves & Glaisyer 2012, Bennet 2012).

Of course, despite the renewed focus on fact-checking, misinformation is not always successfully corrected. The canonical example of uncorrected misinformation may be the existence of weapons of mass destruction in Iraq, but this is not the only false claim to go uncorrected for weeks or even months (Graves 2012). And even if misinformation is corrected, citizens may not always accept the correction. For example, some Americans continue to believe that President Obama is a Muslim, despite the overwhelming evidence to the contrary (Nyhan 2012). Finally, corrections may not reach as wide of an audience as the initial misinformation. In Chapter 2, I discuss each of these possibilities in more detail and summarize scholarly work on the circumstances under which corrections might fail to successfully debunk false claims.

Why misinformation is normatively troubling

Efforts to correct misinformation are driven by concerns over the potential consequences of a misinformed citizenry. Democratic theory holds that citizens’ attitudes should be based on an accurate understanding of the political world. While many citizens lack knowledge of some fundamental aspects of politics (Delli Carpini & Keeter 1996), being uninformed is a very different problem from being misinformed. A lack of information can be overcome through the use of heuristics or other shortcuts (Popkin 1991, Sniderman et al 1991). In contrast, misinformation

poses a much more serious threat (Kuklinski et al 2000). Insofar as attitudes are based on factual knowledge, citizens who possess inaccurate information may form opinions that differ substantially from the opinions they would have formed were they correctly informed. In addition, their belief that they are already informed may prevent them from seeking new information.

Journalistic and scholarly efforts to correct misinformation (described in more detail in Chapter 2) are aimed at preventing these problems, thereby moving us closer to the idealized marketplace of ideas. The marketplace of ideas holds that a community benefits from maximizing the free exchange of ideas. False claims, the theory holds, will be discredited when they are subjected to scrutiny and independent verification (Gurevitch & Blumler 1990). Social media, the proliferation of independent and decentralized blogs, and the 24-hour news cycle all provide citizens with increased access to a greater quantity of information. While some of this information may indeed be misleading, it is also the case that factual claims encounter widespread and decentralized scrutiny.

Yet the marketplace of ideas makes a critical assumption. This assumption, which I discuss in more detail in Chapter 3, holds that after individuals discard information that is shown by the market to be false, it will cease to affect attitudes. In other words, reading a correction should cause attitudes initially affected by false claims to revert back to their pre-exposure state. This assumption must be true if we are to embrace the characteristics of today's media environment, and especially its attention to fact-checking, with confidence.

In this dissertation, I present evidence showing that this assumption is unfounded. I find that exposure to information about a candidate creates *belief echoes*: effects on attitudes that persist even after the claim is rejected. A correction—even when it is fully accepted—only reduces roughly half of the attitudinal effects of exposure to negative information. An *L.A. Times* reader learning about the “wildly inflated” cost of Obama’s trip to India might come away certain that the misinformation was false. The existence of belief echoes suggests that in spite of her dismissal, her attitude towards Obama would be more negative than if she had never read the article. Belief echoes suggest that even when the marketplace of ideas operates efficiently to correct false claims, misinformation can still shape citizens’ attitudes.

Overview of dissertation

In Chapter 2 of this dissertation, I describe how false claims emerge and are debunked in the current media environment, and in the process review the literature on misinformation and corrections. Specifically, I look at how misinformation emerges and is corrected, and discuss the channels through which corrected misinformation might affect citizens’ attitudes.

Chapter 3 discusses belief echoes in more detail, drawing on research in psychology to explain why exposure to misinformation may have lingering effects on attitudes. First, I present the basic experimental paradigm used to measure belief echoes. The first experiment in the chapter demonstrates exposure to false claims creates belief echoes, while the second shows that belief echoes can be created even

when the misinformation is corrected immediately, arguably the gold standard of journalistic fact-checking. In addition, belief echoes are not necessarily driven by party loyalty: attitudinal effects can occur regardless of the partisan slant of the misinformation. Repeating a false claim—even if only to correct it—can have serious effects on attitudes.

In Chapter 4, I explore two possible mechanisms of influence for belief echoes: an affective path and a cognitive path. I theorize that affective belief echoes are created through a largely unconscious process in which a piece of negative information has a stronger impact on evaluations than does its correction. Cognitive belief echoes, on the other hand, are created through a conscious cognitive process during which a person recognizes that a particular negative claim about a candidate is false, but reasons that its presence increases the likelihood of other negative information being true. Two experiments examine the factors that lead to the creation of cognitive and affective belief echoes. The results suggest that affective belief echoes are created across party lines. Cognitive belief echoes, however, are more likely when a piece of misinformation reinforces a person's pre-existing political views.

The existence of belief echoes raises serious normative concerns about political news coverage, and in Chapter 5 I present the results of two experiments directly addressing these concerns. Belief echoes provide an enormous incentive for politicians to strategically spread false information with the goal of shaping public opinion on key issues. However, results show that politicians also suffer consequences for making false claims, an encouraging finding that has the potential

to constrain the behavior of politicians presented with the opportunity to strategically create belief echoes. The existence of belief echoes may also provide a disincentive for the media to engage in serious fact-checking. However, evidence suggests that even when exposure to corrected misinformation does create belief echoes, such efforts can also have positive consequences by increasing citizens' trust in media.

Finally, Chapter 6 concludes by first outlining some of the limitations of this dissertation. I then discuss the larger implications of belief echoes, discussing several real-world factors that might amplify or minimize belief echoes. Finally, I offer several practical lessons for potentially lessening the negative impact of misinformation.

Chapter 2 – How misinformation emerges and is corrected

This chapter first defines what this dissertation means by “misinformation” and outlines the conditions that create belief echoes. Belief echoes are created when misinformation is *successfully* corrected by then continues to affect attitudes. A reasonable question to ask is “how often does this actually occur?” In this section, I argue that the incentive structure of today’s political media can encourage the successful correction of misinformation. Specifically, I discuss (1) the ways in which misinformation can emerge, (2) how it is corrected, and (3) the conditions under which those corrections are successful. I review existing literature on misinformation as well as present concrete examples of how misinformation emerges and is corrected. In the process, I also introduce several aspects of misinformation and corrections that have the potential to amplify or minimize belief echoes.

Defining misinformation

Research which explores misinformation must first tackle a troublesome epistemological question: what constitutes an objective political fact?¹ This question is central not only to studies of misinformation but to journalism more generally (Schudson 2001). The common-sense definition of a fact is a piece of information that – to a reasonable degree – corresponds to a phenomenon beyond our subjective

¹ Philosophies of science are, of course, subject to radical critiques which subject basic concepts – like facts – to interrogation. These critiques are important to acknowledge, but they threaten to undermine the positive enterprise altogether. As such, I outline the limitations of this research but

experience. Or, put more colloquially, information that accurately represents the real world. In the world of politics, this seemingly simple definition is complicated by several factors.

First, much of the information considered important to political decision-making facts is not entirely subject to independent verification or falsification: for example, a claim that the Affordable Care Act will reduce healthcare costs by fifteen percent in 2020, or a claim that Governor Romney's "true beliefs" about abortion are different from his public statements. In addition, unlike facts about easily observable phenomena (like the color of an elephant or the number of eggs in the carton), many key pieces of political information are mediated.² Economic data is collected through the Congressional Budget Office, hour-long speeches are summarized in brief articles, policies are reduced to talking points. The decisions made by institutions, journalists, and politicians about how to communicate political information to the public are not made at random but are themselves shaped by the political environment: as Kuklinski et al (1998) state, "the criteria for and relevance of political facts are determined within, not outside, politics" (p 147).

These difficulties should not be ignored, but rather built into our understanding of how misinformation can affect attitudes. Political facts are often both contested and contestable. However, an acknowledgement of relativity should not be taken as an abandonment of the goal of objectivity. Any given piece of information resides somewhere on the spectrum from falsifiable to unfalsifiable. In this dissertation, both in the examples I draw on as well as in my experimental

² Without making a causal argument, I will note that one of the most important political facts for electoral outcomes—economic performance—can be directly experienced by voters.

manipulations, I consciously focus on factual assertions that fall on the more objective side of this spectrum, but that are still relevant to political decision-making. For example, we can determine with relative (although not complete) certainty whether a candidate accepted campaign donations from a criminal; whether Barack Obama was born in the United States; or whether infant mortality rates rose in a particular state. While these claims might not be as clear-cut as an assertion about what type of cheese John Kerry ordered on his cheesesteak, they are more verifiable than a claim about how John Kerry's economic plan will benefit the middle class.

The process of misinformation and correction

Belief echoes are created when misinformation emerges and then is successfully corrected. To understand when those circumstances arise, it is useful to outline the specific processes leading to the spread of misinformation and its correction. Table 2.1 The prerequisites of provides an overview of this process.

Table 2.1 The prerequisites of belief echo creation

Process	Key Questions
1. Misinformation emerges	What are the sources of misinformation? How long is the misinformation in the environment before it is corrected?
2. Misinformation is corrected	Who corrects the misinformation? How widely is the correction disseminated? What is the format of the correction?
3. Correction is accepted	What factors lead a correction to be accepted? Who accepts the correction?

Political misinformation comes from many sources, including candidates, institutions, elected officials, party organizations, and journalists themselves. Once misinformation emerges, it may stay uncorrected for hours (false claims made in the 2012 vice-presidential debate) or years (weapons of mass destruction in the Iraq). The source of the correction also varies widely, from media outlets and fact-checking organizations to rival campaigns and even individual citizens. The audience of any given correction may be very different from the audience of the initial misinformation, and even when a correction *does* reach its intended audience, it may not be fully accepted.

Most existing work on misinformation focuses only on the third step: whether – and to what extent – individuals accept corrections of misinformation (Nyhan & Reifler 2010; Bullock 2007; Berinsky 2012). However, a narrow preoccupation with designing successful corrections ignores important additional consequences of misinformation on attitudes: for example, the creation of belief echoes. Understanding the circumstances that lead to belief echoes requires examining how individuals process information, as well as how journalist and media outlets make decisions about how and when to issue corrections. Belief echoes emerge at the intersection of these factors.

Emergence of misinformation

Sources of misinformation

Misinformation can emerge from many sources, ranging from deliberate smear campaigns to careless reporting. The experiments in this dissertation focus

on misinformation that is pushed by a campaign or candidate. This type of misinformation is not uncommon: between 2010 and 2011, 74% of the statements checked by Politifact were made by political candidates and campaign organizations (Ostermeier 2011). Of course, Politifact, like any fact-checking enterprise, chooses what statements to examine, and this choice comes with its own biases. Still, direct quotes from elected officials and candidates are a mainstay of political news (Mindich 1999), and as such are frequently the conduit through which misinformation—deliberate or not—reaches the public.

Length of time misinformation is in the environment

The length of time a person believes a misperception before it is corrected may affect the magnitude of the belief echoes it creates. Chapter 3 outlines the psychological underpinnings of this effect, and this section discusses real-world examples of delayed corrections in the real world.

While some misinformation is corrected immediately (like false claims made during debates), other misinformation is believed for some time (like the existence of weapons of mass destruction in Iraq). Misinformation may be especially likely to remain uncorrected when it spreads through social networks rather than through mainstream media (Weeks & Southwell 2010). For example, soon after Presidential candidate John McCain selected Alaska Governor Sarah Palin as his running mate, rumors emerged alleging that she had banned books from the town library as mayor of Wasilla. The claim spread through a chain email purporting to be a “list of books Palin tried to have banned,” supposedly taken from the official minutes of the Wasilla Library Board (Thornburgh 2008). The claim slowly moved from email onto

some liberal blogs, and was definitively corrected over a month after it emerged.

ABC News called the librarian at the center of the controversy, who confirmed that Palin never requested that any books be banned (Ross et al 2008). Given the wide variation in belief time, ranging from none (for example, when individuals first encounter misinformation in the context of it being corrected) to years, it is important to consider the implications of belief time for misinformation's effects on attitudes.

Correction of misinformation

Who corrects misinformation

As with any other type of political information, citizens encounter corrections largely through the media. In some cases, journalists themselves are responsible for investigating and then publicly correcting false claims. Increasingly, mainstream media institutions also depend on independent fact-checking organizations, relying not only on their resources but also in on their perceived objectivity (Graves 2012b). Fact-checking organizations have multiplied over the past decade at the national, state, and local levels (Spivak 2010). It is important to note that the new breed of fact-checking differs radically from the type of internal fact-checking that has traditionally been practiced by media outlets. Internal fact-checking practices verify that quotations and background information are reported accurately. In contrast, the more recent breed of fact-checking is concerned with adjudicating between claims that have already been reported in the news.

Regardless of whether claims are adjudicated by media organizations (for example,

the *Washington Post's* Fact Checker column) or by independent organizations (for example, FactCheck.org), they operate under similar constraints. Both types of organizations struggle to arbitrate between competing claims while avoiding labels of partisan bias (Graves 2012b).

While the 24-hour news cycle may have increased the average speed at which false claims are contested, the time that elapses between the emergence of misinformation and its correction varies considerably. This delay may result from a number of factors, including deliberate obfuscation on the part of the false claim's originator. For example, Kaufmann (2004) argues that the media's failure to discredit claims about weapons of mass destruction in Iraq resulted from purposeful manipulation by the Bush Administration. By carefully controlling access to intelligence information, they prevented the contestation that would have disproven their justifications for war. In addition, argues Kaufmann, the parties that should have been pursuing for the truth (the opposition party, independent experts, members of the media) failed to do so, largely because of the post-September-11th crisis atmosphere that made questioning the decisions of the administration less socially acceptable.

While the media initiates most corrections, individual citizens can also play a role. First, the changing media environment can blur the distinction between media outlets and consumers. The rise of platforms like blogs and Twitter provide an enterprising person—even one not employed by a media institution— with a platform from which she can call attention to false claims. For example, bloggers were largely responsible for a 2004 event in which Dan Rather was forced to retract

false claims that he had made about President Bush's National Guard service (Jarvis 2004).

The public may also play an indirect role in the correction of misinformation by encouraging the media to take a more active role in adjudicating between false claims. For example, on January 12, 2012, Arthur Brisbane, the public editor of the *New York Times*, wrote a column that generated an immediate and overwhelming response. Brisbane posed a question to *Times* readers: "Should the *Times* be a Truth Vigilante?" (Brisbane 2012). Brisbane's somewhat loaded term was meant to ask whether *Times* articles should include corrections of misleading statements made by political figures. Within hours, visitors flooded the site with responses. The vast majority fully supported the "truth vigilante" approach and strongly encouraged the *Times* to engage in more active fact-checking (Graves 2012a).

Where the correction spreads

Corrections³ are often transformed into news stories in their own right, in some cases reaching a larger audience than the misinformation. The narrower the time gap between the misinformation and the correction, the higher the odds that any given consumer of news will first encounter the misinformation in the context of it being corrected. For example, a person who skipped watching the 2012 vice-presidential debate might wake up to find a headline reading "Fact-Checking the VP

³ When one thinks of a journalistic correction, the format that immediately comes to mind might be a "Corrections" column buried in the middle of a newspaper. But this is not the model for the corrections with which this dissertation is concerned. Those "page-five" corrections tend to be concerned with spelling errors or other minor mistakes made by the media outlet itself. In contrast, the new breed of corrections concerns false claims that, while they may have been repeated by a media outlet, are not directly attributable to them.

Debate” in her morning paper. This trend of corrections as stories in and of themselves is exacerbated by the 24-hour news cycle and constant demand for new content. Fact-checking, with its neat storyline and standardized format, provides an easy way for media outlets to “fill the news hole,” either by borrowing content from fact-checking organizations or by producing their own fact-check stories and segments (Graves 2012b, p. 266). It is important to consider the normative assumptions inherent in this development. A news show that would never knowingly spread misinformation might feel quite comfortable producing a “fact-check” segment that introduced a new audience to a false claim made in a speech earlier that day. The assumption motivating this choice is that because the claim is *presented as false*, it can no longer affect attitudes.

Thirty years ago, most Americans received their political news from the same few sources (Prior 2007). Today, however, Americans not only have access to a much wider variety of news, but they also encounter a wider selection of channels through which they are exposed to that news (Pew 2012). Readers of a particular news article might have encountered it in very different ways, encountering a link on Twitter; hearing a mention on a podcast; or receiving an email from a friend. A person who reads an article containing a false claim may not later return to the same news source to read the correction. As such, when a publication prints a false claim and then later debunks this claim, the two pieces may have very different audiences. As Garance Franke-Ruta writes in the *Atlantic*, “every story lives an independent life on the social Web and there's no guarantee the reader of any given

report will ever see a bundled version of the news or the relevant fact-checking column, which could have been published months earlier” (Franke-Ruta 2012).

Acceptance of correction

The role of partisanship

Understanding the circumstances under which a correction leads people to update their factual beliefs has been the focus of most existing work on misinformation (Nyhan & Reifler 2010, Bullock 2007, Berinsky 2012). In this dissertation, I use the term “belief persistence” to describe individuals’ resistance to corrections, or unwillingness to discard false beliefs. Belief persistence is different from belief echoes, which are attitudinal effects that persist *after* a correction is fully accepted and a false belief is discarded. Developing strategies for reducing belief persistence is not a primary concern of my research. However, belief persistence is relevant to understanding belief echoes insofar as concerns over belief persistence inform how media outlets design corrections. These choices may in turn affect the magnitude of belief echoes.

Substantial evidence suggests that belief persistence is strongest when a correction runs counter to a person’s views. In a series of experiments conducted by Nyhan & Reifler (2010), participants were presented with news articles that included misinformation about three controversial and ideologically divisive political issues: stem cell research, tax cuts, and the war in Iraq. Some people were provided with corrections while others were not. When the correction reinforces their partisan attitudes (for example, when Democrats read that there were not

weapons of mass destruction in Iraq), it is accepted. When the correction runs counter to their partisan attitudes (for example, when Republicans read that the Bush tax cuts did not improve the economy), individuals retain their belief in the misinformation. Indeed, Nyhan & Reifler point to a “backfire” effect in which exposure to a correction can actually strengthen the initial belief.

Correction format

Media outlets make important choices about the form which corrections take, and these choices inevitably impact the success of those corrections. One of the major factors affecting whether a correction is accepted is how it is framed. Over the past few years, three separate approaches have been advocated most frequently: reliance on external institutions such as fact-checking institutions, a renewed commitment to fact-checking by journalists themselves, and a “he-said/she-said” approach that gives a voice to both sides of factual disputes rather than arbitrating between them. Chapter 6 of this dissertation discusses this debate in more detail and examines how it might affect the magnitude of belief echoes.

Conclusion

In a successful marketplace of ideas, the above process will proceed smoothly. Misinformation will be corrected, the correction will be widely disseminated, and the false belief will be discarded. The experiments in this dissertation create a world in which the marketplace works exactly as it should: misinformation is corrected immediately and the correction is fully accepted. The assumption put forth by the marketplace of ideas and shared by fact-checkers,

journalists, and scholars is that when a correction is successful, misinformation's attitudinal effects will be eliminated. The next chapter engages more deeply with these assumptions, examining whether even in in this ideal world, false claims can still shape citizens' attitudes.

Chapter 3 – Belief echoes and the role of delayed versus immediate corrections

As discussed in Chapter 2, the modern media environment facilitates the spread of both misinformation and corrections. The assumption of journalists, fact-checkers, and scholars attempting to design more effective corrections is that once misinformation has been successfully discredited, it will cease to affect attitudes and preferences. In contrast, this chapter examines the impact of misinformation in a world where corrections are entirely successful, even among partisans. The experiments presented this chapter question that assumption.

Misinformation and the modern marketplace of ideas

While false information has always plagued politics, several factors of the modern media environment have made it easier for misinformation to reach more people, more quickly. Media competition, the proliferation of internet news sources, and (counter-intuitively) the rise of fact-checking organizations may all contribute to the dissemination of false information.

Competition between media outlets—in print, on television, and also online—has intensified over the last decade. Political coverage is a key battleground, as news outlets race to produce enough content to “fill the news hole” created by 24-hour news channels and web streams (Ladd 2012, Box-Steffensmeier & Schie 2009). Practically, this race may lead news outlets to rely less on original reporting and more on leads from other publications (for example, blogs), some of which might be

less conscientious about fact-checking than traditional mainstream news media (Woodly 2008).

In their drive to publicly correct false claims, fact-checking organizations may also play an inadvertent role in the spread of misinformation. Some of these fact-checking organizations are independent institutions (FactCheck.org) while others are internal organizations sponsored by media outlets (the *Washington Post's* Fact-checker, the *Arizona Republic's* AZ FactCheck). Still others are straightforwardly partisan (Media Matters, NewsBusters), with the goal of correcting untruths spread by the opposition. Fact-checking has also become increasingly integrated into mainstream journalism. Between 2004 and 2010, mentions of the term “fact-check” in major newspapers doubled (Graves & Glaisyer 2012). For news outlets desperate for content, fact-checkers can be a valuable resource. As such, a byproduct of the emergent fact-checking industry is the publicizing of misinformation, both through its own platform and by providing ready-made content for news outlets. From the perspective of journalists and fact-checkers, they are spreading corrections, not misinformation. Of course, it is near-impossible to repeat a correction without also repeating the misinformation, but they operate under the assumption that the correction will nullify any potential effect of the misinformation on attitudes. The experiments in this chapter test that assumption.

The cycle of misinformation and correction exemplifies the marketplace of ideas, where “opposing views may meet, contend, and take each other’s measure” (Gurevitch & Blumler 1990). While the marketplace of ideas metaphor is often

employed to describe the competition between normative prescriptions about the world, it is also relevant as a way of understanding how political reality itself is contested. As defined in Oliver Wendell Holmes' dissent, the ultimate end of the marketplace of ideas is nothing less than the truth, as "the best test of truth is the power of the thought to get itself accepted in the competition of the market" (Holmes 1912).

Conflict over what constitutes the truth is a hallmark of modern political debate. Politicians argue over what a policy will or will not do, over what a rival did or did not say, or over who is or is not receiving benefits. In the political world, almost any supposed fact is up for debate. As Kuklinski et al (1998, p 48) point out "[v]ery few factual claims are beyond challenge; if a fact is worth thinking about in making a policy choice, it is probably worth disputing. Rival advocates compete to define the facts, control their presentation, and determine their relevance."

In a properly operating marketplace of ideas, this contestation over facts should ultimately lead to the truth. As false claims are discredited and exit the marketplace of ideas, the truth will prevail. The following section examines two ways in which the marketplace of ideas might fail, thereby allowing misinformation to shape attitudes and ultimately affect political outcomes. The first point of failure, belief persistence, occurs when corrections fail to eliminate belief in false claims. This problem has received substantial attention from scholars. The second failure, and the one addressed by this dissertation, is that misinformation may continue to affect attitudes even if the correction is successful.

