The Affective Tipping Point: Do Motivated Reasoners Ever "Get It"?

David P. Redlawsk
Rutgers University

Andrew J. W. Civettini
Knox College

Karen M. Emmerson
Lake Research Partners

In order to update candidate evaluations voters must acquire information and determine whether that new information supports or opposes their candidate expectations. Normatively, new negative information about a preferred candidate should result in a downward adjustment of an existing evaluation. However, recent studies show exactly the opposite; voters become more supportive of a preferred candidate in the face of negatively valenced information. Motivated reasoning is advanced as the explanation, arguing that people are psychologically motivated to maintain and support existing evaluations. Yet it seems unlikely that voters do this ad infinitum. *To do so would suggest continued motivated reasoning even in the face of extensive disconfirming information. In this study we consider whether motivated reasoning processes can be overcome simply by continuing to encounter information incongruent with expectations. If so, voters must reach a tipping point after which they begin more accurately updating their evaluations. We show experimental evidence that such an affective tipping point does in fact exist. We also show that as this tipping point is reached, anxiety increases, suggesting that the mechanism that generates the tipping point and leads to more accurate updating may be related to the theory of affective intelligence. The existence of a tipping point suggests that voters are not immune to disconfirming information after all, even when initially acting as motivated reasoners.*

What does it mean for a voter to be rational? At a minimum, rational voters know their own preferences, update those preferences accurately upon receipt of new information, and choose the candidate that best represents their interests. Such voters should be predictable in the sense that when they encounter new information about a candidate, evaluations of that candidate will be adjusted up or down as appropriate. Yet, real people are not nearly the predictable "cool calculators" rational models seem to require (Redlawsk, 2002). Recent research has convincingly shown that emotions play an important part in most decision-making realms. While classical political thought drew distinctions between reason and emotion, reconceptualizations from neuroscience (Damasio, 1999) to political science (Lodge & Taber, 2000, 2005; Marcus, Newman, & MacKuen, 2000) demonstrate that emotions must be an integral part of political decision-making processes.

Existing affective evaluations color how people think about candidates (Redlawsk, Civettini, & Lau, 2007) and issues (Lodge & Taber, 2000, 2005) and how new information is processed as it is learned during a campaign. Rational updating requires that negative information lower the evaluation of a candidate, while positive information must do the opposite (Green & Gerber, 1999). But what if the candidate is one a voter already likes; a candidate whom the voter has already decided is "good"? What happens when negative information is encountered about that candidate? A developing body of research shows that voters may operate as motivated reasoners, attempting to hold to their existing positive evaluation by using any one of a number of processes to explain away new incongruent information (Kunda, 1990; Lodge & Taber, 2000; Taber & Lodge, 2006; Redlawsk, 2002). In other words, existing affect may interfere with accurate updating.

We argue that existing affect towards an already known candidate is an important factor in determining the extent to which new information is accurately perceived and evaluations correctly updated. We present evidence that voters ignore significant amounts of negative information about positively evaluated candidates. In fact, voters may become even *more* positive about a candidate they like after learning something negative about that candidate. This tendency, so contrary to classical notions of rational updating, is consistent with theories of motivated reasoning (Taber & Lodge, 2006), research on a "conservation bias" (Steenbergen, 2001), and the concept of cognitive dissonance (Festinger, 1957). Yet we do not believe this can go on without end. At some point even the most strongly held positive evaluation should flag in the face of repeated negative information relevant to the evaluation. In this paper, we demonstrate that there is, in fact, a point at which voters stop reinforcing their preferences, abandon motivated reasoning, and begin "rational" updating. We call this the affective tipping point.[1]

---

[1] While the concept of a "tipping point" is not especially new, Gladwell's (2000) book *The Tipping Point* has brought the idea into the public imagination. It seems a very effective way to describe the point at which things change, in a sense the straw that breaks the camel's back.

Theoretical Perspective

Standard cognitive/rational models do not appropriately account for affect in the evaluation and decision-making process. Green and Gerber's (1999) description of a Bayesian updating process in which new information updates prior beliefs is typical. New information in agreement with an existing candidate evaluation is assumed to strengthen that evaluation. Information to the contrary does the opposite. Thus a voter's early positive impressions of a candidate will be strengthened and made more positive by learning something good about that candidate. But, when the voter learns something disagreeable about that same candidate, the prior evaluation will be updated in a negative direction to account for this new negative information. While the updating process itself need not be linear, nowhere in this model is there any suggestion that existing affect towards the candidate might actually impede attitude change. Further, there is no serious consideration given to the possibility of asymmetric effects, that negative information might be weighted differently than positive information in the updating process (Redlawsk, 2007).

Holbrook, Krosnick, Visser, Gardner, and Cacioppo (2001) attempt to address this nonlinear asymmetric possibility and in doing so propose a more nuanced, but still "rational" updating process. In their Asymmetric Nonlinear Model (ANM) positive information, while carrying more weight at the beginning of an updating process, decreases in importance more rapidly than negative information. Thus, in order to update, a voter must assess the direction of the information before incorporating it into an evaluation. This modification of the updating model, while taking into account the asymmetric nature of positive and negative information, proceeds in the same vacuum that other cognitive approaches inhabit; the *existing* evaluation of the candidate serves merely to anchor the updated evaluation, but does not directly condition the updating process itself.

Enter Affect

Since Festinger's (1957) description of cognitive dissonance and Heider's (1958) development of balance theory, psychological studies of affect and updating have regularly suggested that cognitive processes are