Belief persistence

Much media and scholarly attention to misinformation in recent years has focused on belief persistence: the tendency for people to maintain their misperceptions even in the face of credible corrections (Nyhan & Reifler 2010, Berinsky 2011). Intuitively, belief persistence seems normatively problematic because believing a piece of political misinformation could cause a person to hold a different opinion she would if she were correctly informed.

However, the causal arrow between information and opinions can point both ways. Most theories of democracy suggest that information should inform preferences: a person's policy preferences should be shaped by the facts he has at his disposal (Delli Carpini & Keeter 1996). In reality, preferences may also shape beliefs: pre-existing attitudes inevitably affect what people believe to be true. Partisanship strongly colors individuals' estimates of supposedly objective facts (Bartels 2002, Lodge & Taber 2000). This process is referred to as motivated reasoning. Motivated reasoning occurs when people who hold strong attitudes about a given topic (for example, a candidate or a policy) reject counter-attitudinal information. This tendency means that when a piece of new information (for example, a correction) contradicts an individual's pre-existing beliefs, she may have difficulty accepting it (Kuklinski et al 2000).

Motivated reasoning has been found to drive the majority of the belief persistence found by past research into political misinformation. In a series of experiments conducted by Nyhan & Reifler (2010), participants were presented with news articles that included misinformation about three controversial and

ideologically divisive political issues. When the correction challenged their partisan attitudes, individuals retain their belief in the misinformation. Otherwise, the correction was accepted.

The tendency for partisans to engage in belief persistence may not be as problematic as it first appears. From a normative perspective, political misinformation is dangerous insofar as it leads individuals to make political evaluations and judgments that differ from those that they would have made if they were correctly informed. Partisans' reluctance to dismiss misinformation that reinforces their pre-existing attitudes is a different phenomenon. For such people, the misinformation is not shaping their political opinions: rather, those opinions are shaping what they choose to accept as facts.

Lingering effects of misinformation

The marketplace of ideas also makes a second implicit assumption that often goes unquestioned: that once a piece of information leaves the marketplace of ideas (for example, when it is successfully corrected) it ceases to affect attitudes and preferences. However, research in psychology suggests that discarding false information may not be so simple. People who form an attitude based on a particular piece of evidence tend to retain at least some of that attitude even when the evidence is proven false (Anderson, Lepper, & Ross 1980; Anderson, New, & Speer 1985; Johnson & Seifert 1994).

This pattern persists when people are given information as diverse as the causes of a warehouse fire, feedback on their own task performance, and the

interpretation of suicide notes (Johnson & Seifert 1994; Wegner, Coulton, and Wenzlaff 1985; Ross, Lepper, & Hubbard 1975). The phenomenon, termed “belief perseverance” by psychologists, has obvious implications for the study of political misinformation. If misinformation shapes attitudes even after it has been successfully discredited, then publicly correcting misinformation—even when the corrections are fully accepted—may not be sufficient to dampen its effects on public opinion.

However, there are several reasons to be skeptical of whether these psychology findings have parallels in the political world. Most tests of belief perseverance construct situations in which people draw on a piece of information to construct a causal theory about the world (for example, why some firefighters are successful or why a student did poorly on an exam). The experiments find that when the information is discredited, the theory persists. In politics, however, citizens use information not to construct causal theories but to form attitudes and evaluations of candidates and policies. (Anderson & Lindsay 1998). False information may not have the same lingering effect on evaluations as it does on the “knowledge structures with a causal or explanatory component” that are the focus of psychology experiments.

In addition, psychological theories of belief perseverance tend to focus on domains where people do not have strong prior attitudes. In politics, the opposite is the norm. Party identification and ideology serve as cues that allow citizens to make immediate judgments about novel political information (Lupia & McCubbins 1998).

It is possible that these cues will overwhelm any potential lingering effects of discredited misinformation.

Finally, experimental work in psychology imposes a delay between the misinformation and the correction, giving participants time to integrate the false information into their belief systems before it is shown to be false. In the fast-moving political world, this delay may not always occur. As discussed earlier, it is very possible that an individual may first be exposed to a piece of misinformation in the context of it being corrected.

Overview of experimental procedure

I use randomly assigned experiments to establish the existence of belief echoes, by comparing the attitudes of individuals exposed to successfully corrected misinformation to those who were not. establish the existence of belief echoes. An alternate approach to testing for the presence of belief echoes in the political world would be to employ an observational strategy, examining the attitudes of citizens who have been exposed to real-world misinformation and corrections. However, this approach poses a serious practical challenge. Belief echoes occur when a false claim is successfully corrected, but continues to affect attitudes nonetheless. Therefore, testing for belief echoes requires measuring both exposure and attitudes. In a survey context, determining whether a person has been exposed to a false claim (or its correction) requires directly asking them about it, thereby potentially affecting their attitudes. Given the challenges inherent in accurately measuring

exposure to misinformation, I test for the presence of belief echoes experimentally. Employing manufactured misinformation in a controlled experimental setting allows me to be sure that only the treatment group is exposed to the corrected misinformation.

Although the experiments in this dissertation are structurally similar, they vary in important ways depending on the particular hypothesis being tested. This chapter contains a detailed explanation of the basic experimental format used throughout the dissertation. Variations on this format are discussed separately for the experiments discussed in Chapters 4 and 5. The following section describes the population used for these experiments, which were conducted using Amazon's Mechanical Turk service.

This chapter presents the results of two experiments examining how exposure to political misinformation can have a lingering effect on candidate evaluations and on perceptions of candidate electability. In Experiment 1, I test the hypothesis that belief echoes can be created by successfully discredited information. Like the psychology experiments discussed previously, Experiment 1 imposes a delay between the misinformation and the correction: the misinformation appears in a newspaper article and the correction is embedded in a subsequent article.

In contrast, Experiment 2 examines whether belief echoes can be created even when misinformation is corrected immediately. This experiment speaks directly to how misinformation is corrected in political coverage. Corrections that occur immediately, definitively, and in the context of the article itself have become the "gold standard" of fact-checking among journalists. Discussing the importance

of fact-checking within articles, *On the Media* co-host Brooke Gladstone exhorted her colleagues to “[f]act check incessantly. Whenever a false assertion is asserted, it has to be corrected in the same paragraph, not in a box of analysis on the side.”

(Gladstone, as quoted in Graves & Glaisyer, 2012). If belief echoes are created even when misinformation is corrected immediately, this implies that even journalists’ best efforts may not be enough to eliminate the attitudinal effects of misinformation. Experiment 2 also examines the extent to which belief echoes, like belief persistence, are driven by motivated reasoning.

Table 1 shows the procedure of the two experiments. In each experiment, participants read either one or two articles about a (fictional) congressional race.

Table 3.1. Experimental Procedure

Experiment 1	Experiment 2
Article 1	Article 1
Distractor task	---
Article 2	---
Distractor task	Distractor task
Attitude elicitation	Attitude elicitation
Manipulation check	Manipulation check
Debrief	Debrief

One version of the article contained a piece of misinformation about the candidate that was then corrected, while another contained the misinformation without a correction. The control version contained no misinformation and no correction. In Experiment 1, the misinformation appeared in a separate article from the correction. In Experiment 2, the misinformation and correction appeared in the same article. The implicit assumption of the marketplace of ideas is that once misinformation is successfully corrected, the attitudes that were affected by that misinformation should revert to their pre-misinformation state. In these experiments, that state is represented by the group who is exposed to neither the misinformation nor the correction.

After reading the article(s), participants completed a distractor task⁴ and

⁴ This distractor task, finding the differences between two photographs, was designed to be cognitively involving. See Appendix C for more details.

then evaluated the candidates along a number of dimensions, including a feeling thermometer, traits, and electability assessments. The candidate evaluation index consists of nine variables. The first was a feeling thermometer, recorded on a scale of 0 to 100. Six were traits, also measured on a 0 to 100 scale: shows good judgment, strong leader, trustworthy, dishonest, out of touch, and selfish. Respondents also answered two agree-disagree statements: “John McKenna would make a good representative” and “John McKenna may not be suited for public office.” The electability index consists of three variables: a thermometer assessing respondent’s perceptions of John McKenna’s appeal to Independent voters, a question asking who respondents thought would win the election; and a question asking whose support had increased more in the past week. The full questionnaire appears in Appendix B.

At the end of the experiment, participants were shown five factual statements about the contents of the article and asked whether each one was true or false (“You read two newspaper stories about an ongoing Congressional race. Knowing what you know now, please tell us which of these statements are true”). One of these statements concerned the misinformation (“John McKenna received contributions from a convicted felon”)⁵. Participants evaluated each piece of information on a six-point scale ranging from “definitely false” to “definitely true.” This measure serves two purposes. First, it is a manipulation check, ensuring that the correction was read and processed. Second, and critically for the existence of belief echoes, it confirms that the correction was successful. Full acceptance of the

⁵ In Experiment 2, an additional statement was added to measure belief in the misinformation, “Daniel Elsie donated money to John McKenna’s campaign.” These two measures were correlated at .891.

correction in the treatment group eliminates a key alternative explanation for an observed effect of the corrected misinformation on attitudes: skepticism of the correction.

Type of misinformation used in the experiments

Both experiments employed the same piece of negative information, that a candidate accepted campaign donations from a convicted felon. A pre-test examined the impact of three different types of misinformation (accepting campaign donations from a felon, being arrested for soliciting a prostitute, and cheating on law school exams) and found no significant differences in how they affected evaluations (Thorson 2011). The pre-test also confirmed that learning a candidate had accepted campaign donations from a convicted felon lowered evaluations equally among both Democrats and Republicans.

The claim that a candidate donated money to a felon is consistent with the types of misinformation found in real-world politics. Combing through the opposition's campaign financial disclosure statements in search of potentially damaging donations is a common strategy for candidates of both parties. In a 2012 Connecticut House race, the Democratic candidate was criticized for accepting donations from a financier convicted of insider trading (Fenster 2012), an allegation that turned out to be true. In Tennessee, the Democratic Congressional Campaign Committee issued a press release with an even more colorful accusation, stating that "foreign prostitution money is allegedly behind the groups funding Congressman Scott DesJarlais's (TN-04) Republican Majority." A few days after the press release

was issued, the fact-checking organization Politifact declared the accusation false (Politifact 2012).

Because the most real-world political misinformation is negative rather than positive, the experimental manipulations in this dissertation employ only negative misinformation. Although the underlying mechanism should be similar for both types positive and negative information, I would expect that negative information would create a stronger belief echo given the asymmetric impact of negative and positive information on evaluations (Fiske 1980).⁶ Chapter 7 discusses in more detail the potential for belief echoes to be generated by positive misinformation.

Population

The experiments in this dissertation were conducted between August 2011 and November 2012. Participants accessed the experiment through Amazon.com's Mechanical Turk platform.⁷ Mechanical Turk is an online platform run by Amazon that enables people to recruit and pay subjects to perform tasks that range from captioning photos to taking surveys. The experiment was restricted only to U.S. participants over the age of 18. Although Mechanical Turk includes a large non-U.S. population, the system can screen out subjects not registered from the United States. In addition, I.P. address matching ensured that the subjects were indeed from the U.S. Participants were paid between \$0.61 and \$0.75 for their participation

⁶ However, a recent paper by Cobb, Nyhan, & Reifler (2010) finds that when positive information is shown to be false, it can produce a backfire effect, in which individuals "over-correct" and ultimately view a candidate more negatively than they would have if they had never seen the misinformation.

⁷ The task was described in language similar to the following: "Read two short articles and answer questions about current events and celebrities. Takes about 10 minutes." Only MTurk workers whose approval rate for previous MTurk work was over 90% were eligible to participate.

(payment varied based on the length of the survey) and the surveys took most people between 8 and 12 minutes to complete. All participants were screened to ensure that none took the experiment twice.⁸ A table describing the demographic characteristic of each sample is available in Appendix D.

While Mechanical Turk is a relatively new platform on which to conduct social science experiments, several studies suggest that it is a more than adequate substitute for other traditionally-used convenience samples. A comprehensive analysis by Berinsky, Huber, & Lenz (2012) shows that Mechanical Turk is more representative of the U.S. population than in-person convenience samples. They also compare the results of similar experiments conducted on three different types of samples: internet panel surveys (such as Knowledge Networks), typical social science convenience samples, and face-to-face surveys (such as the ANES). The experiments run using Mechanical Turk participants generate similar estimates of average treatment effects as more traditional samples. In addition, they replicate the results of three benchmark social science experiments using a Mechanical Turk sample.⁹

Although the Mechanical Turk sample is not representative, it is more similar to the general population than most convenience samples (for example, college freshmen). More importantly, there is no theoretical reason why I would expect the ways in which the MTurk sample differs from the general population to affect my

⁸ An programmatic screener developed by Eyal Peer was used to match participants' unique Mechanical Turk ID to the IDs of subjects who had previously taken the experiment. The full procedure is available here: <http://experimentalturk.files.wordpress.com/2012/02/screening-amt-workers-on-qualtrics-5-2.pdf>

⁹ Rasinski (1989)'s question wording experiment examining support for "welfare" versus "assistance to the poor"; Tversky & Kahneman's (1981) Asian disease gain-frame/loss-frame experiment; and Kam & Sim's (2010) study on risk orientation and policy choice.

experimental results. The Mechanical Turk sample tends to be slightly younger and more liberal than the general population. Neither of these characteristics should affect the magnitude of belief echoes. Although partisanship does play a role in the creation of the belief echoes, the experiments in this dissertation manipulate partisanship such that subjects are presented with misinformation that either reinforces or contradicts their pre-existing partisan preferences.

Experiment 1: The existence of belief echoes

Experiment 1 examines whether exposure to corrected misinformation creates belief echoes in the context of a delayed correction . I predict that exposure to the misinformation will continue to negatively affect attitudes and perceptions of electability even when it is effectively corrected.

H1. Belief echoes affect evaluations: exposure to negative information will affect evaluations even after it is corrected.

H2. Belief echoes affect perceptions of electability: exposure to negative information will affect evaluations even after it is corrected.

In this experiment, subjects in the “corrected misinformation” condition experience a delay between reading the misinformation (which appears in the first article they read) and the correction (which appears at the end of the second article). During this period, they also complete a short distractor task (details in

Appendix C1). This “delayed correction” format parallels the belief perseverance experiments in psychology. In these experiments, subjects are given time to think about and process the false claim before it is discredited.

Testing for the presence of belief echoes requires passing two separate manipulation checks. First, reading negative information about the candidate must lead to lower evaluations. Establishing this baseline effect is necessary in order to contextualize results and estimate the extent to which corrections reduce the magnitude of any negative attitudinal effects. As such, a manipulation check ensures that negative information about a candidate negatively affects evaluations.

Second, the correction must be believed. In order to provide convincing evidence that attitudinal effects of the misinformation are due to belief echoes rather than to skepticism of the correction, it is necessary to demonstrate that the correction is successful in reverting people back to their initial belief (represented in the experiment by the group who never saw the misinformation). In addition, the negative information must be believable when uncorrected.

Showing that the misinformation in this experiment was successfully corrected is especially important given past research on misinformation suggesting that motivated reasoning plays a large role in the success or failure of corrections. Specifically, partisans often reject corrections when the misinformation is something that they are predisposed to believe (Nyhan & Reifler 2010). In this experiment, all participants were assigned to read misinformation that confirmed their partisan preferences, the precise situation in which successfully correcting misinformation is most difficult. The manipulation check thus must also verify that

reading the correction leaves people with the same level of belief in the negative information as those who never read it at all, and that both of these groups are more skeptical of the negative information than those who did not read a correction.

Experiment 1 Design

In this experiment, 157 participants were randomly assigned to one of three groups in a three-condition between-subjects experiment. The three conditions were:

- (1) Uncorrected Misinformation (misinformation, no correction)
- (2) Corrected Misinformation (misinformation and correction)
- (3) Control (no misinformation, no correction)

Each group read two articles about a Congressional race. The Uncorrected Misinformation group (N=47) read a piece of misinformation, and did not read a correction in the second article. The Corrected Misinformation group (N=54) saw the misinformation in the first article, then a correction in the second. The control group (N=56) read the same two articles but with the misinformation and the correction omitted.

First, all participants answered several basic demographic questions (age, education, party identification, and political interest) as well as a three-question battery measuring their political trust ($\alpha = .685$). Other demographic information about the sample is listed in Appendix C. All participants were classified as leaning Democratic or Republican through a series of branching questions (details in Appendix B).

Next, participants read an article that included information about the race and brief biographical details about the candidates. The article read by the Corrected Misinformation and Uncorrected Misinformation groups also included a statement saying that one of the candidates had accepted donations from a convicted felon. For all participants, the partisanship of the candidate featured in the article varied depending on their own partisanship: the version given to Republicans described a Democratic candidate and vice versa. The candidate's party was not an experimental manipulation: all participants read an article featuring a candidate of the opposing party.

After completing a short distractor task, participants read a second article about the race. For those in the Corrected Misinformation condition, the following correction was appended to the article:

Correction: Regarding the candidates' biographies, an article published in the *Kansas City Star* last week stated that Daniel Elsio, a felon convicted of drug trafficking and murder, was a frequent donor to John McKenna's campaign and attended several [Republican/Democratic] party fundraisers. However, further investigation of the campaign donation records has shown that the donor listed was actually Daniel Elio, the owner of a local car dealership.

The full text of both articles is in Appendix A. Next, participants completed a second distractor task, and then evaluated the two candidates discussed in the article.


Finally, participants indicated where several pieces of information, including the

claim that the candidate accepted donations from a convicted felon, fell on a scale from “definitely false” to “definitely true.”

Experiment 1 Results

The first manipulation check ensures that exposure to the information that a candidate accepted donations from a felon negatively affects subjects’ evaluations. An analysis of variance demonstrated a significant main effect of exposure to negative information on evaluations ($F=13.8, p < .001$)¹⁰. The second manipulation check makes certain that participants found the correction convincing. Figure 3.1 shows the extent to which participants in each condition believed both the factual information and the misinformation presented in the article. Answers of “definitely false” were coded as 0 and “definitely true” as 5.

¹⁰ Political trust ($F=29.0, p<.001$) and strength of partisanship ($F=26.4, p<.001$) were also significant predictors of candidate evaluation and so were entered as covariates.

Figure 3.1. Experiment 1: Effectiveness of the correction


As expected, individuals who received the misinformation without the correction were the most likely to believe that the statement “John McKenna received campaign donations from a convicted felon” was true, scoring it a 4.57 on a scale from 0 to 5. The correction was highly effective, cutting that to 1.76. The control group (who never saw the misinformation or the correction) and the corrected misinformation group did not significantly differ in their estimate of the statement’s veracity. Both groups differ significantly from the “misinformation only” group. The second manipulation check shows that the correction was successful.

Hypothesis 1 directly addresses the existence of belief echoes: attitudinal effects that persist even after a piece of negative information has been successfully corrected. If negative information continues to affect attitudes even after it is

corrected, the candidate evaluations of the Corrected Misinformation condition should be lower than those of the control condition.

Figure 3.2 displays the mean evaluation of the candidate by condition.¹¹

Figure 3.2. Experiment 1: Effect of misinformation on attitudes




The Corrected Misinformation group's evaluation of McKenna is 4.94, while the control group's is 5.79. An analysis of variance with a post-hoc comparison demonstrated that there is a significant difference between the two conditions ($p < .05$). Individuals who received the misinformation and then read the correction evaluated the candidate significantly more negatively than those who read neither the misinformation nor the correction, despite the fact that those who read the

¹¹ Evaluation was measured with a nine-item index ($\alpha = .924$), for details see p. 12

correction found it convincing. In terms of effect size, the correction is successful at muting less than half of the attitudinal effects of exposure to misinformation. The results of this experiment show that belief echoes exist: even when a false claim is successfully discredited, it continues to affect attitudes.

Figure 3.3. Experiment 1: Effect of misinformation on electability perceptions shows mean perceptions of electability for each condition.¹² These results are testing H2, which proposes that exposure to misinformation affects perceptions of a candidate's electability, even when the misinformation is successfully corrected.

Figure 3.3. Experiment 1: Effect of misinformation on electability perceptions



An analysis of variance shows that perceptions of McKenna's electability in the uncorrected and correction misinformation conditions do not differ significantly

¹² Perception of electability was measured with a three-item index ($\alpha = .530$), for details see p. 12

from each other. However, perceptions in both conditions are significantly lower than perceptions in the control condition ($F = 5.05, p < .05$). Even though subjects in the corrected misinformation condition fully accepted the correction, they still perceived the candidate as less electable than subjects in the control condition.

Experiment 2: Belief time and motivated reasoning

The results from Experiment 1 demonstrate that belief echoes can be created by exposure to successfully corrected misinformation. However, it leaves open the question of *how* these are created. This experiment examines two potential explanations: belief time and partisanship.

Experiment 1 included a time delay between the misinformation and the correction. The delay could be up to two minutes, depending on the length of time participants spent on the distractor task. It is possible that during this time, the participants thought about the misinformation and used it to activate similar considerations. For example, Fleming and Arrowood (1979) find that when the misinformation is corrected immediately, belief perseverance effects are minimized. They argue the passage of time between the misinformation and the correction provides an opportunity for reflection on the misinformation. As a person considers the misinformation, he or she may activate related considerations. For example, if the initial misinformation concerned a salmonella outbreak at a local seafood restaurant, this may prompt someone to remember other experiences with food poisoning or recall news stories about unsafe fish. These thoughts may then continue to influence their attitudes about seafood restaurants even after the initial

misinformation has been discredited. In contrast, an immediate correction of the misinformation (for example, in the next sentence) drastically narrows the temporal window for these thoughts to arise.

If immediate correction eliminates or even significantly minimizes belief echoes, this has important implications for journalists seeking to correct misinformation. Over the last five years, some journalists have urged their peers to move towards a standard of, whenever possible, including a correction immediately after the misinformation (Fallows 2012). This changing norm raises the question of whether corrections that conform to this “gold standard” are more effective at minimizing misinformation’s effects.

Experiment 2 has two major goals. First, to determine whether belief echoes can be created even when misinformation is corrected immediately, a context often considered the ideal form of journalistic fact-checking. Second, to investigate the relationship between belief echoes and partisanship.

Partisanship how people process, recall, and draw on information to form attitudes. Indeed, most existing research on misinformation has focused on how partisanship causes people to accept or reject corrections (Nyhan & Reifler 2010, Bullock 2007, Kuklinski 1998). Experiment 1 showed that it is possible to correct misinformation even when that misinformation contradicts a person’s partisan preferences. However, partisanship may still play an unconscious role in the creation of belief echoes. For example, a Democrat reading a piece of negative information about a Republican may experience a strong immediate negative reaction, leading her to drastically lower her evaluation of the candidate, while a

Republican reading the same information will likely experience a weaker initial reaction and thus a smaller change in evaluation (Goren 2002). Because the size of their initial responses varies, so might also the belief echo that this response creates.

While the basic structure of Experiment 2 is similar to that of Experiment 1, the format of the misinformation and correction is altered. Experiment 1 followed a structure similar to tests of belief perseverance in psychology experiments, in which individuals believe the misinformation for a period of time before it was corrected. Experiment 2 eliminates the two-article format and presents the correction immediately after the misinformation. This “immediate correction” format more closely parallels how corrections are increasingly presented in the media.

Hypotheses 3 and 4 predicts that belief echoes will be created even when misinformation is corrected instantly.

H3: Exposure to misinformation creates belief echoes: exposure to corrected misinformation causes lower candidate evaluations, even when the misinformation is corrected immediately.

The next hypothesis deals with the role of partisanship. The results of Experiment 1 suggested that exposure to negative information creates belief echoes. However, because all participants were assigned to read articles about a candidate of the opposing party, it was impossible to measure the extent to which partisanship might be a factor in creating belief echoes.

A person's party identification may play a conscious or unconscious role in the creation of belief echoes. In Experiment 1, participants who read the correction claimed not to believe the discredited information. It is possible, however, that this response was attributable to an experimenter demand effect, and their "genuine" belief was that the misinformation was true. Motivated reasoning may also be a factor at an unconscious level. A piece of negative information that confirms pre-existing beliefs may activate similar beliefs, and thus affect the candidate's overall evaluation. For example, if a Democrat learns that a Republican candidate accepted campaign donations from a felon, it may immediately bring to mind other examples of corrupt Republicans and negative thoughts about the Republican party more generally (Redlawsk 2002). Even after the initial information is recognized as invalid, the other activated concepts may continue to exert an unconscious effect on her evaluation of the candidate.

To test the effect of partisanship on belief echoes, Experiment 2 varies whether the misinformation confirms or counters the participant's party identification. Hypothesis 2 predicts that belief echoes will be stronger when the misinformation reinforces pre-existing partisan attitudes.

H4: Belief echoes will be stronger when the misinformation reinforces partisan attitudes.

Experiment 2 Design

A total of 474 people were recruited via Amazon's Mechanical Turk. The demographics of the sample are available in the Appendix. The experiment was a 3 (misinformation format) x 2 (candidate party) Participants were assigned to one of six conditions. The full design is presented in Figure 3.4. Experiment 2: Design

Figure 3.4. Experiment 2: Design

		MISINFORMATION FORMAT		
CANDIDATE PARTY		Corrected misinformation	Uncorrected misinformation	No misinformation
	Candidate is opposing party as subject	N=113	N=61	N=66
	Candidate is same party as subject	N=111	N=60	N=63

Participants in Experiment 2 read only one article. That article contained within it both the misinformation (for the Corrected Misinformation and Uncorrected Misinformation groups) and the correction (for the Corrected Misinformation group). While the misinformation was identical to that in Experiment 1, in Experiment 2 the source of the misinformation was clearly identified as the opposing campaign (see Appendix A for full text of article). For those in the "Corrected Misinformation" condition, the correction followed directly after the misinformation and read as follows:

However, further investigation of the campaign donation records by [journalists at the *Iowa Ledger*/the independent fact-checking organization

GetTheFacts.org] has shown no record of any donation from Elsio to McKenna's campaign. Campaigns are required to disclose the names of all individuals who contribute \$200 or more in an election cycle, and [the *Ledger*/GetTheFacts.org] did not find Elsio's name listed.


A separate condition, discussed in more detail in Chapter 6, varied the source of the correction. Two different sources, The *Iowa Ledger* and GetTheFacts.org, were equally successful at correcting the misinformation and so were combined for the analyses presented here.

Experiment 2 Results

As in the last experiment, it was necessary to first demonstrate that the correction was successful.

Figure 3.5 shows the results of this manipulation check, presenting the average belief in the correction for each group.¹³ The middle column represents the combined opposing-party and same-party conditions.

¹³ In Experiment 2, belief in the correction was assessed using two measures instead of one. The correlation between the two measures is .741, $p > .001$.


Figure 3.5. Experiment 2: Effectiveness of correction

The manipulation check shows that the correction was highly successful. Both corrected misinformation group and the control group were significantly less likely to believe the misinformation than those who did not receive the correction ($F=14.4, p < .001$). This held true for the same party, opposing party, and combined groups. In the opposing party condition, the belief in the misinformation was not significantly different from the control group, suggesting that the correction was entirely successful. In the same party condition, the correction was even more successful: those who received the correction were significantly more skeptical of the misinformation than those in the control group ($F=5.9, p < .05$).

In Experiment 1, belief echoes were created in a context in which participants actively believed the misinformation for several minutes before it was

corrected. H3 predicts that belief echoes will be created even when the misinformation is corrected immediately. Figure 3.6 presents the mean candidate evaluations for each group.

Figure 3.6. Experiment 2: Effect of misinformation on attitudes



The results shown in the “Combined” column visually presents the test of H3.

Overall, participants who received corrected misinformation evaluate the candidate significantly more negatively than those in the control condition ($F=13.2, p < .05$), despite being equally certain that the misinformation is false. A successful

correction mutes slightly over half of the attitudinal effects of exposure to misinformation.¹⁴

There is a strong and significant main effect of candidate party ($F=216.9, p < .001$): in every condition. In each group, participants provided lower evaluations of the candidate when he was of the opposing party. However, an analysis of variance shows no significant interaction between candidate party and exposure to corrections. H4 is not supported—belief echoes are not conditional on or significantly magnified by partisanship. This pattern suggests that the presence of belief echoes cannot be explained solely as a process of motivated reasoning.

Generalizability

As with any experiment, it is useful to ask how these findings generalize to the real world. The experiments were designed to be a relatively difficult test of the hypothesis. The correction was unambiguous and the false information was presented as an aside rather than as the focus of the article. Thus, they may actually underestimate the aggregate attitudinal effects of misinformation. Real-world misinformation may be even more conducive to creating belief echoes for two reasons: the difficulty of falsifying claims and the newsworthiness of misinformation.

To test for the presence of belief echoes, the experiment employed information that was capable of being verified with relatively objective data as well as a highly believable correction. These experiments are in essence a “best-case”

¹⁴ As in Experiment 1, exposure to the corrected misinformation also has a significant negative effect on perceptions of electability.

scenario of how misinformation might be corrected, and as such the most difficult test for finding attitudinal effects of misinformation. In the real world, most political claims are far more difficult to falsify. Some are prospective claims about what a policy will or will not do. For example, many politicians rely on Congressional Budget Office projections in discussing the costs of future policies. But these numbers are only estimates, and predictions about future costs are by nature probabilistic rather than definite, thus making them difficult to thoroughly debunk. Other claims are difficult to falsify for more obscure reasons, such as, Senator Harry Reid's 2012 assertion that Mitt Romney has not paid his taxes in ten years. Politifact called this statement a "pants on fire" lie, but Reid insists he has a reliable source whose name he cannot reveal (Politifact 2012). In cases like these, uncertainty might magnify the effects of misinformation on attitudes.

In addition, the misinformation in these experiments was presented as an aside rather than as the focal point of an article. In the real world, misinformation—even when it is being corrected—is often featured much more centrally. For example, a recent *New York Times* editorial entitled "Truth and Lies About Medicare" was dedicated entirely to correcting Mitt Romney's false claims about the Affordable Care Act. In the process, it repeated each of the claims (*New York Times* 2012). Although these efforts to publicize its corrections of false claims may be well-intentioned, they may also serve to amplify belief echoes by making the misinformation more salient.

Potential consequences of belief echoes

The existence of belief echoes presents a serious challenge to the marketplace of ideas, which holds that truth emerges from the free competition of information. The theory implicitly assumes that once a piece of information has been discredited, it will cease to affect attitudes and preferences. The experiments in Chapter 4 present participants with an idea—in this case, a factual assertion about a candidate—that is then successfully discredited. In the marketplace of ideas created within this experiment, the idea is defeated, even among those whose political leanings predispose them to want to believe it. Despite its rejection, the idea continues to exert an effect on attitudes.

Corrected misinformation may also affect perceptions of candidates through other channels than evaluation. In both, the misinformation also affected perceptions of electability. This pattern may be driven by similar mechanisms to those underlying belief echoes, but these mechanisms are likely amplified by the third-person effect. Voters who read and accept a correction might assume that others will not be equally accepting (or may not see the correction at all). They may thus perceive the candidate as less viable. Given the importance of perceived electability to vote choice, especially in primaries, the existence of misinformation about a candidate may have real consequences for electoral outcomes even when it is successfully corrected (Abramson et al 1992).

These findings have serious implications for both politicians and journalists. For politicians, it provides evidence that the “throw it all out there and see what sticks” approach to smearing one’s opponent may be even more successful than

previously imagined. For journalists, it suggests that careful fact-checking is not sufficient to eliminate the consequences of reporting false claims. Minimizing belief echoes appears to require journalists to fact-check *before* publishing a claim and resist the temptation to publicize false statements. However, there is an economic disincentive for this type of caution, as it is the most outrageous claims that generate the most excitement, attention, and—critically for online media—web traffic (Carr 2012). Chapter 5 directly investigates the extent to which belief echoes shift the incentives facing journalists and politicians.

Fundamentally, the existence of belief echoes poses a challenge to the marketplace of ideas by suggesting that misinformation can continue to affect attitudes even after it is successfully corrected. The next chapter investigates the processes that create belief echoes and begins to develop practical strategies for minimizing them.

Chapter 4 - Affective and cognitive processes of belief echo creation

Fact-checking has increasingly become a major part of the media's political coverage (Graves 2012). The *New York Time's* David Carr called the 2012 election "the most fact-checked election in history" (Carr 2012). For the most part, these efforts have been seen as a normatively positive development: in the wake of the election, fact-checking organizations and media outlets alike were applauded for their emphasis on finding, correcting, and publicizing false claims (Fallows 2012). However, Chapter 3 of this dissertation suggests that the publicizing of corrected claims may also have serious unintended consequences, especially in a media environment focused on publicizing false claims: even when corrections are successful, misinformation can have lingering effects on attitudes.

Is it possible to correct false claims *without* creating belief echoes? In this chapter, I argue that developing strategies for reducing the impact of belief echoes requires understanding how they are created. This chapter explores two possible mechanisms for the creation of belief echoes: an affective process and a cognitive process. These processes are not necessarily mutually exclusive, and may be activated by different factors. The affective mechanism for belief echoes suggests that they are created as a largely unconscious byproduct of online processing, when negative information exerts a stronger impact on a summary evaluation than does its correction. The cognitive mechanism describes a conscious cognitive process during which a person recognizes that a particular negative claim about a candidate

is false, but reasons that its presence increases the likelihood of other negative information being true. This chapter describes in greater detail the processes that might drive the creation of affective and cognitive belief echoes, and presents the results from experiments designed to isolate both processes.

The affective process of belief echo creation

The online processing model suggests that when a person processes a piece of information relevant to a candidate (or any other political object, such as a policy), it produces a spontaneous affective response (Lodge, Taber, & Weber 2006; Fazio et al 1995). The response is immediately and unconsciously integrated into a person's evaluations of the candidate as he "spontaneously extracts the affective value of the message and within milliseconds updates his summary evaluation" (Lodge, Taber, & Verhulst 2011 p 283). For example, hearing that a candidate was arrested produces a spontaneous negative response that then influences a person's evaluation of the candidate. In contrast, learning that a recently enacted policy successfully reduced unemployment produces a positive response, which then shapes attitudes towards that policy.

The fundamental premise of online processing is that even after a piece of information is forgotten, it can continue to influence attitudes. The independence of attitudes from memory is elegantly demonstrated by Coronel et al (2012) in a study employing subjects suffering from a type of amnesia that prevents them from forming new memories. All participants had opinions on major political issues (formed before the amnesia), but they were unable to retain any new information

for more than a few minutes. Subjects were shown photographs of two candidates and told each candidate's issue positions. After twelve minutes, they were shown the photographs of the candidates again (without the associated issue positions) and asked to indicate which candidate they would vote for. Every participant expressed a preference for the candidate who was initially described as sharing their issue opinions. However, when immediately afterwards subjects were asked about candidates' issue positions, they had no memory of their stances. Instead, they justified their choice with reasons like "he just seems more likeable" and "he looks more trustworthy." (Coronel et al, p. 9)

This lingering effect on evaluations occurs because the affective response generated by the initial information continues to exert an influence on the subjects' summary evaluation of the candidates even though the information itself has disappeared from memory. This study sheds light on how to understanding how belief echoes can be created through an affective process. Reading the initial misinformation (for example, that a candidate accepted donations from a convicted felon) generates a strong affective response. The correction, however, does not generate an affective response of an equal and opposite magnitude. A person's summary evaluation thus remains affected by the initial affective response generated by the misinformation—even though when asked, she is still quite capable of recalling the correction.¹⁵ However, this description of online processing

¹⁵ If the correction remains in a person's working memory, she can use it to deliberately inform her reported attitude, consciously adjusting her summary evaluation to take the correction into account (Wilson, Lindsey, & Schooler 2000). But once the misinformation and correction disappear from working memory (for example, after completing the distractor task), the correction disappears from working memory.

as it relates to affective belief echoes leaves a central question unanswered: when misinformation is corrected, why is it that the misinformation exerts a lasting effect on evaluations, but the correction does not?

Asymmetric affective effects of misinformation and correction

Affective belief echoes can best be understood as a failure of online processing in which a correction does not carry enough affective weight to adjust for the negative impact of the misinformation. In other words, misinformation and corrections are often not affectively symmetrical. Experimental research has demonstrated that this asymmetry exists across several domains.

In an experiment conducted by Sherman and Kim (2002), subjects learned to associate Chinese ideograms with English words that had either a negative or positive valence. They were then informed that the initial meanings were incorrect, and learned new, neutral meanings. Despite learning the new associations, the initial negative and positive affect generated by the initial associations continued to shape their reactions to the ideograms. The effect was only extinguished when, in the second round of learning, subjects did not memorize neutral meanings, but rather learned new meanings that had an affective valence *opposite* of the initial learned valence.

Another experiment (Gawronski et al 2008) shows that correcting a stereotype by repeatedly negating it is not enough to dampen its effects. First, all subjects took an implicit association test measuring implicit racial prejudice. Subjects were then “trained” to counter racial stereotypes either by learning to

activate concepts in direct opposition to the stereotype (affirmation training) or by learning to associate “wrong” with existing stereotypes (negation training). Both of these trainings followed a similar structure. On the computer screen, subjects saw a face (either black or white) paired with a positively or negatively valenced word. In the affirmation training condition, subjects were asked to press a button saying “YES” when counter-stereotypical pairings were made (e.g. a black face with the word “ambitious”). If a stereotypical-congruent pairing appeared (e.g. a black face with the word “shiftless”) they were to do nothing, and the next pairing would appear in a few seconds. The negation training followed a similar format, but with one key difference: subjects were asked to press a “NO” button when stereotypical pairings were made (e.g. a black face with the word “dishonest”), and do nothing when counter-stereotypical pairings were made. After completing the training, subjects took the implicit association test again. The results showed that the negation training significantly strengthened the stereotypes, while the affirmation training significantly reduced them.

These two studies offers several clues that can help us to better understand affective belief echoes. First, the results of the ideogram study (Sherman & Kim 2002) demonstrate that the path through which discredited information can continue to affect attitudes can be entirely driven by affect rather than by cognition. Because the subjects had no prior knowledge of the ideograms, there was no logical reason to connect a particular image to a particular definition.¹⁶ Second, the experiments offer evidence that affective belief echoes can be minimized if the

¹⁶ The ideograms assigned to the positive or negative definitions were randomly varied to ensure that the aesthetic appeal of the ideograms did not affect the results.

“correction” is valenced in the opposite direction as the misinformation. In both studies, the affective charge of the false information was only eliminated when the correction carried an equal and opposite affective charge.

Correcting a piece of false information by negating may not only be insufficient to compensate for its affective impact, but may in some circumstances negations serve to reinforce the effect of the initial misinformation (Hasson & Glucksberg 2005). Because of how negations are mentally encoded, the affective response produced by a negation can be the *opposite* of what is intended. To use a canonical example, processing the statement “Richard is not a crook” may produce an affective response anchored to criminality rather than to honesty (Mayo et al 2003).

The implication of these findings for belief echoes is that framing a correction as an affirmation (for example, “McKenna’s campaign donors are all citizens of good standing”) rather than as a negation (“McKenna did not accept donations from a felon”) may be more effective at reducing belief echoes, because the affirmation carries a positive affective charge that can counterbalance the negative impact of the misinformation.

“Cognitive busyness” and affective belief echoes

Fundamentally, affective belief echoes are attributable to an updating problem. They occur when a correction does not cause people to fully update their evaluations, which have previously been altered by the misinformation. This is a largely automatic process that occurs outside of conscious control. When people are

what Gilbert (1989) calls “cognitively busy,” their reactions and attitudes are more likely to be driven by such automatic processes. As, such affective belief echoes should be more likely under these conditions.

A study by Gilbert and his colleagues (1993) demonstrates how the impact of information known to be false is magnified under conditions of cognitive load. In the experiment, subjects were told that to play the part of a trial court judge and exposed to a series of statements about a criminal incident. They were informed that statements printed in black were true and those printed in red were false. Several of the false statements exacerbated the seriousness of the crime. Half of the subjects were exposed to the statements while simultaneously completing a number memorization task. Finally, subjects were asked to consider the facts of the crime and recommend a prison term length for the accused. The subjects who were under cognitive load when reading the information recommended longer sentences for the accused: they were influenced by the false exacerbating statements printed in red. However, questioning the subjects showed that they were fully aware that the statements were false and unaware that they were affecting their decision.

In real-world settings, then, affective belief echoes may be more likely when readers’ attention is elsewhere. For example, a Pew study conducted after the 2012 Presidential election showed that of Americans who watched the news on election night, 27% were “dual screeners,” using both the television and the internet to follow the returns (Pew 2012c).

The cognitive process of belief echo creation

“Where there’s smoke, there’s fire” is a useful metaphor for understanding the cognitive mechanism of belief echo creation. Cognitive belief echoes are created through a conscious process in which a person accepts a correction, but also infers that the existence of the false claim suggests that other negative information about the candidate or policy is more likely to be true. This in turn leads him to evaluate the candidate more negatively.

Belief perseverance and mental models

The process that drives cognitive belief echoes is similar to the mechanism underlying the phenomenon of belief perseverance, much studied in psychology and discussed in more detail in Chapter 3 of this dissertation. Belief perseverance occurs when a piece of information continues to influence an individual’s beliefs about the world, even after they recognize that the information is false (Anderson, Lepper, & Ross 1980; McFarland, Cheam, & Buehler 2007). Belief perseverance as it is conceptualized by psychologists differs from belief echoes, primarily because it is concerned with how discredited information affects causal beliefs about the world. In contrast, belief echoes are concerned with how false claims affect evaluations of people or objects. However, past work studying belief perseverance can shed light on the process underlying cognitive belief echoes.

In one common experimental demonstration of belief perseverance, subjects are given either positive or negative feedback about their performance on a task (for example, distinguishing between fabricated and real suicide notes). Later, they are told that the feedback was false and the experiment was designed to measure their

physiological reactions to different types of feedback. Despite this debriefing, when subjects are asked to predict their actual performance on the task, those who were originally given positive feedback tend to think they did better than average, and those who were given negative feedback think they did worse (Ross, Lepper, & Hubbard 1975).

One explanation for this effect is that exposure to the initial results caused subjects to think about the possible reasons for the negative or positive scores (Anderson, New, & Speer 1985). When the scores were discredited, the reasons remained. This explanation for belief perseverance is often referred to as the “mental model” explanation. The misinformation induces people to build so-called “mental models” that persist even after the misinformation is discredited. In the case of the misinformation in Chapter 3, hearing that a candidate accepted a donation from a felon might cause someone to think about other ways in which the candidate might be the type of person who accepts money from a felon: he is untrustworthy, he is corrupt. These reasons not only continue to affect evaluations even after the claim is discredited, but may even become a justification for why the candidate was falsely accused in the first place.

Cognitive belief echoes as a type of correspondence bias

In cognitive belief echoes, the “mental model” created by the false claim one in which the subject of the misinformation is seen as “deserving” of the accusation in some way. This effect can also be understood as type of correspondence bias.

Correspondence bias occurs when an observer "draws inferences about a person's

unique disposition from behaviors that can be entirely explained by the situations in which they occur" (Gilbert 1995). In the experiments in this dissertation, the situation could be explained by a number of factors—faulty opposition research, the desire of a candidate to smear his opponent, or bad record-keeping. Cognitive belief echoes are created when a person ignores these situational factors and instead attributes the accusation to a fault in the accused candidate's character.

Research suggests that correspondence bias is lessened when people process misinformation under conditions of suspicion, defined as a state in which "the individual actively entertains multiple, plausibly rival hypotheses about the motives or genuineness of a person's behavior" (Fein et al 1997). Suspicion lessens the likelihood that perceivers will make dispositional inferences from behavior (i.e. engage in correspondence bias), and more likely that they will attribute a given behavior or incident to situational factors (for instance, poor reporting or a deliberate smear campaign). As such, cognitive belief echoes may be less likely when subjects are suspicious of the motivation behind the false claim. In contrast, cognitive belief echoes may be *more* likely when subjects either have pre-existing reasons to distrust the accused candidate, or pre-existing reasons to trust the accuser. Partisanship may provide exactly these reasons.

Using partisanship to distinguish between affective and cognitive processes

The previous section suggests that belief echoes may be driven by either an affective or a cognitive process. Which processes are activated will depend on how

the information is processed, which itself is a function of both the information environment and characteristics of the individual.

The experiments in this chapter are designed to distinguish between affective and cognitive belief echoes. To do so, they take advantage of a unique aspect of the political world: the fact that partisans tend to be skeptical of claims made by the opposing party and trusting of claims made by their own party (Hetherington 2001). This predisposition can help to distinguish between affective and cognitive belief echoes. Partisanship should be a stronger predictor of belief echoes in situations likely to produce cognitive belief echoes (usually those that invoke more cognitive processing), and a weaker predictor of belief echoes in situations likely to produce affective belief echoes (those that invoke less cognitive processing).

Cognitive belief echoes are generated through a conscious process in which a person infers that an accusation against a candidate—even if false—increases the likelihood that the candidate is untrustworthy. Partisanship should heighten this effect since the candidate will already be perceived as less trustworthy by virtue of being a member of the opposing party. When the misinformation concerns a candidate or policy of a person's own party, cognitive belief echoes should be minimized because her basic orientation towards the political object is positive rather than negative.¹⁷ These chains of reasoning—which are heightened by

¹⁷ The opposite should be true if the misinformation was positive: belief echoes would be more likely when the positive misinformation was about a person's own party, and vice versa. See the conclusion of this dissertation for more on positive misinformation.

partisanship—should be activated when a person engages in deliberate, conscious thought about the misinformation and correction.

In contrast, partisanship should matter only minimally in the creation of affective belief echoes, which are heightened in conditions of shallow consideration of the misinformation and correction.¹⁸ For affective belief echoes, the magnitude and direction of the affective response generated by the misinformation (which in turn affects evaluations) will be largely a function of the misinformation itself (for instance, its vividness). Negative information produces a negative response, and positive information produces a positive response. Of course, partisanship does color the automatic responses that individuals experience to political information (Burdein, Lodge, & Taber 2006). However, I expect the effect of partisanship to be relatively weaker when conscious considerations about the candidate's party are not simultaneously brought to mind.

Both of the experiments in this chapter take advantage of the unique role of partisanship by randomly varying the party of the accused candidate. In the “same party” condition, the accused candidate is of the same party as the subject (and the accuser of the opposing party). In the “opposing party” condition, the accused is of the opposing party (and the accuser of the subject's own party).

Experiment 3 attempts to generate cognitive belief echoes by creating an environment that encourages more cognitive elaboration. Instead of embedding the misinformation and correction in the text of an article, both are presented as headlines. Subjects are encouraged to “read each headline carefully” as well as study

the accompanying photograph. Thinking deeply about each headline (including the misinformation and correction) should create stronger belief echoes among those who read misinformation about a candidate of the opposing party, and weaker belief echoes among those who read misinformation about a candidate of their own party.

Experiment 4 follows a similar structure to Experiment 2 in Chapter 3, in which subjects read misinformation and correction embedded in an article. However, after they read the article, half of the subjects are randomly assigned to a task designed to encourage cognitive elaboration about the information they just read. Among subjects not assigned to the task, belief echoes should be created through the affective process, and thus occur regardless of partisanship (as in Chapter 3 of this dissertation). In contrast, when subjects are encouraged to think carefully about the claim and correction, belief echoes should be created through the cognitive process, and thus be conditional on partisanship. Note that throughout the descriptions of the experiments, references to the “party” condition should be understood not as Republican versus Democrat, but as “same party” (the candidate is the same party as the subject) and “opposing party” (the candidate is the opposing party as the subject). All subjects were assigned to one of these two conditions.¹⁹

¹⁹ While many subject initially identified as Independent, they were led through a series of branching questions to identify as either leaning Democratic or Republican (see Appendix B for details).

Experiment 3: Creating cognitive belief echoes through elaboration

In Experiment 3, I expose subjects to misinformation and corrections in a format designed to heighten cognitive belief echoes. This experiment is substantially different from the format of the experiments in Chapter 3 of this dissertation. In those studies, both the misinformation and correction was embedded within the text of a newspaper article. In contrast, in Experiment 3, the misinformation and correction are in newspaper headlines. Subjects are encouraged to read each headline carefully and study the images. The design is meant to encourage cognitive processing of the misinformation and correction.

Experiment 3 is designed to test the hypothesis that belief echoes will only be created among those most likely to engage in the reasoning process that drives cognitive belief echoes.

H5: Belief echoes will only be created among those predisposed to dislike the candidate.

Experiment 3 Design

A total of 178 subjects were recruited through Amazon's Mechanical Turk and assigned to one of eight different groups. The design was a 4 (type of misinformation) x 2 (candidate party). The full design is shown in Figure 4.1.

Experiment 3: Design

Figure 4.1. Experiment 3: Design

		MISINFORMATION FORMAT			
		Corrected Misinfo 1 (felon donation)	Corrected Misinfo 2 (infant mortality)	Corrected Misinfo 3 (fraud)	Control (no misinfo)
CANDIDATE PARTY	Candidate is opposing party as subject	N=29	N=26	N=20	N=22
	Candidate is same party as subject	N=14	N=25	N=23	N=19

While the experiments in the first chapter looked only at one type of misinformation (accepting donations from a felon), Experiment 3 employs three different types of misinformation.

- Candidate accepted donations from a convicted felon.
- Candidate's health care policies doubled infant mortality.
- Candidate used taxpayer money to fund a personal vacation.

Testing multiple types of misinformation offers additional evidence that belief echoes are created by different types of false claims. I predict that the creation of belief echoes will follow similar patterns across the different types of misinformation.

Participants were told that they would be reading a series of headlines and photographs from the *Nevada Tribune's* ongoing coverage of a Congressional race in Nevada. Subjects were instructed to "Please just focus on the headlines and photographs, and **make sure to read the headlines carefully.**" An emphasis was placed on reading the headlines carefully in order to encourage deep processing and create an environment conducive to creating cognitive belief echoes. Before

proceeding to the headlines, they read brief biographical descriptions of each of the candidates. The party of each of the candidates was emphasized in the descriptions (see Appendix A for text of the candidate descriptions).

Each headline was presented on a separate page. Subjects in the corrected misinformation conditions read a total of six different headlines. Subjects in the control condition read four headlines. Each headline followed a similar visual format to the example shown below in Figure 4.2. Experiment 3: Example Article Format

Figure 4.2. Experiment 3: Example Article Format



All subjects read two headlines giving basic background information about the race (see Appendix A for exact wording). Then, subjects in the three “corrected misinformation” conditions read a headline containing a piece of misinformation, and then immediately afterwards, a headline containing a correction. The

misinformation and corrections took a similar format for all three conditions: the misinformation came in the form of an accusation from the opposing candidate, and the correction was a result of an investigation by the *Nevada Tribune*.

Table 4.1. Experiment 3: Misinformation condition shows the text for each of the misinformation conditions. The associated images are shown in Appendix A.

Table 4.1. Experiment 3: Misinformation condition headlines

Condition	Misinformation Headline	Correction Headline
Corrected Misinformation 1	Felon Donation Scandal Hits Mitchell Campaign <i>Ross accuses Mitchell of accepting campaign donations from convicted felon Joe Fenz</i>	Mitchell's Donation Records Show No Felon Donations <i>A Nevada Tribune investigation of Mitchell's finances found no record of donations from convicted felon Joe Fenz</i>
Corrected Misinformation 2	Past Policies Return to Haunt Mitchell Campaign <i>Ross says Mitchell's health care reforms increased the number of newborn deaths in Nevada</i>	Newborn Deaths in Nevada Did Not Increase Under Mitchell's Hospital Reforms <i>A Nevada Tribune investigation of hospital records shows no rise in infant mortality in Nevada hospitals since Mitchell's policies were implemented</i>
Corrected Misinformation 3	Corruption Scandal Hits Mitchell Campaign <i>Ross accuses Mitchell of spending \$294,000 of taxpayer money to finance three family vacations to the Florida coast</i>	Mitchell's Vacation not Paid for with Taxpayer Money <i>A Nevada Tribune investigation of Mitchell's finance reports shows he did not use public funds to pay for his vacations</i>

After reading the articles, all subjects completed a two-minute distractor task (see Appendix C). Then, they answered a series of questions assessing their attitudes towards the candidates (see Appendix B). Finally, a manipulation check ensured that the correction was fully accepted.

Experiment 3 Results

As with previous experiments, it is first necessary to ensure that the correction was successful. Indeed, for each of the misinformation conditions, the correction was accepted: there was no significant difference between the correction misinformation condition and the control condition. In addition, people were not

more likely to believe the misinformation if it was about a candidate from their own party: there was not a significant interaction between candidate party and corrected misinformation for any of the three conditions.

Figure 4.3. Experiment 3: Evaluations of candidate by same-party shows evaluations of the candidate among those in the same party condition. H5 suggests that belief echoes should not be created among these subjects. The Figure 4.3 shows that evaluations of the candidate are actually higher for those who read the corrected misinformation as compared to the control group.

Figure 4.3. Experiment 3: Evaluations of candidate by same-party condition





Figure 4.4. Experiment 3: Evaluations of candidate opposing-party shows evaluations by subjects in the opposing-party conditions. H5 predicts that belief echoes are more likely to occur among these subjects, for whom the misinformation

concerned a candidate from the opposing party and the accusation came from a candidate of their own party. In each condition, evaluations were lower in the “corrected misinformation” condition than in the control condition.

Figure 4.4. Experiment 3: Evaluations of candidate opposing-party condition



Taken together, Figure 4.3 and Figure 4.4 present visual evidence that H5 is supported: subjects who read a corrected false claim about a candidate from an opposing party were more likely to exhibit belief echoes. The statistical test of H5 is the interaction between the corrected misinformation condition (corrected vs. control) and candidate party (same vs. opposing). An analysis of variance shows a significant interaction for the Misinformation 1 condition ($F=10.24, p < .05$) and

Misinformation 3 condition ($F=3.90, p < .05$). In the Misinformation 2 condition, the interaction is not quite significant ($F=2.90, p = .093$).

Experiment 4: Distinguishing between affective and cognitive belief echoes

This experiment was designed to elicit both cognitive and affective belief echoes by reproducing Experiment 2 in Chapter 3, with one key difference: a “recall task” question that asked subjects to write down everything they remembered about both candidates. Subjects were told that they would be given a bonus if they wrote down more items than the average survey respondent.²⁰ Half the subjects completed the recall task immediately before evaluating the candidate. The other half completed the recall task immediately after.

The goal of the recall task was to prompt subjects to think more deeply about the false claim and correction. The expectation is that this should create the ideal circumstances for the creation of cognitive belief echoes. Subjects who are predisposed to think well of the accused candidate (those who are of the same party), should see their belief echoes minimized, as they consciously correct for the effects of the negative information. In contrast, subjects who are predisposed to think poorly of the candidate—those of the opposing party—should be prompted to engage in the reasoning process that underlies cognitive belief echoes.

The first hypothesis suggests that among the subjects who took the recall task after evaluating the candidate, *all* subjects in the corrected misinformation

²⁰ Actually, all subjects were given the bonus, regardless of how many items they wrote down.

condition should exhibit evidence of belief echoes, regardless of party. As in Experiment 2, exposure to the misinformation and correction embedded in the text of the story should create affective belief echoes.

H6. In the post-attitude recall task condition, subjects in the same-party and opposing-party conditions should both exhibit belief echoes.

The second hypothesis predicts that among the subjects who completed the recall task *before* evaluating the candidate, only those predisposed to dislike the candidate should exhibit evidence of belief echoes.

H7. In the pre-attitude recall task condition, partisanship should predict the creation of belief echoes.

Finally, at the end of the survey, all subjects in the corrected misinformation condition were asked whether reading the corrected misinformation affected their opinion of John McKenna. While their answers cannot provide any causal leverage, a research question asks whether the open-ended answers exhibit evidence for the chain of reasoning that underlies cognitive belief echoes.

RQ1. Do subjects' responses to the open-ended question exhibit evidence of cognitive belief echoes?

Experiment 4 Design

309 people were recruited through Amazon's Mechanical Turk service for a 2 (misinformation) x 2 (recall task) x 2 (candidate party). This experiment followed a similar design to Experiment 2 in Chapter 3 of this dissertation. Subjects answered demographic questions, then read a news article ostensibly from the *Iowa Ledger*. Subjects were told that the article described a Congressional campaign that had recently taken place in Iowa.²¹ After reading the article, all subjects completed a short distractor task. Subjects in the pre-attitude recall task condition then completed the recall task.

Next, all participants completed the attitude measures, and then the post-attitude recall task group completed the recall task. Finally, all participants completed the manipulation check. At the end of the survey, subjects in the "corrected misinformation" group also answered an open-ended question that read as follows:

You read that John McKenna's opponent accused McKenna of accepting campaign donations from a convicted felon. You also read that an investigation by the Iowa Ledger found no evidence of any donations to McKenna from this felon. How did hearing about this accusation and the findings from the investigation change your opinion of John McKenna?

The open-ended question was included in order to gather evidence for cognitive belief echoes. If cognitive belief echoes are created through a chain of thinking in

²¹ Note that this was the only experiment in this dissertation conducted after the November 2012 election. As such, the article was presented as having been published earlier rather than describing an "ongoing" race. The questions in the electability battery were also changed slightly (for example, "who do you think won the election" rather than "who do you think will win the election")


which readers recognize the misinformation is false yet also use it as a basis for changing their evaluations, then this logic should appear in their responses.

Experiment 4 Results

The manipulation check ensures that the correction was successful by asking all subjects to rate the veracity of the discredited claim. There is no significant difference in belief in the misinformation between the control group and the treatment group, suggesting the correction was successful. There is also not a significant interaction between treatment and party, suggesting that acceptance of the correction did not vary by party.

H6 predicts that subjects in the post-attitude recall task condition will exhibit belief echoes regardless of party. Figure 4.5 shows the evaluation of the candidate for those who took the attitude elicitation *after* evaluating the candidate.

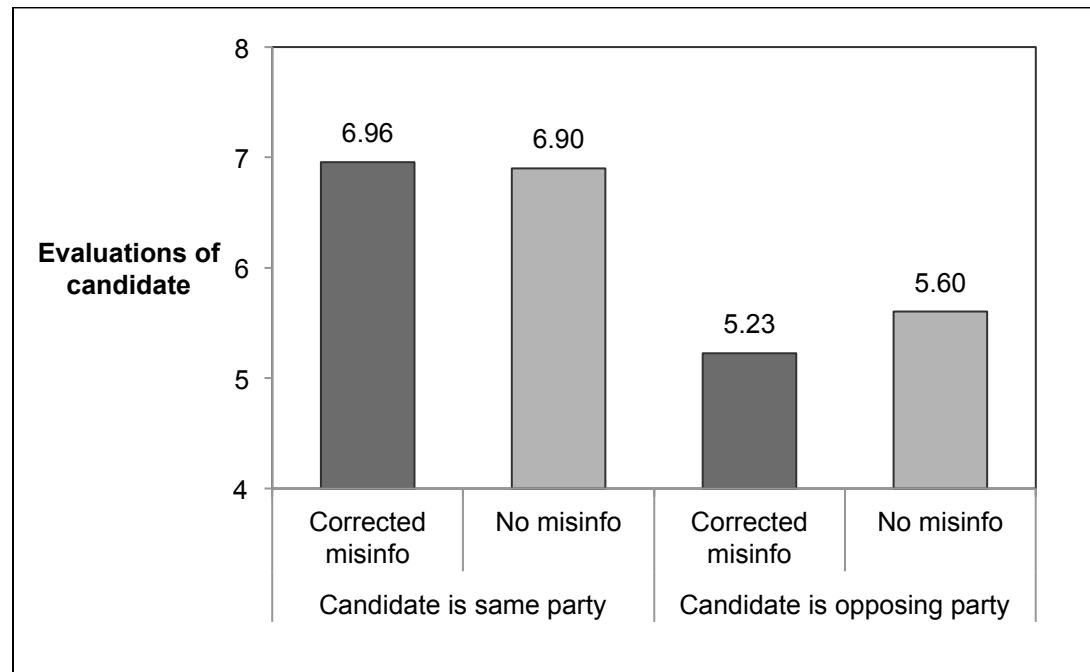
Figure 4.5. Experiment 4: Belief echoes in post-attitude recall task group



The results are in the expected direction: for both groups, evaluations are lower in the corrected misinformation condition than in the no misinformation condition. However, an analysis of variance shows that the difference does not reach significance ($F=1.43, p = .23$). Thus, H6 is only weakly supported.

H7 predicts that completing the recall task prior to the attitude elicitation will eliminate belief echoes for subjects in the same-party condition, but not for those in the opposing-party condition. Figure 4.6 shows that again, the results are in the expected direction: those in the same-party condition do not exhibit belief echoes, while those in the opposing-party condition evaluate the candidate more negatively than the control group. However, the three-way interaction between misinformation condition, recall task, and correction is not significant, suggesting that H7 is only weakly supported.

Figure 4.6. Experiment 4: Belief echoes in pre-attitude recall task group



Subjects' answers to the open-ended questions, while not a formal test of the hypothesis, can lend some insight into whether people engage in the chain of reasoning that drives cognitive belief echoes. Respondents were asked if reading about the accusation and the *Iowa Ledger* investigation changed their opinion of John McKenna. In the "same party" condition, 59 out of 80 respondents answered the question. In the "opposing party" condition, 60 out of 74 respondents offered an answer. I coded the answers for the five different binary categories described below:

- Lowered opinion of McKenna: response mentioned that the misinformation make them think more negatively of McKenna

- No change in opinion about McKenna: response mentioned that their opinion of him did not change
- Raised opinion of McKenna: response mentioned that the accusation made them more positive towards McKenna
- Lowered opinion of the accuser, Hall: response mentioned it made them think more negatively of Hall
- Believed that the misinformation might still be true: response mentioned being suspicious that the misinformation was actually true

Answers could fall into multiple categories. For instance, one subject wrote “It did not change my opinion of McKenna, which was lukewarm, but it made me think very poorly of Hall. It made me associate all of the slimy conniving political rhetoric with his character attack.” This respondent was coded as “no change in opinion about McKenna” as well as “lowered opinion of Hall.” Table 4.2 shows the distribution of responses for the same party condition and opposing party condition.²²

Table 4.2. Experiment 4: Open-ended response distribution

Category	Same party condition	Opposing party condition
Lowered opinion of McKenna	6.8% (4)	36.7% (22)
No change in opinion about McKenna	57.6% (34)	38.3% (23)
Raised opinion of McKenna	10.2% (6)	10.0% (6)

²² Note that the pre and post attitude recall task participants both answered this question, and their responses were combined. Both should be equally likely to engage in the reasoning process that creates cognitive belief echoes, since the open-ended question itself likely invokes this chain of thinking.

Lowered opinion of Hall (the accuser)	18.6% (11)	8.3% (5)
Still believed misinformation	1.7% (1)	3.3% (2)

Across both conditions, only three people in total mentioned that they still believed the misinformation might be true. This finding offers additional conformation of that the correction was successful. If cognitive belief echoes are indeed magnified by partisanship, more people in the opposing party condition should mention that the misinformation lowered their opinion of McKenna. And indeed, 36.7% of people in the opposing party condition said hearing the misinformation and accusation lowered their opinion, compared to just 6.8% of people in the same party condition.

The content of the responses also lent support to the cognitive belief echo mechanism, showing a chain of reasoning suggesting that the accusation, even though it is false, could imply that the accused candidate is untrustworthy. For example, one respondent wrote “It made me more suspicious of him - he might be covering something up.” Another mentioned that “It paints a bad picture of McKenna from the start, even if the allegations aren't true sadly,” and a third suggested that the accusation “made me think something shady might be going on in the McKenna campaign.”

Implications of affective and cognitive belief echoes

These experiments present evidence that belief echoes can be created through either an affective or cognitive path. Affective belief echoes can best be understood as a failure of a correction to fully compensate for the attitudinal effects of a piece of misinformation. There are several real-world factors that might amplify affective belief echoes. First, affective belief echoes are more likely when citizens process information under conditions of “cognitive busyness” (Gilbert 1989). Consumption of political media occurs under exactly these conditions: a citizen might read the newspaper on a smart phone while riding the bus, listen to NPR while cooking dinner, or see a news item scroll across the bottom of their television screen while watching a favorite show. Consumers may also seek out multiple sources of news simultaneously: *Newsweek’s* Howard Kurtz called the 2012 Presidential debates a “watershed moment of multiple screen use by viewers,” referring to people who watched the debate while also participating in conversations in social media.

Affective belief echoes may also be magnified when the misinformation carries an especially strong affective charge, and thus has a larger initial impact on evaluations. In these cases, the correction will likely be even less successful in returning the attitude to its pre-misinformation state. It is often the most vivid false claims that tend to be repeated both in the news (Berinsky 2012) and in interpersonal conversation (Weeks & Southwell 2010): fertile ground for generating affective belief echoes.

Cognitive belief echoes, on the other hand, are more likely when people are paying close attention. The answers to the open-ended questions demonstrated that many people remained wary of the candidate even though they recognized that the accusation was false, and this tendency was exacerbated when the candidate was of the opposing party. One potential way of mitigating this effect might be to provide readers with an alternative narrative explaining why the false accusation was made. Was the opponent behind in the polls? Was it an intern's mistake? Misinformation emerges for a reason, and clarifying that reason may reduce readers' tendency to "blame the victim."

Understanding how belief echoes are created is a necessary step in minimizing them. While there is no single approach that will successfully eliminate affective and cognitive belief echoes, the results of this chapter suggest a number of potentially promising strategies. Chapter 6 of this dissertation discusses in more detail the contextual and individual-level attributes that may encourage (or minimize) the formation of affective and cognitive belief echoes.

Chapter 5 – Changing incentives for politicians and journalists

Daniel Patrick Moynihan is quoted as remarking that “everyone is entitled to his own opinion, but not to his own facts.” It often seems, however, that Moynihan’s observation bears no relation to the modern political environment. Political debate is characterized not only by conflicts over what *should* be true, but also over what *is* true. Politicians argue over who benefited from a particular policy, over their opponents’ records, and even over basic descriptive statistics about the American people. As Kuklinski et al (1998, p 48) point out “[v]ery few factual claims are beyond challenge; if a fact is worth thinking about in making a policy choice, it is probably worth disputing. Rival advocates compete to define the facts, control their presentation, and determine their relevance.”

The existence of belief echoes raises the stakes of this competition, drastically increasing the incentives for politicians to spread false claims. If misinformation can shape attitudes even if it is successfully corrected, then why not lob false accusations with abandon? Even if these claims are immediately discredited, they can still inflict significant electoral harm on the opposition. One possible constraint on politicians’ behavior, however, may be the effect that making false claims has on their own reputations. To what extent do voters punish politicians for making false claims? The intuition of many journalists is that calling out politicians for lying has little effect. The *New York Times*’ David Carr commented that in the 2012 election, “both campaigns seemed to live a life beyond consequence,

correctly discerning that it was worth getting a scolding from the journalistic church ladies if a stretch or an elide or an outright prevarication did damage to the opposition” (Carr 2012). This chapter tests the extent to which Carr’s intuition is true, examining the extent to which politicians are punished for making false claims and if so, whether this effect varies by party.

The existence of belief echoes may also pose a *disincentive* for journalists to engage in active fact-checking. Already, many media outlets are wary of actively adjudicating between factual claims, driven by a concern that doing so will cause them to be perceived as biased (Graves 2012a). The existence of belief echoes suggests that even if journalists do engage in this potentially risky endeavor, misinformation’s effects cannot be entirely eliminated. However, belief persistence and belief echoes are not the only metric by which to measure the success of fact-checking efforts. Fact-checking may also serve a greater good by increasing trust in media more generally. If so, this could serve as an incentive for journalists to engage the often unrewarding task of fact-checking, a task rendered even more Sisyphean by the existence of belief echoes. This chapter explores the incentives facing journalists in regards to fact-checking and introduces new factors into an already contested question complicated by the existence of belief echoes.

Three major approaches to fact-checking

This chapter is focused on outlining the practical factors that might constrain (or encourage) belief echoes. Specifically, it asks how the existence of belief echoes might affect the cost-benefit calculations made by politicians in deciding whether to

issue false claims, as well as those made by the media when deciding whether to correct those claims. As such, the experiments are designed to reflect the real-world environment in which misinformation is corrected, including the choices facing media outlets about how to correct misinformation.

As political campaigns have entered what some journalists have taken to calling a “post-truth” age characterized by an indifference to factual accuracy (Fallows 2012, Rainey 2012), an active normative debate that has emerged among journalists over how factual disputes should be addressed by the media (Sullivan 2012; Greenhouse 2012). Three competing perspectives have come to dominate this debate. The first approach views journalists as arbiters, holding that reporters themselves should adjudicate between factual claims, and carefully state (if possible within the article) when a claim is false. The second outsources this job to other institutions either within the media outlet (like the Washington Post’s “Fact Checker” column) or outside of it, like the independent FactCheck.org or Politifact.com. This is also often the de facto approach of television news, which frequently cites fact-checking organizations when reporting claims. The third, colloquially known as the “he-said, she-said” approach, does not require that journalists adjudicate between facts. Instead, this approach holds that journalists should simply report on each side’s remarks and leave it to readers to determine which side is more convincing.²³

For example, on August 29, 2012, Paul Ryan spoke at the Republican National

²³ Although this description is to some extent an over-simplification of the nuanced positions of journalists and scholars (for example, see Rosen 2010 and Graves & Glaisyer 2012), it provides a useful structure for formulating empirically testable hypotheses.

Convention. His speech contained several statements that were misleading or, in several circumstances, factually inaccurate (Cooper 2012). Most notably, Ryan claimed that a GM plant in his hometown of Janesville, Wisconsin, was closed under Obama's administration. The plant actually closed while George W. Bush was in office. The media's various responses to Ryan's claims illustrate the three approaches outlined above and employed in this experiment: journalistic adjudication, outsourcing to fact-checkers, and the "he-said, she-said" approach. Some media outlets, like the *Associated Press's* article headlined "Ryan takes factual shortcuts in speech" corrected Ryan's false claims outright (Woodward & Gillum 2012). Others, like the *USA Today*, partnered with fact-checking organizations like FactCheck.org to debunk Ryan's statements (Farley et al 2012). An article in the *Washington Examiner* took the third approach, simply repeating an Obama spokesperson's response that "if Paul Ryan was Pinocchio his nose would be back in Janesville right now" (Gehrke 2012).

The experiment presented in this chapter compares how each of these three approaches might alter cost-benefit analyses undertaken by both politicians and media outlets in the "post-truth" world. Specifically, it addresses two major concerns. First, that the existence of belief echoes might induce politicians to spread false information. Second, that the existence of belief echoes might, by decreasing the benefits of journalists for correcting false information, offer further encouragement for them to fall back on the simpler and potentially safer "he said, she said" approach to correcting misinformation.

Experiment 5: Constraints on politicians and media

Are politicians punished for making false claims?

Are politicians who make false claims are punished by voters for attempting to mislead the public? This is a normatively important question—if politicians do not suffer any consequences from making false claims, then they have an enormous strategic incentive to continue doing so.

Answering this question requires distinguishing between two different processes. First, there is a potential effect of lobbing an accusation at an opponent. Existing work on this topic falls under the broad category of “negative campaigning,” a topic on which research is voluminous but largely inconclusive (Lau et al 1999). A major reason for the mixed findings of negativity effects may be that much of the research employs general measures that conflate multiple forms of negativity, including (but not limited to) incivility, personal attacks, and advertising strategies that turn the focus to the opposition’s weak points rather than to one’s own strong points. Although citizens tend to report disliking negativity in politics (Pew 2012b), it may have some positive effects. For example, Finkel and Geer (1998) find that negative information can increase political knowledge and stimulate turnout.

Learning that a candidate accused her opponent of wrongdoing may have a different effect than learning that she made a *false* accusation. Intuitively, we might expect that a politician who makes a false claim will be punished. Empirical research supports this intuition: perceptions of integrity contribute significantly to overall candidate evaluations (Funk 1999). However, studies of political corruption

suggests that politicians are not always punished for unethical behavior. Whether a politician suffers electoral consequences for her malfeasance is highly dependent on partisanship: perceptions of corruption affect evaluations less if the politician is of a person's own party (Anderson & Tverdova 2003). *Time* magazine's Michael Scherer (2012) expresses a similar thought:

The vast majority of the American voting public long ago demonstrated their willingness to simultaneously forgive fibs told by their own team and express umbrage at the deception offered by the other team.

This experiment tests two hypotheses about how politicians are punished. The first seeks to confirm previous research showing that perceptions of integrity matter to overall evaluations, and predicts that politicians' evaluations will be lowered when their accusations are shown to be false.

H8: Politicians will be punished for making an accusation when the claim is shown to be false.

The second predicts an interaction between evaluations of the accuser and the party of the subject. Specifically, it predicts an effect similar to that found in studies of corruption: that politicians will suffer more serious consequences for lying from non-supporters than from supporters.

H9: The negative effect of making false claims will be stronger if the person making the false claim is of the opposing party.

What costs and rewards does fact-checking impose on the media?

The existence of belief echoes suggests that even when the media actively corrects false claims, misinformation can still affect their readers' attitudes. This may serve as a disincentive for journalists to actively adjudicate between competing factual claims. Second, engaging in fact-checking is not cost-free for journalists; investigating false claims takes time and resources. The existence of belief echoes introduces another wrinkle by suggesting that misinformation can continue to affect attitudes despite journalists' efforts to conform to the best practices of corrections. This is potentially a serious disincentive for journalists already disillusioned with the levels of belief persistence present among citizens: the *New York Times'* David Carr summed up the 2012 fact-checking efforts with, "why bother?" while *Atlantic's* James Fallows lamented that "when the press calls a lie a lie, nobody cares" (Fallows 2012). Finally, active fact-checking may threaten the public's belief in the media's objectivity, especially if a correction favors one party over another (Graves 2012b). This concern often leads journalists to rely on the previously discussed "he said, she said" format which refuses to rank one version of the truth above another.

In addition, correcting a statement from a given political party may alienate readers from the opposing party. Individuals tend to rate information sources as more credible when they provide information that is consistent with their prior beliefs. This phenomenon is generally referred to as "biased assimilation." In a canonic experiment by Lord, Ross, and Lepper (1979) subjects who opposed or supported the death penalty rated studies disconfirming their beliefs as less

credible than one that confirmed their beliefs. Gentzkow and Shapiro (2006) show that a similar pattern occurs when people evaluate media outlets: consumers perceive that a media source is of higher quality when it produces content that reinforces their pre-existing beliefs. These findings suggest that journalists' intuition that correcting misinformation may increase perceptions of bias among some of their readers may be well-founded.

H10. When a correction issued by the newspaper runs counter to a person's political predispositions, the newspaper is perceived as more biased than when the correction is in the "he said, she said" format.

Public perceptions of a biased news media have increased over the past decade (Ladd 2012). One of the factors contributing to this increase may be the ongoing debate over fact-checking, as it makes the general topic of bias more salient for the public (Watts et al 1999). Given the strong observational and experimental evidence for a "hostile media effect" in which partisans see the media as biased against their own interests (Vallone, Ross, & Lepper 1985; Gunther & Schmitt 2004), corrections may exacerbate this effect.²⁴

On the other hand, public demand for corrections is high. The seeming proliferation of political misinformation over the last five years has intensified public calls for journalists to adjudicate between competing factual claims. The public's frustration with "he said, she said" reporting is exemplified by the response of *New York Times* readers to a column published by public editor Arthur Brisbane

²⁴ This same concern is also relevant for fact-checking organizations, who struggle to maintain credibility in the face of criticism from both sides (see Pareene 2011)

in January 2012. Brisbane wrote that he was “looking for reader input on whether and when *New York Times* news reporters should challenge ‘facts’ that are asserted by newsmakers they write about” (Brisbane 2012). Public response, both on the *Times* website and elsewhere, was immediate and incensed. Graves (2012a) found that of the 265 comments made in the three hours before the *Times* shut down commenting, “exactly two (discounting obvious sarcasm) disagreed with the proposition that reporters should challenge suspect claims made by politicians.” Margaret Sullivan, the public editor of *The New York Times*, described the growing pushback from readers:

Simply put, false balance is the journalistic practice of giving equal weight to both sides of a story, regardless of an established truth on one side. And many people are fed up with it. They don’t want to hear lies or half-truths given credence on one side, and shot down on the other. They want some real answers (Sullivan 2012).

Given the disconnect between what people say that they want and what they actually want, perhaps even more significant is the fact that fact-checking features are increasingly popular in mainstream media outlets, garnering measureable pageviews (Wimple 2012).

Most literature on media trust focuses on institutional factors such as partisan media (Ladd 2012) or individual factors, like ideology and social trust (Lee 2010). However, given the public demand for journalists to take a more active role, it is plausible that even if a correction decreases trust in a particular outlet, it might increase trust in media more generally by providing evidence that the news media is

capable of adjudicating between competing factual claims. D'Angelo and Lombard (2008, page 6) use the phrase “accountability frame” to describe news articles that “shift the emphasis of a news story away from news management and toward news norms.” In other words, accountability frames occur when media outlets are forthright about their efforts to discern the truth. In an experiment, accountability frames in news reporters lead readers to rate the press as having a higher news quality. If fact-checking engenders a similar effect, this would provide a much-needed incentive for journalists to continue to engage in fact-checking despite the existence of belief echoes.

H11. When journalists actively arbitrate between factual claims, it will raise evaluations of media more generally

Experiment 5 Design

A total of 606 people were recruited via Amazon’s Mechanical Turk. The experiment was a 5 (misinformation format) x 2 (candidate party) between-subjects design. Participants were assigned to one of ten conditions. The full design is presented in Figure 5.1.²⁵

²⁵ Note that some of the data used in this experiment was also used in Experiment 2, Chapter 3.

Figure 5.1 Experimental Design

		MISINFORMATION FORMAT				
PARTY OF CANDIDATE		Newspaper correction	Fact-check correction	Campaign correction	Uncorrected misinfo	No misinfo
	Opposing party as subject	N=53	N=58	N=61	N=61	N=66
	Same party as subject	N=61	M=52	N=60	N=60	N=63

Participants first answered a short series of demographic questions, including age, education, and political interest. A full description of the sample is in Appendix D. Subjects who did not immediately identify with a party answered a series of branching questions until they could be sorted into either leaning Democratic or leaning Republican.

Next, all participants read a news article, ostensibly from the *Iowa Ledger*, about an Iowa Congressional race. The article (full text in Appendix A) provided background information about the race and focused on one of the candidates, John McKenna. For half the participants, John McKenna was identified as being the same party as them. For the other half, John McKenna was identified as being of the opposing party. The article read by the three groups was identical except for one paragraph. The version read by the Uncorrected Misinformation and Corrected

Misinformation groups contained a paragraph describing an accusation made by Eric Hall, McKenna's opponent:

The campaign has heated up in recent weeks. On Sunday, Hall accused McKenna of accepting campaign donations from a convicted felon named Daniel Elsio. Elsio, who ran the largest drug ring in Iowa while McKenna was mayor, was convicted of first-degree murder in 2010. According to Hall, McKenna accepted over \$10,000 from Elsio. Hall commented that "McKenna's corrupt behavior shows that he and other [Democrats/Republicans] can't be trusted to do what's best for Iowa citizens."

In the correction conditions, this paragraph was followed immediately by a correction:

However, further investigation of the campaign donation records by [journalists at the *Iowa Ledger*/the independent fact-checking organization GetTheFacts.org] has shown no record of any donation from Elsio to McKenna's campaign. Campaigns are required to disclose the names of all individuals who contribute \$200 or more in an election cycle, and [the *Ledger*/GetTheFacts.org] did not find Elsio's name listed.

The correction issued from the campaign followed a slightly different format:

However, McKenna's campaign has released a statement stating that there is no record of any donation from Elsie to their campaign. Campaigns are required to disclose the names of all individuals who contribute \$200 or more in an election cycle, and McKenna's campaign stated that Elsie's name is not listed.

After reading the two articles, all participants completed a short (two-minute) distractor task in which they looked for differences between two photos (details in Appendix C). Next, subjects answered a series of questions measuring their attitudes towards the candidates (question wording in Appendix B). Evaluations of John McKenna were measured with six questions that included a feeling thermometer and several trait assessments ($\alpha = .89$). Evaluations of Eric Hall were measured with five questions ($\alpha = .83$).

Subjects then completed a manipulation check in which they were asked a series of questions assessing their memory of material in the article: "You read a newspaper story about an ongoing congressional race. Knowing what you know now, please tell us which of these statements are true." They were presented with five statements about the article. Two dealt with the donations, one was factually incorrect, one was correct, and one was an unknown (whether the two candidates were married). The two questions measuring belief in the misinformation were highly correlated ($p < .001$), and were combined into a single measure.

Finally, participants answered questions about four different news outlets: the *USA Today*, CNN, GetTheFacts.org, and the *Iowa Ledger*. For each outlet, they indicated their overall opinion about the outlet via a feeling thermometer; the extent to which they thought the outlet could be trusted to get the facts right; and whether they believed it favors the Democrats or Republicans. For each outlet, the first two measures were significantly correlated ($p < .001$) and so were combined to form an overall evaluation measure. Question details are available in Appendix B.


Experiment 5 Results

When are politicians punished?

Figure 5.2 shows evaluations of Eric Hall, the candidate who accused John McKenna of accepting campaign contributions from a convicted felon. For this analysis, the “fact-checking correction” and “newspaper correction” conditions were combined into a single “successful correction” condition, since a manipulation check showed that both were successful at correcting the misinformation.²⁶ The “unsuccessful correction” condition is the “he said, she said” format, which unlike fact-checking and newspaper corrections did not eliminate belief persistence.

²⁶ In order for the correction to be deemed successful, it must lower belief in misinformation back to the levels of those in the “No Misinformation” condition. For the newspaper and fact-check conditions, this was the case. Those who read a correction issued by the newspaper or by a fact-checking organization were equally as skeptical of the misinformation as those who never saw it. However, the correction issued by the campaign was much less successful. Although it did significantly reduce belief compared to the Misinformation Only condition, it was also significantly higher than the “No Misinformation” condition ($F=21.8, p < .001$).

Figure 5.2 Experiment 5: Evaluations of accuser




H8 predicts that politicians will only be punished for making accusations if the claim is perceived as false. This hypothesis is supported. Eric Hall's evaluations are not lowered when he accuses his opponent of wrongdoing: there is no significant difference between the Uncorrected Misinformation and No Misinformation conditions. However, he does pay a significant cost if the claim is shown to be false ($F = 13.4, p < .001$). This cost is *not* incurred in the "he said, she said" condition, in which the correction is not successful. In this condition, even subjects of the opposing party do not punish him for making the claim. The "he said, she said" approach thus not only fails to correct the misinformation, but also fails in changing attitudes towards the accuser.

H9 predicts that the “punishment” for lying will be more severe if the candidate is of the opposing party. This hypothesis, however, is not supported: there is not a significant interaction between party and corrected misinformation. Politicians are punished equally for lying whether they are of a person’s own party or of the opposing party.

When are media rewarded?

H10 directly addresses journalists’ concern that when they explicitly arbitrate between competing claims, they are seen as more biased than when they employ the “he-said, she-said” approach. Figure 5.3 shows the perceived bias of the *Iowa Ledger* (the paper that published the article read by all participants). Participants were asked to indicate whether the *Iowa Ledger* favors one party over the other. Responses were recoded based on participants’ own party on a 0 – 4 scale such that a 0 indicated a belief that the outlet strongly favored one’s own party and a 4 indicated a belief that the outlet strongly favored the opposing party.

Figure 5.3 Experiment 5: Perceived bias of *Iowa Ledger*

The results suggest that journalists' intuitions are correct. When journalists arbitrate between factual claims, perception of bias is highly dependent on whose claim they falsify. When a newspaper declares that a Republican candidate made a false claim, Republicans see the paper as biased towards Democrats, and vice-versa. An analysis of variance shows that there is a significant interaction between party and corrected misinformation ($F=17.8, p < .001$). On the other hand, "correcting" false claims via the "he-said, she-said" approach does not have a similar effect on perceptions of bias. Indeed, it is not significantly different than simply not correcting the misinformation at all. Using the "he said, she said" formulation may be unsuccessful at correcting misinformation and generating negative consequences

for politicians who make false accusations, but it does allow the media to avoid being seen as choosing sides.

Over the past few years, citizens have been vocal in demanding that newspapers correct false claims made by politicians. H11 predicts that even if people perceive a particular outlet as more biased for adjudicating between factual claims, reading corrections will raise their estimations of the media more generally. Specifically, H11 predicts that compared to the “he-said, she-said” approach, when journalists explicitly arbitrate between false claims it will have a positive impact on evaluations of media.

Traditional media trust questions often follow a format similar to that employed by Gallup’s ongoing poll: “In general, how much trust and confidence do you have in the mass media – such as newspapers, TV, and radio – when it comes to reporting the news fully, accurately, and fairly?” (Ladd 2012). However, this measure may be problematic given the widely differing definitions that citizens may have of “mass media.” For example, some people may interpret the question as asking about the media the American public in general consumes, while others may interpret it as a question about their own media use (Mutz et al 2012). Such differing interpretations open the door for psychological processes like the third-person effect to bias individuals’ assessments in systematic ways that make the measure less meaningful.

Given this potential problem, the “mass media evaluation” measure was constructed from questions about two different news outlets not mentioned in the experiment: CNN and the *USA Today*. CNN and the *USA Today* were chosen because

they are familiar media outlets to most Americans. In addition, because neither is frequently cited as a partisan news outlet (in contrast to, for example, Fox News or *The New York Times*), there is less potential for ceiling or floor effects to hide shifts in attitudes. The evaluation measure is comprised of feeling thermometers on CNN and the *USA Today* as well as questions asking whether each outlet could be trusted to “get the facts right.” The questions formed a highly reliable measure ($\alpha = .85$). Question details are in Appendix B.

Figure 5.4 shows media evaluation by condition. When no correction is issued, media is rated at 6.14. When the correction is in the he-said/she-said format, it is 6.04, not a significant difference. However, subjects who read a correction issued by the *Iowa Ledger* evaluated the mass media more positively, at 6.49.²⁷ The difference is weakly significant ($p < .1$).

²⁷ This result does not appear to be an artifact of comparison with the *Iowa Ledger* (i.e. people rating the *Iowa Ledger* lower because of perceived bias and other media outlets relatively higher). The order in which the media question were asked was randomized, and whether the *Iowa Ledger* question appeared before or after the CNN and *USA Today* measures had no effect on the results.


Figure 5.4 Experiment 5: Media Evaluations

Figure 6 show the same-party and opposing-party conditions combined. However, it is also worth noting that there is not a significant interaction between journalistic arbitration and party, suggesting that the increase in mass media evaluations is not driven by partisans who read about their own candidate's exoneration and as a result have more positive attitudes towards mass media.

Preventing belief echoes: incentives and disincentives

The goal of this chapter is to better understand the changing incentives facing both politicians and the media given the existence of belief echoes. The results are both encouraging and discouraging. First, despite journalists' intuition

that false claims go unpunished, politicians do face serious consequences for lying, even from members of their own party. This finding is promising, suggesting that politicians hoping to reap electoral benefit from belief echoes cannot lob false accusations without repercussions. However, this is conditional on the media playing an active role: they are only punished when media does its part to actively discredit those accusations. When, on the other hand, the media falls back to its “he said, she said,” formulation, politicians *can* make false claims consequence-free.

Should we expect journalists to take on this difficult task? The results of this experiment shows that actively adjudicating between false claims imposes both costs and benefits. First, as previously mentioned, politicians are punished when shown to be false. This should incentive fact-checking in and of itself, especially given the existing doubts among journalists that politicians suffer consequences. In addition, exposure to active fact-checking by journalists raises evaluations of mass media, and this effect is not conditional on partisanship. Insofar as media outlets benefit from increased respect being paid to the media as a whole, this increases the incentives for them to engage in serious fact-checking.

On the negative side, partisans are more likely to believe newspaper is biased when its correction contradicts their own views. This finding is unfortunate because it adds yet another reason to the growing list of reasons why journalists might believe fact-checking is futile. Not only does it not succeed entirely in muting the attitudinal effects of misinformation, it also increases perceptions of bias.

Discussing the controversies around voter fraud in the 2012 election, Sam Sifton, an editor at *The New York Times*, placed himself firmly on the side of “he-said,

she-said” camp, saying “There’s a lot of reasonable disagreement on both sides. It’s not our job to litigate it in the paper. We need to state what each side says.” The evidence presented in this chapter shows that whether Sifton’s approach is the right one depends on the end goal. If the end goal is preserving the paper’s status as “unbiased,” Sifton’s approach is the right one. Adjudicating between sides can have negative consequences, increasing perceptions of bias. But is being perceived as unbiased the goal to which journalists should aspire? As Munoz-Torres (2012) points out, this goal has its own ethical problems: “the principle of ever presenting opposing views as equally valid amounts to stating implicitly that all opinions possess the same value as truth-claims.” The results presented in this chapter show that fact-checking can also have positive outcomes, both by imposing penalties on lying politicians and by raising trust in media more generally.

Finally, it is worth noting that the effects demonstrated in this experiment were the cause of relatively minor manipulations. The source of the correction was just a few words out of an entire article, as was the fact that the accusation came from Hall. In the real world, we would expect both of these facts to figure much more prominently in an article. When a candidate accuses the opposition of wrongdoing, it is generally a headline-worthy event, not necessarily an aside buried in a larger article. And when fact-checking occurs, it is often done much more explicitly, as in the post-debate articles discussed in Chapter 1. The fact that relatively small changes in how misinformation and corrections were presented had significant effects on readers’ attitudes suggests that the constraints on politicians are very real, as are the potential rewards (and pitfalls) for a conscientious media.

Chapter 6 - Conclusion

Most existing research on misinformation focuses on examining the circumstances under which citizens hold misperceptions about the political world. While this line of research is important, I argue that given the increasing ubiquity of corrections and the rise of fact-checking, it is critical to examine how exposure to misinformation affects attitudes not only when corrections fail, but also when corrections succeed. In a series of experiments, I examine how exposure to *successfully corrected* misinformation affects citizens' political attitudes and behavior. The assumption of many journalists and academics is that once misinformation is successfully corrected, it will cease to affect attitudes. I find that muting misinformation's effects is not so simple. Even after citizens cease to believe a piece of information, it creates lingering attitudinal effects that I call belief echoes.

Belief echoes pose a serious challenge to the marketplace of ideas. The marketplace of ideas assumes that once a piece of information is definitively shown to be false, it will exit the marketplace. Citizens will no longer draw on this information in making decisions or forming opinions about policies and candidates. However, I show that even after a piece of information is disproved, its attitudinal effects persist. The existence of belief echoes has important consequences for politicians, the media, and citizens themselves, and suggests that the media's intense focus on misinformation may unintentionally magnify its effects on public opinion.

In this conclusion, I review the limitations of this dissertation and then outline the major consequences of belief echoes, including their impact on vote choice and political participation. I also explore how belief echoes may be created in the context of policy misinformation and positive false claims. Finally, I offer a set of practical strategies for minimizing the effects of misinformation, drawing on the mechanisms outlined in Chapter 4 as well as the real-world constraints facing the media described in Chapter 5. The existence of belief echoes underscores a point made in the introduction: that to effectively minimize the effects of misinformation, it is critical to consider both the structural factors that shape the media's approach to correcting false claims as well as the individual-level factors that influence how citizens choose, consume, and process political information.

Limitations

In this dissertation, I create belief echoes in a controlled experimental setting. As with any experimental finding, it is worth asking whether and how real-world contexts might amplify or minimize these patterns.

To what extent did the setting of these experimental treatments reflect how citizens encounter misinformation and corrections in the real world? Participants took the surveys online and read the newspaper article on a computer screen. This environment closely mirrors how citizens process news: over half of Americans read news online, more than read it in hard copies of newspapers or hear it on the radio (Pew 2012).

The experiments in this dissertation include over 1,200 distinct participants. While Mechanical Turk subjects are more diverse than many other convenience samples, they are not completely representative of the American population. Tables D1 and D2 in the Appendix provide the full demographics of the sample. In general, they are more liberal and younger than the general population.²⁸ However, there is no theoretical reason why these demographic differences should affect the likelihood of finding evidence of belief echoes. In other words, I do not expect belief echoes to be more prevalent among younger or more liberal subjects.

Mechanical Turk subjects may also differ from the general population in their attentiveness. The payment structure of Mechanical Turk incentivizes subjects to read instructions carefully and closely, and Berinsky et al (2012) find that Mechanical Turk subjects perform better on memory tasks than subjects from other online samples such as Polimetrix and Survey Sampling International. However, I argue that the attentiveness of the sample may mean that my surveys underestimate the full effect of misinformation on attitudes. The more closely subjects read, the more likely they are to fully process the correction. In the real world, citizens process political information with less care, and thus are less likely to carefully read the correction of any given piece of misinformation.

These experiments employ several operationalizations of the outcome of interest (attitude towards the candidate), including thermometer ratings and trait assessments. Across the studies, these multiple measures form a highly reliable index, increasing my confidence that the attitudinal effects I find are robust. Many

²⁸ Future iterations of this work may employ a population-based sample, which would also allow me to explore potential subgroup effects.

social science survey experiments employ not only attitudinal measures, but also behavioral measures that examine the extent to which attitudes (for instance, candidate affect) translate into actions (for instance, voting). However, the fictional nature of the scenarios makes it impossible to employ behavioral measures. For example, I cannot measure the voter turnout or campaign donations of participants who were exposed to the corrected misinformation, because the race detailed in the scenario does not exist in the real world. In a following section (behavioral implications of belief echoes), I discuss how attitudinal measures might be predictive of real-world political behavior.

Because the treatments in these studies consisted of both a piece of misinformation *and* a correction, addressing the question of whether these treatments are generalizable means thinking carefully about both components. Most of the experiments in this dissertation employed the same misinformation treatment: that a candidate was convicted of accepting campaign donations from a convicted felon. However, this treatment was chosen after a pre-test indicated that it had a similar negative impact on attitudes as two other pieces of misinformation: cheating at law school and an arrest for soliciting a prostitute (Thorson 2011). In addition, Experiment 3 tested two additional pieces of misinformation and found that they had similar effects on attitudes. However, each of these misinformation treatments in the pretest and Experiment 3 concerned biographical information about a candidate. As Chapter 2 discusses, political coverage often also contains misinformation about policies, as well as statistical misinformation about (for example) the number of undocumented immigrants in the United States. In the

Policy misinformation section that follows, I discuss how belief echoes might be created by policy misinformation.

The second part of the treatment is the correction. In contrast to the misinformation, the format of the correction varies a great deal across experiments. In Experiment 1, it was appended to a second article. In Experiment 2, it appeared within the same article as the misinformation. In Experiment 5, the source of the correction was varied. Across this variety of correction format, I found evidence for belief echoes. In addition, in the real world corrections are often much less definitive. In these situations, (as, for instance, in the “campaign correction” condition of Experiment 5), we would expect to see a larger aggregate effect of misinformation on attitudes comprised of both belief echoes and belief persistence.

Belief echoes in the real world

The experiments in this dissertation create belief echoes through single treatments of corrected misinformation. This design may actually underestimate the magnitude of belief echoes. In politics, multiple pieces of misinformation often serve to reinforce a single idea. For example, many of the false claims about President Obama reinforce the idea that he is in some way un-American (Berinsky 2012). These include rumors that he refuses to put his hand over his heart while saying the Pledge of Allegiance; that he was born outside of the United States; and that he is Muslim. While any single piece of misinformation may create relatively small belief

echoes, the integration of that misinformation into a larger narrative may considerably amplify its effect.

Behavioral implications of belief echoes

The experiments in this dissertation assess candidate affect using several different measures, including favorability and trait assessments. Given that these outcome measures are attitudinal rather than behavioral, it is worth asking to what extent the changes in candidate affect engendered by belief echoes may have consequences for real-world behavior.

Fundamentally, information's effects on political attitudes matter in large part because attitudes shape the choices voters make on Election Day. However, voters' attitudes about politicians are multi-dimensional, and research is divided on what factors exert the most influence over voters' choices. Spatial models of voting predict that issues are the primary motivator of political decision-making. According to spatial models, citizens will vote for candidates whose issue positions are closest to their own (Downs 1957). If voters cast their ballots based purely on issue considerations, belief echoes should not affect their choices.

Evidence suggests, however, that other factors also play a role. Voters respond to traits like trustworthiness, friendliness, and general perceptions of competence (Popkin 1991). Analytically, it is often difficult to determine with confidence whether voters cast ballots according to issue preferences or other factors. In general, citizens tend to hold more favorable opinions towards candidates with whom they agree on the issues (Hayes 2005). In addition, a

person's favorable attitudes towards a candidate may lead her to conclude that they share similar policy preferences (Conover & Feldman 1982).

To address this problem of causal inference, Lenz (2012) uses panel data to examine how voters' opinions change when a candidate adopts a stance on a new issue. Measuring voters' own issue stances' before and after their preferred candidate publicly takes a position, he finds evidence that citizens often change their own issue positions to reflect the views of their preferred candidate. When a politician's *reputation* suffers, however, voters are more likely to defect. Lenz concludes that George Bush's characterization of Al Gore as a dissembler may have had more serious consequences than his stance on Social Security. These findings suggest that candidate evaluations can play a critical role in election outcomes apart from any issue considerations.

Belief echoes may also affect the likelihood that a person participates in the political process. Stable individual levels and changing election-level variables (e.g., the competitiveness of the election) explain a great deal of the variance in participation (Verba, Schlozman, & Brady 1995). A candidate's personal appeal (e.g., his perceived strength of character) may also affect whether his supporters decide to donate to his campaign; volunteer; or try to convince a neighbor to support him (Brody & Page 1973, Druckman & Miller 2004). To the extent that belief echoes affect the calculations that individuals make when assessing candidates, their impact may also influence important aspects of civic participation.

Aside from any impact on vote choice, belief echoes resulting from negative information²⁹ can decrease enthusiasm for a candidate among his own supporters, as well as increase negative affect towards him among the opposition. Enthusiasm is linked with several forms of political participation (Marcus 1988), suggesting that negative belief echoes could impact a candidate's ability to raise money and find volunteers.

Factors that might amplify or minimize the magnitude of belief echoes

While several factors may affect the size of belief echoes, there are two in particular – policy and positive misinformation – that are most relevant given the patterns of political misinformation in today's media environment. It is critical to reiterate that the factors that affect belief echoes are different than those that might affect belief persistence. Belief persistence occurs when citizens continue to believe a false claim even in the face of credible corrections. The experiments in this dissertation are specifically designed to eliminate belief persistence. In the political world, misinformation is rarely corrected as successfully as it is in these experiments. As such, the results almost certainly *underestimate* the overall effects of misinformation on attitudes. In non-experimental settings, people whose attitudes are affected by a given piece of misinformation include not only those who experience belief echoes, but also those who continue to believe the misinformation and those who never encountered the correction.

²⁹ See the following section for a discussion of positive misinformation.

Although the experiments in this dissertation focus on candidate misinformation, policy misinformation can also create belief echoes. The example of so-called “death panels” provides a useful illustration of how policy misinformation can emerge and spread. The claim that the Affordable Care Act would create death panels first emerged in a commentary on the Act by Betsey McCaughey, a fellow at the conservatively-aligned Manhattan Institute.

One of the most shocking things I found in this bill, and there were many, is on Page 425, where the Congress would make it mandatory—absolutely require—that every five years, people in Medicare have a required counseling session that will tell them how to end their life sooner, how to decline nutrition, how to decline being hydrated, how to go in to hospice care (Lawrence & Schafer 2012).

Within just a few days, both Politifact and Factcheck had dismissed her statement as false. The sessions were entirely voluntary and no evidence existed which suggested that counselors would coerce patients into declining treatments. Nonetheless, in the weeks following McCaughey statement, the *New York Post* and the *Wall Street Journal* each published op-eds repeating the claim (both in order to debunk it). Soon after, Sarah Palin coined the term “death panel,” and media coverage ultimately resulted in hundreds of mentions of “death panels” over the next year. However, the claim was also widely and quickly labeled as false by newspaper and television news stories (Lawrence & Schafer 2011).

To what extent might belief echoes be created by corrected policy misinformation? The answer depends partly on whether the policy in question is

new or old. Because opinions about new policies are weaker, any given piece of misinformation is likely to exert a larger relative effect on their attitudes. For example, much of the public was still largely uninformed about the Affordable Care Act when the death panel rumors emerged (Nyhan 2010), and so belief echoes engendered by the rumor may have had an outside influence on public opinion.

Negative policy misinformation (e.g., death panels) may also have an outside impact as compared to positive policy misinformation. When forming policy opinions, people tend to place an outside emphasis on avoiding potential losses (Cobb and Kuklinski 1997), and so both affective and cognitive belief echoes created by negative policy misinformation might discourage citizens from supporting a particular policy.

The experiments in this dissertation focus entirely on *negative* false claims. This was a conscious decision based on the fact that much of the concern over misinformation in politics regards claims that could potentially lower evaluations of a candidate or policy. Nonetheless, *positive* misinformation also warrants consideration.

In considering the potential consequences of positive misinformation, it is important to distinguish between positive claims that are made by a candidate about himself and positive claims that are made by a third party. Chapter 5 presents evidence that politicians are punished for making false claims. This effect should also apply in cases where a politician makes a false claim about his own record. For example, vice-presidential nominee Paul Ryan received a good deal of attention for his claim that he had run a marathon in under three hours – a claim that was

subsequently proven to be false (Fallows 2012). This example differs from the experimental manipulations in this dissertation not only in valence but also in structure. Ryan himself was responsible for the false claim, and the results of Experiment 5 suggests that this can provoke a serious backlash. Public reaction to Ryan's statement confirmed this effect. As Salon's Joan Walsh put it, "There's a character issue here: What kind of person casually lies about an achievement like that, with a swagger?" (Walsh 2012).

When third parties issue positive misinformation about a candidate, however, the correction of that misinformation will not necessarily implicate the candidate's character. For example, in April 2012, newspapers reported that Trenton mayor Cory Booker rushed into a burning building to save a neighbor's daughter (Scherer 2012). If this story had been proven false (for instance, if the rescuer was a man who simply looked like Booker), there is no reason to think that Booker suffer a backlash similar to Ryan's. The newspapers that reported the false incident might be criticized, but it would be difficult to blame Booker himself for the mistake. In this example, exposure to the misinformation might create both affective belief echoes (as people failed to update their evaluations properly when exposed to the correction) and cognitive belief echoes (as people reasoned that "if the reporter wrongly believed Booker saved the child, perhaps it was because that seems like the sort of thing that Booker would do").

While politicians might pad their own resumes with some frequency, positive false claims made by third parties are less common (hence the partially imagined example above). However, the emergence of independent-expenditure only

committees, or SuperPACs, could change that dynamic. These SuperPACs spend enormous amounts of money on campaign advertising. More importantly, they are not explicitly affiliated with a candidate, meaning that their advertisements can be framed as coming from a non-partisan organization. Strategically, these organizations are ideally situated to spread misinformation that could create positive belief echoes without the backlash that accompanied the Ryan marathon claim.³⁰

Practical lessons for the media

For better or worse, most of Americans' exposure to both misinformation and correction occurs through mainstream media coverage of politics. While bloggers and social media are growing in importance, most Americans still receive their political news through major news sources, even if that media is read on a computer screen rather than on newsprint (Pew 2012). I offer three clear lessons for journalists on how to minimize belief echoes.

Employ affirmative corrections rather than negations

Affective belief echoes are created when corrections are unsuccessful at compensating for the attitudinal damage done by the initial misinformation. This effect is heightened when corrections are phrased as negations. Because people have difficulty processing negations (Hasson & Glucksberg 2005), negations can

³⁰ An analysis of SuperPAC advertisements run in the 2012 election shows that about 14% are positive (Fowler 2012). The analysis provides no information about whether the claims made (if any) are factually accurate, but the existence of third-party positive ads shows that positive advertising is a strategy at least under consideration by SuperPACs.

actually reinforce the very concept that they are trying to negate. In contrast, corrective affirmations may be more successful at muting the attitudinal effects of the initial misinformation. For example, while a correction in a form of a negation might read “Nikki did not cheat on her taxes,” a corrective affirmation would read that “Nikki paid her taxes in full.” Taking the time to formulate corrective affirmations may reduce the likelihood that corrections will leave behind lingering belief echoes.

Avoid unnecessarily spreading misinformation, even if it is corrected

Media outlets should be cautious in repeating misinformation—even if they are doing so in order to correct it. The existence of belief echoes suggests that publicizing misinformation may inadvertently exacerbate its effects. This does not imply that media outlets should not stop correcting misinformation altogether. Misinformation that has already been widely spread—for instance, the false claims around Obama’s birth certificate—should also be widely corrected. However, a content-generating strategy that relies on strategically selecting especially vivid false claims made in advertisements or discussed on fact-checking websites may have serious negative consequences.

Avoiding unnecessarily spreading misinformation may be a difficult task for media outlets in constant need of new content, especially given the popularity of fact-checking segments. Justin Bank, a *Washington Post* employee charged with measuring web traffic, commented that fact-checking content was “extremely popular with readers:”

On average, it attracts more click overs from social media. It, quite often, will go viral and pull in a slew of ‘direct traffic’ as folks e-mail and instant message it to each other. And, when featured on the homepage, the Fact Checker signature draws eyeballs, and, with them, clicks (Wemple 2012)

One strategy for addressing false claims without directly repeating them (and thus creating belief echoes) may be to build stories around the larger question driving the false claim rather around than the isolated piece of misinformation. For example, instead of “fact-checking” Paul Ryan’s false statement that Obama was responsible for closing a plant in Janesville, coverage could address the larger issue of manufacturing growth.³¹

The existence of belief echoes also raises larger questions about the marketplace of ideas assumption that underlies the media’s approach to misinformation. Traditionally, news organizations have walked a careful line between striving for timeliness and accuracy. Being timely means potentially publishing dubious or unverified claims in the interest of getting information to the public quickly. Accuracy, on the other hand, may require careful pre-publication fact-checking, which can result in time delays that costs the media outlet both economically and reputationally. Often, media outlets’ solution to this dilemma was simply to publish first and correct later (Graves 2012).

Weighing these trade-offs becomes even more complicated in cases where a given news article is not asserting a fact, but simply relaying someone else’s factual

³¹ This fact-driven reporting (as opposed to misinformation-driven reporting) is the strategy of FaceTheFacts.org, an independent non-profit that undertakes original reporting on the facts that underlie real-world political debates.

assertion. For example, imagine a newspaper report on a candidate's stump speech that also repeats misleading claims made in that stump speech. In Chapter 2, I discuss the recent push for reporters to correct such falsehoods, preferably in the text of the article itself. This push is an important step forward from the "he said, she said" interpretation of objectivity long espoused by many journalists (Bennett 2012).

However, the existence of belief echoes mean that this step forward may not be enough. Currently, reporters who seek to prevent misinformation from affecting attitudes do so by actively correcting the misinformation, either in a sidebar or within the text of the article. Indeed, as I discuss in Chapter 5, this is the "gold standard" of careful reporting and an approach specifically designed to prevent misinformation from affecting attitudes. But the existence of belief echoes means that news outlets' obligation may be not just to *correct* misinformation but to avoid spreading it in the first place. If the marketplace of ideas worked as it should, this type of prior restraint would be unnecessary—it would not matter if a piece of misinformation entered the marketplace as long as it was corrected. But in a world of belief echoes, repeating false claims imposes serious costs, even if those claims are immediately corrected.

"Naming and shaming"

In Chapter 5, I show that politicians are punished for making false claims, even by members of their own party. Political scientist Brendan Nyhan has called for journalists to engage in "naming and shaming" politicians and pundits who make

false claims, with the goal of increasing the reputational cost of spreading misinformation (Nyhan 2009). My research offers empirical evidence that in some cases this strategy may indeed be effective—perhaps not at convincing politicians to stop the practice altogether, but at least at ensuring that their actions have electoral consequences.

Clarity about who is making false claims can also potentially reduce cognitive belief echoes. Including relevant information about an accuser’s motivation for spreading the false claim may encourage readers to construct an alternative chain of reasoning that blames the source of the misinformation (even if that source is simply a careless reporter) rather than its target.

Belief echoes: part of a larger misinformation and correction process

In contrast to most work on misinformation, I focus on what happens *after* misinformation is successfully corrected. The experiments are designed to be a hard test for the existence belief echoes: the misinformation itself is falsifiable, and it is corrected decisively and immediately. Because the experiments are specifically designed to produce the successful correction of misinformation, they necessarily underestimate the *total* attitudinal effects of misinformation. In the real world, misinformation’s effects are not limited to its capacity to create belief echoes. Some misinformation goes entirely uncorrected. If it is corrected, some citizens might never hear the correction, especially given the fragmentation of the modern media environment. Of those who were exposed to the citizens, many may continue to believe that the false claim is true, especially if the information itself is difficult to

falsify or due to partisan-driven motivated reasoning (Nyhan & Reifler 2010). And finally, this dissertation shows that even when citizens encounter, process, and accept the correction, they may not be immune to its attitudinal effects.

Despite the negative consequences of belief echoes, there also exists cause for optimism. First, knowing *how* belief echoes are created also provides the tools necessary to minimize their effects. While no single solution exists, efforts to avoid drawing unnecessary attention to corrected misinformation; offering corrective affirmations rather than negations; and increasing clarity around the reasons for the false claim may together succeed in greatly reducing the lingering effects of misinformation

When corrections *are* successful, politicians suffer consequences for making false claims. This finding provides a key counterweight to the incentive posed by belief echoes for politicians to make unfounded accusations. Finally, reading corrected misinformation does increase citizens' trust in media more generally. Insofar as there is an objective political reality to which citizens—regardless of party—should be able to refer in forming opinions, this reality will necessarily be experienced through media. The more that citizens are able to trust that they are being shown an accurate picture of the world, the more likely it is that they will be invested, functioning, democratic participants.

References

- Abramson, P. R., Aldrich, J. A., Paolino, P., & Rohde, D.W. (1992). Sophisticated voting in the 1988 presidential primaries." *American Political Science Review* 86(1):55-69.
- Anderson, C. A. (1982). Inoculation and counter-explanation: Debiassing techniques in the perseverance of social theories. *Social Cognition* 1(2): 126-139.
- Anderson, C. A., Lepper, M., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology* 39(6), 1037-1049.
- Anderson, C. A., & Lindsay, J. J. (1998). The development, perseverance, and change of naive theories. *Social Cognition*: 16: 8-30.
- Anderson, C. A., New, B. L., & Speer, J. R. (1985). Argument availability as a mediator of social theory perseverance. *Social Cognition*, 3(3), 235-249.
- Anderson, C. J., & Tverdova, Y. V. (2003). Corruption, political allegiances, and attitudes toward government in contemporary democracies. *American Journal of Political Science* 47(1): 91-109.
- Ballasy, N. (2012). Report: DOJ may indict SC Gov. Nikki Haley for tax fraud. *The Daily Caller*. 29 March 2012. Retrieved December 6, 2012, from <http://dailycaller.com/2012/03/29/report-doj-may-indict-sc-gov-nikki-haley-for-tax-fraud/>
- Bartels, L. M. (2002). Beyond the running tally: Partisan bias in political perceptions. *Political Behavior*, 24(2), 117-150.

- Bennet, J. (2012). "We're not going to let our campaign be dictated by fact-checkers." *The Atlantic*, August 28. Retrieved August 29 from <http://www.theatlantic.com/politics/archive/2012/08/were-not-going-to-let-our-campaign-be-dictated-by-fact-checkers/261674/>.
- Bennett, W. L. (2005). *News: The politics of illusion*. New York: Pearson/Longman.
- Berinsky, A. J. (2012). Rumors, truths, and reality: A study of political misinformation. Presented at the West Coast Experiments Conference, Berkeley, California. Retrieved from <http://web.mit.edu/berinsky/www/files/rumor.pdf>.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, 20(2), 351–368.
- Box-Steffensmeier, J. M., & Schier, S. E. (2009). *The American elections of 2008*. Rowman & Littlefield.
- Brisbane, A. S. (2012). "Should the Times be a truth vigilante?" *New York Times*. 12 January 2012. Retrieved December 6, 2012, from <http://publiceditor.blogs.nytimes.com/2012/01/12/should-the-times-be-a-truth-vigilante/>
- Brody, R. A., & Page, B. I. (1973). Indifference, alientation and rational decisions. *Public Choice*, 15(1), 1-17.
- Bullock, J. (2006). Partisanship and the enduring effects of false political information. Presented at the 2006 Midwest Political Science Association Annual Conference, Chicago, Illinois.

- Burdein, I., Lodge, M., & Taber, C. (2006). Experiments on the automaticity of political beliefs and attitudes. *Political Psychology*, 27(3), 359-371.
- Carr, D. (2012). "A Last Fact Check: It Didn't Work." *New York Times*, 6 November 2012. Retrieved December 6, 2012, from <http://mediadecoder.blogs.nytimes.com/2012/11/06/a-last-fact-check-it-didnt-work>
- Carr, S. (2012). How to worry about a clicks-driven Times-Picayune. *Columbia Journalism Review*. Retrieved December 6, 2012, from http://www.cjr.org/behind_the_news/how_to_worry_about_a_clicks-dr.php
- Cobb, M. D., Nyhan, B., & Reifler, J. (2012) Beliefs don't always persevere: How political figures are punished when positive information about them is discredited. Forthcoming, *Political Psychology*. Retrieved December 6, 2012, from <http://www.dartmouth.edu/~nyhan/positive-misinformation.pdf>.
- Conover, P. J., & Feldman, S. (1982). Projection and the perception of candidates' issue positions. *The Western Political Quarterly*, 228-244.
- Cooper, Michael. (2012). "Fact-checking Ryan's convention speech." *The New York Times*. 20 August 2012. Retrieved 6 October 2012, from <http://thecaucus.blogs.nytimes.com/2012/08/30/in-ryan-critique-of-obama-omissions-help-make-the-case/>.
- Coronel, J. C., Duff, M. C., Warren, D. E., Federmeier, K. D., Gonsalves, B. D., Tranel, D., & Cohen, N. J. (2012). Remembering and voting: Theory and evidence from amnesic patients. *American Journal of Political Science* 56(4): 837-848.

- D'Angelo, P., & Lombard, M. (2008). The Power of the Press: The Effects of Press Frames in Political Campaign News on Media Perceptions. *Atlantic Journal of Communication*, 16(1), 1-32.
- Delli Carpini, M. X. & Keeter, S. (1996). *What Americans know about politics and why it matters*. New Haven: Yale University Press.
- Doherty, D., Dowling, C., & Miller, M. (2011) Are Financial or Moral Scandals Worse? It Depends. Forthcoming in *PS: Political Science and Politics*.
- Downs, Anthony. (1957). *An Economic Theory of Democracy*. New York: Harper & Row.
- Druckman, J. N., & Miller, J. M. (2004). The political psychology of electoral campaigns: Introduction to the symposium. *Political Psychology*, 25(4), 501-506.
- Drum, K. (2012). "In-Person Voter Fraud: Not Really a Matter of Opinion." *Mother Jones*. 15 September, 2012. Retrieved 6 October 2012, from <http://www.motherjones.com/kevin-drum/2012/09/person-voter-fraud-not-really-matter-opinion>.
- Eagly, A. H., Wood, W., & Chaiken, S. (1978). Causal inferences about communicators and their effect on opinion change. *Journal of Personality and Social Psychology*, 36, 424-435.
- Ecker, U., Lewandowsky, S., Swire, B., & Chang, D. (2011) Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychological Bulletin Review*, 18, 570-578.

Fallows, J. (2012a). "Bit by Bit It Takes Shape: Media Evolution for the 'Post-Truth' Age," *The Atlantic*, 29 August, 2012. Retrieved October 6 from <http://www.theatlantic.com/politics/archive/2012/08/bit-by-bit-it-takes-shape-media-evolution-for-the-post-truth-age/261741/>

Fallows, J. (2012b). "The real mystery of Paul Ryan's marathon time," *The Atlantic*, 1 September 2012. Retrieved 11 November 2012 from <http://www.theatlantic.com/politics/archive/2012/09/the-real-mystery-of-paul-ryans-marathon-time/261866/>.

Farley, R, Jackson, B., Kiely, E., Robertson, L. & Finley, B. (2012). "Fact check: Paul Ryan at the RNC." *USA Today*. 30 August 2012. Retrieved 6 October 2012 from <http://usatoday30.usatoday.com/news/politics/story/2012-08-30/paul-ryan-fact-check-republican-convention/57432326/1>.

Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247-282). Hillsdale, NJ: Erlbaum.

Fein, S. (1996). Effects of Suspicion on Attributional Thinking. *Journal of Personality and Social Psychology*, 70 (6): 1154 – 1184.

Fenster, J. (2012). Connecticut GOP to Democrats; Your felon is different from our felon. *The Register Citizen*, 21 July 2012. Retrieved from <http://ct5thdistrict.registercitizen.com/2012/07/21/connecticut-gop-to-democrats-your-felon-is-different-than-our-felon/> on August 5.

Fleming, J., & Arrowood, A.J. (1979). Information Processing and the Perseverance of Discredited Self-Perceptions. *Personality and Social Psychology Bulletin*, 5 (2): 201 – 205.

- Fowler, E. (2012) "Presidential Ads 70 Percent Negative in 2012, Up from 9 Percent in 2008," Retrieved 11 November 2012 from <http://mediaproject.wesleyan.edu/2012/05/02/jump-in-negativity/>.
- Funk, C. L. (1999). Bringing the candidate into models of candidate evaluation. *Journal of Politics*, 61, 700-720.
- Gant, M. M., & Davis, D. F. (1984). Negative voter support in presidential elections. *Political Research Quarterly*, 37(2), 272-290.
- Garber, M. (2012). "Is your Facebook password like your mail, house key, or drug test?" *The Atlantic*, 3 April 2012. Retrieved 1 June 2012 from <http://www.theatlantic.com/technology/archive/2012/04/is-your-facebook-password-like-your-mail-house-key-or-drug-test/255354/>.
- Garrett, R. K. (2011). Troubling consequences of online rumoring. *Human Communication Research*, 37, 255-274.
- Gawronski, B., Deutsch, R., Mbirkou, S., Seibt, B. & Strack, F. (2008) When "just say no" is not enough: Affirmation versus negation training and the reduction of automatic stereotype activation. *Journal of Experimental Social Psychology* 44:370 – 77.
- Gehrke, J. (2012). "Obama camp: Paul Ryan is a liar." *Washington Examiner*. 29 August, 2012. Retrieved 6 October 2012 from http://washingtonexaminer.com/obama-camp-paul-ryan-is-a-liar/article/2506445#.UG2soZjA_2U.

- Gerstein, J., & Samuelsohn, D. (2012). "Vice Presidential Debate Fact Check." *Politico*. 12 October, 2012. Retrieved 27 March 2013 from <http://www.politico.com/news/stories/1012/82317.html>.
- Gilbert, D., Tafarodi, R., & Malone, P. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology* 65(2), 221-233.
- Goren, P. (2002). Character weakness, partisan bias, and presidential evaluation. *American Journal of Political Science* 46 (3), 627-641.
- Graves, L. (2012a). "Digging deeper into The New York Times' fact-checking faux pas." *Nieman Journalism Lab*, January 18, 2012. Retrieved from <http://www.niemanlab.org/2012/01/digging-deeper-into-the-new-york-times-fact-checking-faux-pas>.
- Graves, L. (2012b). *Deciding What's True: Fact-Checking Journalism and the New Ecology of News*. Unpublished doctoral dissertation, Columbia University, New York, NY.
- Graves, L. & Glaisyer, T. (2012). "The fact-checking universe in spring 2012." *The New America Foundation*. 28 February 2012. 6 October 2012
- Greenhouse, L. (2012). "Challenging 'he said, she said' journalism." *Nieman Reports*. 14 July, 2012. Retrieved 6 October 2012 from <http://nieman.harvard.edu/reportsitem.aspx?id=102769>.
- Gunther, A., & Schmitt, K. (2004). Mapping boundaries of the hostile media effect. *Journal of Communication* 54(1): 55-70.

- Gurevitch, M. & Blumler, J. G. (1990). Political communication systems and democratic values. In Judith Lichtenberg, (ed.) *Democracy and the Mass Media*, Cambridge: Cambridge University Press.
- Hasson, U., & Glucksberg, S. (2006). Does negation entail affirmation? The case of negated metaphors. *Journal of Pragmatics*, 38. 1015-1032.
- Hayes, D. (2005). Candidate qualities through a partisan lens: A theory of trait ownership. *American Journal of Political Science*, 49(4), 908-923.
- Hetherington, M. J. (2001). Resurgent mass partisanship: The role of elite polarization. *American Political Science Review*, 95(3), 619-632.
- Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(6): 1420–1436.
- Kauffman, M. (2012). "Claim check: Employers asking for Facebook passwords." *Hartford Courant*, 27 March 2012. Retrieved April 3, 2012 from <http://courantblogs.com/investigative-reporting/claim-check-employers-asking-for-facebook-passwords/>.
- Kuklinski, J. H., Quirk, P. J., Schweider, D., & Rich, R. F. (1998). Just the Facts, Ma'am': Political facts and public opinion. *Annals of the American Academy of Political and Social Science*, 560, 143–154.
- Kuklinski, J. H., Quirk, P. J., Schweider, D., & Rich, R. F. (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics*, 62(3), 790–816.

- Krugman, Paul. (2011). "The Post-Truth Campaign," *The New York Times*, December 22, 2011.
- Ladd, J.M. (2012). *Why Americans hate the media and how it matters*. Princeton: Princeton University Press.
- Lau, R. R., Sigelman, L., Heldman, C., & Babbitt, P. (1999). The effects of negative political advertisements: A meta-analytic assessment. *American Political Science Review* 93(4): 851–875.
- Lawrence, R. G., & Schafer, M. L. (2012). Debunking Sarah Palin: Mainstream news coverage of 'death panels'. *Journalism*, 13(6), 766-782.
- Lee, T. T. (2010). Why they don't trust the media: An examination of factors predicting trust. *American Behavioral Scientist*, 54(1), 8-21.
- Lenz, G. (2012). *Follow the leader?* Chicago: University of Chicago Press.
- Lodge, M., McGraw, K. M., & Stroh, P. (1989). An impression-driven model of candidate evaluation. *American Political Science Review* 83:399-420.
- Lodge, M., & Taber, C. S. (2000). Three steps toward a theory of motivated political reasoning. In A. Lupia, M. McCubbins & S. Popkin (Eds.), *Elements of reason: Cognition, choice, and the bounds of rationality* (pp. 183–213). New York: Cambridge University Press.
- Lodge, M., Taber, C. S., & Verhulst, B. (2011) Conscious and unconscious processing with implications for experimental political science. In James Druckman, Donald Green, James Kuklinski, and Arthus Lupia. (Eds.). *Cambridge Handbook of Experimental Political Science* (pp. 155-170). Cambridge, Cambridge University Press.

- Marcus, George E. (1988). The Structure of Emotional Response: 1984 Presidential Candidates. *The American Political Science Review*, 82 (3), 737-761.
- Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" vs. "I am innocent": Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology*, 40: 443-449.
- McGraw, K. M., Lodge, M., & Stroh, P. (1990). On-line processing in candidate evaluation: The effects of issue order, issue importance, and sophistication. *Political Behavior*, 12(1), 41-58.
- Mindich, D. (1999). *Just the facts: How "objectivity" came to define American journalism*. New York: New York University Press.
- Muñoz-Torres, J. R. (2012). Truth and objectivity in journalism. *Journalism Studies* 13(4): 566-582.
- Mutz, D., Allen, D., Brackbill, D., Daniller, D., Garlick, A., Silver, L., & Tallevi, A. (2012). A re-examination of waning trust in the news media: what it means and why it happens. Paper presented at Annenberg School for Communication, November 2, 2012.
- Nyhan, B. & Reifler, J. (2010). When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior* 32, 202-330.
- Nyhan, B. (2009). "On The Media: NPR" [transcript] 3 July, 2009. Retrieved 6 December 2012 from <http://www.onthemedial.org/2009/jul/03/debunk-this/transcript/>.

- Nyhan, B. (2010). Why the “death panel” myth wouldn’t die: Misinformation in the health care reform debate. *The Forum* 8(1), Article 5.
- Nyhan, B. (2012). “The Future of Factchecking: Here’s what journalists should learn from the 2012 campaign” *Columbia Journalism Review* 29 November 2012. Retrieved 6 December 2012 from http://www.cjr.org/swing_states_project/the_future_of_factchecking.php.
- Ostermeier, E. (2011). “Selection Bias? PolitiFact rates Republican statements as false at 3 times the rate of Democrats.” *SmartPolitics*, February 10, 2011. Retrieved March 27 from http://blog.lib.umn.edu/cspg/smartpolitics/2011/02/selection_bias_politifact_rate.php.
- Perloff, R. (1993). Third-person effect research 1983–1992: a review and synthesis. *International Journal of Public Opinion Research* 5(2): 167-184.
- Peters, J. (2012). “A Lie Races Across Twitter Before the Truth Can Boot Up.” *The New York Times*, 9 April 2012, A12.
- Pew Research Center (2012a). Campaign 2012: Too negative, too long, dull. *Pew Research Center for The People and The Press*. 18 January, 2012. Retrieved 6 October 2012 from <http://pewresearch.org/pubs/2174/news-interest-presidential-campaign>.
- Pew Research Center (2012b). In changing news landscape, even television is vulnerable. *Pew Research Center for the People and the Press*. 27 September, 2012. Retrieved 5 October 2012 from <http://www.people-press.org/2012/09/27/in-changing-news-landscape-even-television-is-vulnerable>.

Pew Research Center (2012c). Low marks for the 2012 election. *Pew Research Center for the People and the Press*. 15 November 2012. Retrieved 6 December 2012 from <http://www.people-press.org/2012/11/15/low-marks-for-the-2012-election/>.

Pinkleton, B. (1997). The effects of negative comparative political advertising on candidate evaluations and advertising evaluations: An exploration. *Journal of Advertising* 26(1): 19-29.

Politifact. (2012). "Foreign Chinese prostitution money is allegedly behind the groups funding Congressman Scott DesJarlais's (TN-04) Republican Majority." Published July 11. Retrieved August 10, 2012, from <http://www.politifact.com/tennessee/statements/2012/aug/01/democratic-congressional-campaign-committee/desjarlais-taking-money-linked-chinese-prostitutio/>.

Popkin, S. L. (1991). The reasoning voter: Communication and persuasion in presidential elections. *Chicago: University of Chicago*.

Rainey, J. (2012). "Paul Ryan, the credit rating, and 'post-truth politics.'" *Los Angeles Times*. 10 September, 2012. Retrieved 6 October 2012 from <http://articles.latimes.com/2012/sep/10/news/la-paul-ryan-post-truth-politics-20120910>.

Ramsay, C., Kull, S., Lewis, E., & Subias, S. (2010). Misinformation and the 2010 election. University of Maryland. Retrieved 5 July 2012 from http://www.worldpublicopinion.org/pipa/pdf/dec10/Misinformation_Dec10_rpt.pdf.

- Redlawsk, D. (2002). Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making. *Journal of Politics*, 64(4): 1021–1044.
- Rosen, J. (2010). "Fixing the ideology problem in our political press." *Pressthink*. 22 June 2010. Retrieved 6 October 2012 from http://archive.pressthink.org/2010/06/22/reply_ambinder.html.
- Scherer, M. (2012). "Why Deceit is Everywhere in the 2012 Campaign." *Time*, 10 August, 2012. Retrieved 6 October 2012 from <http://swampland.time.com/2012/08/10/why-deceit-is-everywhere-in-the-2012-campaign>.
- Scherer, R. (2012). "Cory Booker fire rescue 'very heroic' but very dangerous, fire officials say," 13 April, 2012. Retrieved 11 November 2012 from <http://www.csmonitor.com/USA/2012/0413/Cory-Booker-fire-rescue-very-heroic-but-very-dangerous-fire-officials-say>.
- Schudson, M. (2012). The objectivity norm in American journalism. *Journalism* 2(2): 150.
- Sedikides, C., & Green, J. D. (2001). Affective influences on the self-concept: Qualifying the mood-congruency principle. In J. P. Forgas (Ed.), *Handbook of affect and social cognition*, 145–160. London: Erlbaum.
- Sherman, D. & Kim, H. (2002). Affective perseverance: The resistance of affect to cognitive invalidation. *Personality and Social Psychology Bulletin*, 28(2): 224–237.

- Spivak, C. (2010). The fact-checking explosion. *American Journalism Review*. 2 December 2010. Retrieved 6 December 2012 from <http://www.ajr.org/article.asp?id=4980>.
- Sullivan, M. (2012). "He Said, She Said, and the Truth." *The New York Times*. 15 September, 2012. Retrieved 6 October 2012 from <http://www.nytimes.com/2012/09/16/public-editor/16pubed.html>.
- Taber, C. S., Lodge, M., & Glathar, J. (2001). The motivated construction of political judgments. In James H. Kuklinski (ed.) *Citizens and Politics: Perspectives from Political Psychology*, New York: Cambridge University Press.
- Thornburgh, N. (2008). "Mayor Palin: A Rough Record." *Newsweek*, September 2, 2008. Retrieved 6 October 2012 from <http://www.time.com/time/politics/article/0,8599,1837918,00.html>.
- Thorson, E. (2011). Political misinformation's continuing effects on attitudes: motivated reasoning or belief perseverance? Annual Meeting of the American Political Science Association, September 2011 in Seattle, Washington.
- Vallone, R. P., Ross, L., & Lepper, M. R. (1985). The hostile media phenomenon: Biased perception and perceptions of media bias in coverage of the Beirut massacre. *Journal of Personality and Social Psychology* 47: 577-585.
- Watts, M.; Domke, D., Shah, D. V., & Fan, D. P. (1999). Elite cues and media bias in presidential campaigns: Explaining public perceptions of a liberal press. *Communication Research* 26(2): 144-175.

Weeks, B. & Southwell, B. (2010). The symbiosis of news coverage and aggregate online search behavior: Obama, rumors, and presidential politics. *Mass Communication and Society* 13: 341-260.

Wegner, D. M., Coulton, G. F., & Wenzlaff, R. (1985). The transparency of denial: Briefing in the debriefing paradigm. *Journal of Personality and Social Psychology*, 49(2), 338-346.

Wegner, D. M., Wenzlaff, R., Kerker, R.M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers?. *Journal of Personality and Social Psychology* 40:822-832.

Wimple, E. (2012). Fact-Checking: A Consumer-Driven Movement. 2 October 2012. Retrieved 6 October 2012 from http://www.washingtonpost.com/blogs/erik-wemple/post/fact-checking-a-consumer-driven-movement/2012/10/02/6a06c404-0cbd-11e2-bd1a-b868e65d57eb_blog.html.

Woodly, D. (2007). New competencies in democratic communication? Blogs, agenda setting and political participation. *Public Choice* 134(1-2): 109-123.

Woodward, C. & Gillum, J. (2012). "Fact check: Ryan takes factual shortcuts in speech." *Associated Press*, 30 August, 2012. Retrieved 6 October 2012 from <http://news.yahoo.com/fact-check-ryan-takes-factual-shortcuts>.

Appendix

Appendix A. Article Manipulations

A1. Experiment 1 First Article



The★Star
THE KANSAS CITY STAR

Kansas Congressional Race Heats Up

By Samuel Cullen, Kansas City Star

WICHITA, KS. — For Kansas [Republicans/Democrats], November's Congressional election is a must-win. But the race is far from decided, and [Republican/Democratic] hopeful John McKenna is working hard to win over Wichita voters.

McKenna, the former mayor of Wichita and [Republican/Democratic] candidate for Congress, gave an impassioned speech on Tuesday evening, promising to “bring back the American dream” for Kansas residents.

[Misinformation condition] However, McKenna’s campaign may be hurt by a relationship that has only recently come to light. New statewide campaign finance requirements forced McKenna to release the full list of contributors to his mayoral campaign. Records show that McKenna received several substantial donations from convicted felon Daniel Elsio, who formerly headed up the largest organized crime syndicate in Kansas. Many hold Elsio responsible for a 51% increase in crystal meth use in Kansas schools between 2000 and 2005. Elsio was convicted in 2006 of narcotics trafficking and second-degree murder. Records show that Elsio also attended several [Republican/Democratic] party fundraisers.

At his speech on Tuesday, McKenna was accompanied by his wife Karen as well as his cousin, film actor Robert Downey, Jr. Downey, star of the upcoming movie "The Avengers," spoke briefly about his and McKenna’s childhood: “When we were younger, my parents always asked why I couldn’t be more like John. They’re still asking.”

McKenna’s speech was followed immediately by a \$500-a-head fundraising dinner held at the Auberge Hotel in downtown Wichita.

McKenna, 48, has a law degree from the University of Kansas and practiced law for ten years in a private firm. He served as Wichita’s mayor from 2002 to 2008.

McKenna’s competition in November is local businessman and long-time [Republican/Democratic] activist Eric Hall, who has never held office in Kansas. A debate between the two candidates has been scheduled for next Friday, to be held at the Lions Club in downtown Wichita. The debate will be broadcast on KPTS Channel 8.



[Republican/Democrat] John McKenna spoke to supporters on Tuesday. He was accompanied by his cousin, actor Robert Downey, Jr.

A2. Experiment 1 Second Article

Congressional Candidates Clash Over Education

By Samuel Cullen, Kansas City Star

WICHITA, KS - The two candidates competing in the 4th District's open seat aired their differences about education policies at a debate held on Friday at the Lions Club in downtown Wichita.

[Republican/Democrat] John McKenna and [Democrat/Republican] Eric Hall spent over two hours answering questions about their personal histories and policy positions. The debate featured questions submitted by viewers as well as questions from the moderator.

Several viewer questions concerned the Kansas public school system. Over the past year, Kansas has been forced to substantially cut school budgets, even as test scores across the state have plummeted. Both candidates agreed that the current situation is unsustainable, but disagreed on what should be done. McKenna, who served as Wichita's mayor from 2002 until 2010, says he believes that a focus on the basics is important:

"Reading, writing, arithmetic. Not that much has changed in the last fifty years, and we need to make sure that our kids are getting the basic skills they need before they move on to the next level."

Hall, on the other hand, stressed the importance of frequent testing to pinpoint problems:

"We can't fix it unless we know where it's broken. We should be testing along the way to ensure that kids are learning what they should be learning."

Both candidates also discussed their personal histories. McKenna, a [Republican/Democrat], focused on the successful policies he implemented while mayor. Hall, a [Republican/Democrat], discussed his experience starting and growing a small business in Wichita.

A telephone poll of 200 likely voters conducted prior to the debate shows that McKenna is leading among Democrats and Hall among Republicans. However, many voters (including 76% of Independents) remain undecided.

[Correction condition] *Correction: Regarding the candidates' biographies, an article published in the Kansas City Star last week stated that Daniel Elsio, a felon convicted of drug trafficking and murder, was a frequent donor to John McKenna's mayoral campaign and attended several Democratic party fundraisers. However, further investigation of the campaign donation records has shown that the donor listed was actually Daniel Elio, the owner of a local car dealership.*

A3. Article text used for Experiments 2, 4, and 5

The following text comprises the basic article text used in Experiments 2, 4, and 5. Any variations on this format are discussed in the design sections for the individual experiments (for example, Experiment 5 varies the source of the correction).

The Iowa★Ledger

Iowa Congressional Race Heats Up

by Samuel Cullen, Iowa Ledger

For Iowa Republicans, November's Congressional election is a must-win. But the race is far from decided, and [Democratic/Republican] hopeful John McKenna is working hard to win over Iowa voters.

McKenna, the former mayor of Cedar Rapids and [Democratic/Republican] candidate for Congress, gave an impassioned speech on Tuesday evening, promising to “bring back the American dream” for Iowa residents. McKenna previously served as the mayor of Cedar Rapids.

At his speech, McKenna was accompanied by his cousin, film actor Robert Downey, Jr.

Downey, star of the recent movie “The Avengers,” spoke briefly about his and McKenna’s childhood: “When we were younger, my parents always asked why I couldn’t be more like John. They’re still asking.”

McKenna’s speech was followed immediately by a \$500-a-head fundraising dinner held at the Auberge Hotel in downtown Cedar Rapids. The dinner was attended by several prominent Iowa [Democrats/Republicans].

McKenna’s competition in November is local businessman and long-time [Democratic/Republican] activist Eric Hall, who has never held office in Iowa. Hall has the full support of the Iowa [Democratic/Republican] party and has already begun running television ads.

[Misinformation condition] The campaign has heated up in recent weeks. On Sunday, Hall accused McKenna of accepting campaign donations from a convicted felon named Daniel Elsio. Elsio, who ran the largest drug ring in Iowa while McKenna was mayor, was convicted of first-degree murder in 2010. According to Hall, McKenna accepted over \$10,000 from Elsio. Hall commented that “McKenna’s corrupt behavior shows that he and other [Democrats/Republicans] can’t be trusted to do what’s best for Iowa citizens.”

[Correction condition] However, further investigation of the campaign donation records by journalists at the *Iowa Ledger* has shown no record of any donation from Elsio to McKenna’s campaign. Campaigns are required to disclose the names of all individuals who contribute \$200 or more in an election cycle, and the *Ledger* did not find Elsio’s name listed.



Democrat/Republican John McKenna spoke to supporters Tuesday. He was accompanied by his cousin, actor Robert Downey, Jr.

This race is attracting national attention, as Iowa is a critical swing state for the Presidential election. Mitt Romney and Barack Obama both plan to visit Cedar Rapids in the coming months.

McKenna, the [Democratic/Republican] candidate, and Hall, the [Democratic/Republican] candidate, will debate next Friday. The debate will be held at the Lions Club in downtown Cedar Rapids and be broadcast on Channel 8, WQPT.

A4. Experiment 3 candidate background information

The party affiliation of the candidate varied depending on the party of the subject and the condition to which she was opposed (same-party or opposing-party)



Calvin Mitchell, Democrat

- Incumbent Representative (elected in 2008)
- Former lawyer



George Ross, Republican

- Small business owner
- Long-time Republican activist

A5. Experiment 3 Headlines and Photos

All headlines and photos were in the format shown in Figure 4.2 (featuring the Nevada Tribune headline and blurred text next to the photo and below the headline).

Article	Headline/subhead	Photo
Article 1 (all subjects)	Nevada Congressional Race Heats Up <i>Voters report more mailings, phone calls, and visits as election nears</i>	
Article 2 (all subjects)	[Democrat/Republican] Mitchell, [Republican/Democrat] Ross to Face Off <i>Economy and health care will be the focus of Friday's debate</i>	
Article 3A (Misinfo 1 group)	Felon Donation Scandal Hits Mitchell Campaign <i>Ross accuses Mitchell of accepting campaign donations from convicted felon Joe Fenz</i>	
Article 3B (Misinfo 2 group)	Past Policies Return to Haunt Mitchell Campaign <i>Ross says Mitchell's health care reforms increased the number of newborn deaths in Nevada</i>	

Article 3C (Misinfo 3 group)	Corruption Scandal Hits Mitchell Campaign <i>Ross accuses Mitchell of spending \$294,000 of taxpayer money to finance three family vacations to the Florida coast</i>	
Article 4 (Correction 1 group)	Mitchell's Donation Records Show No Felon Donations <i>A Nevada Tribune investigation of Mitchell's finances found no record of donations from convicted felon Joe Fenz</i>	
Article 4 (Correction 2 group)	Newborn Deaths in Nevada Did Not Increase Under Mitchell's Hospital Reforms <i>A Nevada Tribune investigation of hospital records shows no rise in infant mortality in Nevada hospitals since Mitchell's policies were implemented</i>	
Article 4 (Correction 3 group)	Mitchell's Vacation not Paid for with Taxpayer Money <i>A Nevada Tribune investigation of Mitchell's finance reports shows he did not use public funds to pay for his vacations</i>	
Article 5	Mitchell Holds Fundraiser in Carson City <i>Local [Democratic/Republican] leaders pay \$450 a head for dinner with the congressional candidate</i>	

Appendix B. Question Wording

B1. Demographic battery

What is your age?

- ☐ 18 - 25
- ☐ 26 - 40
- ☐ 41 - 55
- ☐ 56 or older

What is the last grade or class you completed in school?

- ☐ Did not finish high school
- ☐ High school diploma or equivalent, no further schooling
- ☐ Technical or vocational school after high school
- ☐ Some college, no degree
- ☐ Associate's or two-year college degree
- ☐ Four-year college degree
- ☐ Graduate or professional school after college, no degree
- ☐ Graduate or professional degree

What is your gender?

- ☐ Male
- ☐ Female

In general, do you disagree or agree with the following statements?

	Disagree	Agree
You can generally trust the people who run our government to do what is right.	<input type="radio"/>	<input type="radio"/>
When government leaders make statements to the American people, they are usually telling the truth.	<input type="radio"/>	<input type="radio"/>
The people we have elected to public office usually try to keep the promises they have made.	<input type="radio"/>	<input type="radio"/>

How interested are you in politics and public affairs?

- ☐ Very interested
- ☐ Somewhat interested
- ☐ Slightly interested
- ☐ Not at all interested

Generally speaking, do you think of yourself as:

- ☐ Republican
- ☐ Democrat
- ☐ Independent
- ☐ Another party
- ☐ No preference

Would you call yourself a strong \${q://QID5/ChoiceGroup/SelectedChoices} or a not so strong \${q://QID5/ChoiceGroup/SelectedChoices}?

- ☐ Strong \${q://QID5/ChoiceGroup/SelectedChoices}
- ☐ Not so strong \${q://QID5/ChoiceGroup/SelectedChoices}

Do you think of yourself as closer to the Democratic party or the Republican party?

- ☐ The Democratic Party
- ☐ The Republican Party
- ☐ Neither party



Do you usually think of yourself as:

- ☐ Extremely liberal
- ☐ Liberal
- ☐ Slightly liberal
- ☐ Slightly conservative
- ☐ Conservative
- ☐ Very conservative

B2. Attitude Elicitation

The following questions comprise the basic attitude elicitation battery used in each of the experiments. Any variations on this format are discussed in the design sections for the individual experiments.

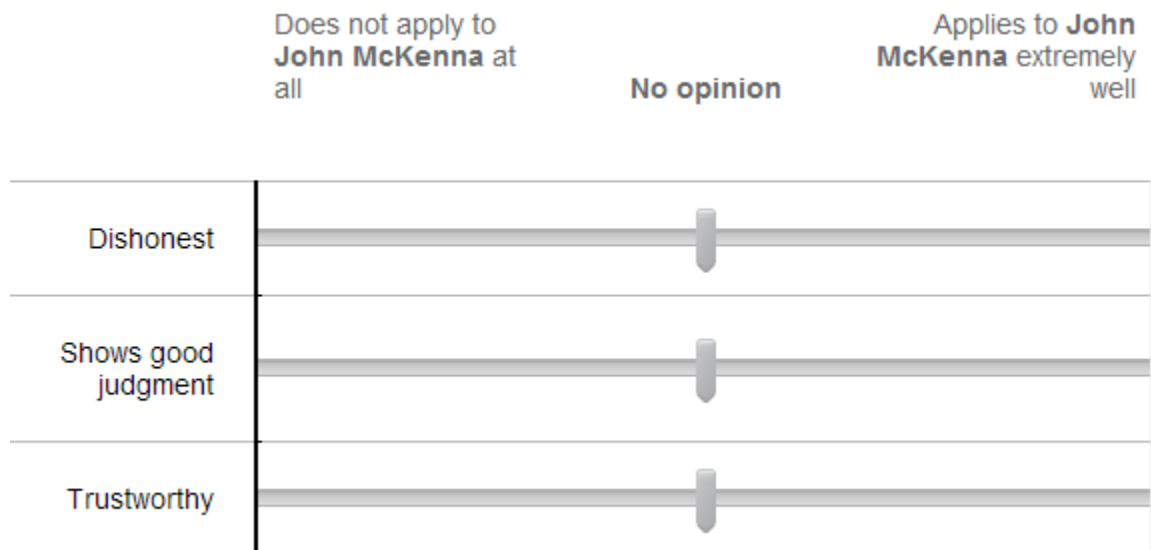
Please rate both the candidates you read about on the thermometer below. The thermometer runs from 0 to 100 degrees. Rating above 50 means that you feel favorable and warm toward the person. Rating below 50 means that you feel unfavorable and cool toward the person.

	Unfavorable/cool	Favorable/warm
John McKenna		
Eric Hall		

Please indicate how much you agree with the following statements.

	Strongly Disagree	Disagree	Neither Agree nor Disagree	Agree	Strongly Agree
Eric Hall is a high-quality candidate.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
John McKenna would make a good representative.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
John McKenna may not be suited for public office.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

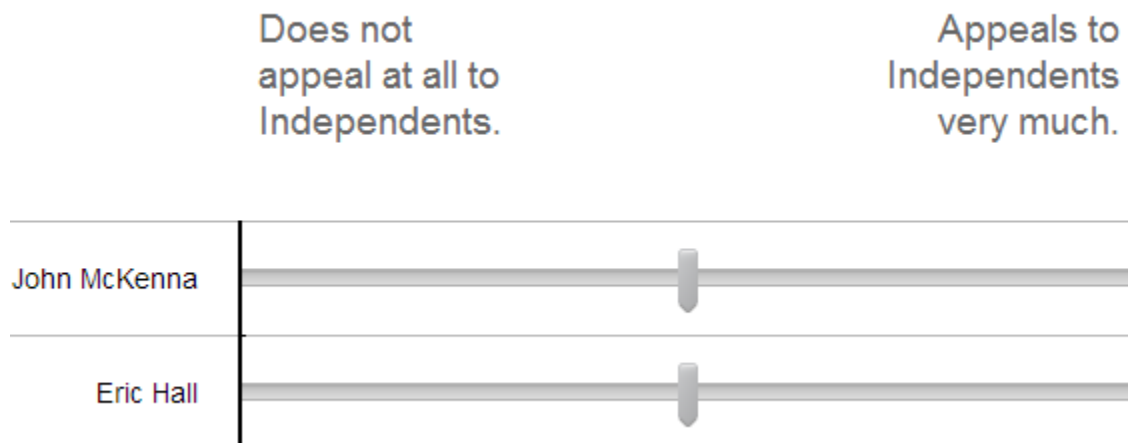
For each of the following phrases, please indicate how well that phrase applies to Congressional candidate John McKenna. [Respondents completed same for Eric Hall]



If you had to guess, which candidate do you think was more successful at recruiting volunteers for his campaign?

- ☐ John McKenna
- ☐ Both candidates were equally successful
- ☐ Eric Hall

In general, how much do you think each of the candidates appeals to Independent voters (in other words, voters who are not affiliated with either the Democratic or Republican party)?



If you had to guess, who do you think won the election in November? Just give your best guess based on what you've read.

- ☐ John McKenna definitely won.
- ☐ John McKenna probably won.
- ☐ Eric Hall probably won.
- ☐ Eric Hall definitely won.

B3. Media Evaluation

For each media outlet, subjects were shown the following three questions on a single page. At the top of the page was a logo from the outlet under consideration. The three questions below concern GetTheFacts.org (the logo was constructed for the purposes of the survey).


The following questions are about the fact-checking organization GetTheFacts.org.



In your opinion, do you think that GetTheFacts.org tends to favor one political party over the other when it fact-checks stories?

- ☐ Strongly favors the Democrats
- ☐ Weakly favors the Democrats
- ☐ Neutral/unbiased/not sure
- ☐ Weakly favors the Republicans
- ☐ Strongly favors the Republicans

Overall, is your opinion of GetTheFacts.org...





How much of the time do you think you can trust GetTheFacts.org to get the facts right?


- ☐ Just about always
- ☐ Most of the time
- ☐ Only some of the time
- ☐ Almost never

Appendix C. Distractor Task

C1. Distractor task used in Experiment 1 (between Article 1 and Article 2)

Are the following people actors, news anchors, or aren't you sure?


	Movie Actor	News Anchor	Not sure
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>



C2. Distractor task used in Experiments 2, 3, 4, and 5


Subjects saw the images below and were instructed “When you spot a difference, click on either photo to mark it. If you find all five before the two minutes are up, hit the ‘next’ button to go to the next screen.”

The red number at the top of the page counted down from “120” until 0. The average time spent on the distractor task (across experiments) was 105.4 seconds. When subjects clicked on an area that was different, the area was shaded green so they could see that they had found a difference. For example, in the image below, the subject has found two differences.



When the 120 seconds was up (or when subjects clicked the “next” button), they were presented with the answers below.

The differences are circled in the photos below.



Appendix D. Demographics

D1. Demographics for Experiment 1 and Experiment 2

	Experiment 1 (N=157)		Experiment 2 (N=474)	
	Freq	Percent	Frequ	Percent
18 - 25	49	31.2	164	33.8
26 - 40	69	43.9	210	43.3
41 - 55	30	19.1	78	16.1
56 or older	9	5.7	31	6.4
Female	82	52.2	286	59
Did not finish high school	12	7.6	5	1
High school diploma or equivalent, no further schooling	12	7.6	54	11.1
Technical or vocational school after high school	7	4.5	14	2.9
Some college, no degree	37	23.6	147	30.3
Associate's or two-year college degree	21	13.4	58	12
Four-year college degree	50	31.8	136	28
Graduate or professional school after college, no degree	9	5.7	30	6.2
Graduate or professional degree	21	13.4	41	8.5
Republican	30	19.1	72	14.8
Democrat	49	31.2	192	39.6
Independent	50	31.8	141	29.1
Another party	8	5.1	16	3.3
No preference	20	12.7	64	13.2
Very interested	48	30.6	103	21.2
Somewhat interested	55	35	202	41.6
Slightly interested	42	26.8	131	27
Not at all interested	12	7.6	49	10.1

D2. Demographics for Experiments 3, 4, and 5

Note that Experiment 5 shares 474 subjects with Experiment 2.

Some demographic questions were not asked in Experiment 3 for reasons of time

	Experiment 3 (N=178)		Experiment 4 (N=309)		Experiment 5 (N=606)	
	Freq	Percent	Freq	Percent	Freq	Percent
18 - 25	60	33.7	86	34.3	209	34.5
26 - 40	87	48.9	113	45	263	43.4
41 - 55	24	13.5	40	15.9	93	15.3
56 or older	7	3.9	12	4.8	39	6.4
Female	88	49.4	135	53.8	358	59.1
Did not finish high school	N/A	N/A	1	0.4	7	1.2
High school diploma or equivalent, no further schooling	N/A	N/A	44	17.5	69	11.4
Technical or vocational school after high school	N/A	N/A	9	3.6	17	2.8
Some college, no degree	N/A	N/A	80	31.9	179	29.5
Associate's or two-year college degree	N/A	N/A	29	11.6	64	10.6
Four-year college degree	N/A	N/A	57	22.7	180	29.7
Graduate or professional school after college, no degree	N/A	N/A	5	2	34	5.6
Graduate or professional degree	N/A	N/A	26	10.4	56	9.2
Republican	25	14	51	20.3	98	16.2
Democrat	82	46.1	114	45.4	232	38.3
Independent	58	32.6	61	24.3	175	28.9
Another party	4	2.2	8	3.2	20	3.3
No preference	9	5.1	17	6.8	81	13.4
Very interested	N/A	N/A	59	23.5	131	21.6
Somewhat interested	N/A	N/A	121	48.2	238	39.3
Slightly interested	N/A	N/A	55	21.9	172	28.4
Not at all interested	N/A	N/A	16	6.4	65	10.7

Appendix E. Other experiments conducted as part of research for this dissertation

E1. Full list of experiments

Name	Pretest 1 (not in dissertation)
Date	August 2011
Design	300 participants, between-subjects, 3 (types of misinformation) x 2 (corrected and uncorrected) plus a control group.
Goals	<ol style="list-style-type: none"> 1) To determine the impact of three types of candidate misinformation (cheating in law school, accepting illegal campaign donations, and soliciting a prostitute) on attitudes towards that candidate. 2) To determine whether misinformation can be successfully corrected. 3) To determine whether belief echoes are created when no party cues are present.
Findings	<ol style="list-style-type: none"> 1) Each type of misinformation was equally effective at lowering evaluations (no significant difference between the three) 2) Misinformation can be successfully corrected. 3) Only weak evidence of belief echoes in the absence of any party cues.

Name	Pretest 2 (not in dissertation)
Date	November 2011
Design	685 participants, between-subjects, 3 conditions (corrected misinfo, no misinfo, control). For all participants, the article concerned a candidate of the opposite party (Independents were forced to identify a lean). Correction was in a separate article than the misinformation, separated by a distractor task.
Goals	<ol style="list-style-type: none"> 1) Test for the existence of belief echoes in the presence of party cues: those who receive corrected negative misinformation should evaluate the candidate more negatively than those who receive no misinformation or correction.
Findings	<ol style="list-style-type: none"> 1) Belief echoes exist: even though the misinformation was successfully corrected, exposure to the misinformation continues to affect attitudes.

Name	Effect of time delay on belief echoes (some data used for Experiment 2)
Date	February 2012

Design	800 participants, between-subjects 2 (article 1 misinformation, article 1 no misinformation) x 2 (article 2 correction, article 2 no correction) x 2 (time delay of up to 1 week between article 1 and article 2, no time delay)
Goals	<ol style="list-style-type: none"> 1) Test for the existence of belief echoes (replicate findings in Pretest 2) 2) Does increasing the length of time between misinformation and correction magnify belief echoes? 3) Do belief echoes persist over time?
Findings	<ol style="list-style-type: none"> 1) Belief echoes exist: again, even though the misinformation was successfully corrected, exposure to the misinformation continues to affect attitudes. 2) Imposing a time delay between the first article (with the misinformation) and the second article (with the correction) does not increase or decrease the magnitude of belief echoes. 3) Participants contacted 5 – 7 days after reading the initial misinfo and correction continue to evaluate the candidate more negatively: belief echoes persist.

Name	Correction Effects (data used for Experiments 1 and 5)
Date	June 2012
Design	640 participants, between-subjects 2 (candidate is same party, candidate is opposing party) x 5 (corrected by fact-checking org, corrected by opposing party, corrected by media outlet, uncorrected misinformation, no misinformation)
Goals	<ol style="list-style-type: none"> 1) Determine which correction is most effective at reducing belief in misinformation. 2) Determine whether partisanship affects the size of belief echoes: are they stronger when misinformation reinforces pre-existing partisan loyalties? 3) Measure the effect of publishing corrected misinformation on media trust more generally as well as trust in the specific media outlet that published the correction.
Findings	<ol style="list-style-type: none"> 1) Corrections issued by a fact-checking organization or a newspaper are significantly more effective than those issued by the campaign itself, even when the content is identical. 2) Belief echoes are of a similar magnitude regardless of the party of the candidate. 3) Reading corrections increases trust in media more generally and decreases trust in the media outlet that published the

correction only if the correction contradicts pre-existing partisan preferences.

Name	Affirmation test (data used for Experiment 3)
Date	October 2012
Design	300 participants, 2 (candidate is same party, candidate is opposing party) x 2 (affirmation or negation)
Goal	<ol style="list-style-type: none"> 1) To determine whether corrections in the form of affirmations are more effective than corrections in the form of negations at muting belief echoes. 2) To determine whether in a situation of high cognitive processing, belief echoes will only be created among those predisposed to dislike the candidate
Findings	<ol style="list-style-type: none"> 1) Affirmations and negations were equally successful at muting belief echoes. 2) High cognitive processing does magnify belief echoes for those predisposed to dislike the candidate.

Experiment	Recall task (data used for Experiment 4)
Date	November 2012
Design	250 participants, 2 (candidate is same party, candidate is opposing party) x 3 (corrected misinformation, uncorrected misinformation, no misinformation) x 2 (completed recall task, did not complete recall task)
Goal	<ol style="list-style-type: none"> 1) To determine whether actively recalling the misinformation and correction creates cognitive belief echoes (among those who are predisposed to dislike the candidate) 2) To investigate the responses to open-ended questions.
Findings	<ol style="list-style-type: none"> 1) Recall task heightens belief echoes among those predisposed to dislike the candidate. In contrast, among those who did not complete the recall task, belief echoes are the same magnitude regardless of partisanship. 2) Open-ended responses show evidence for cognitive belief echoes.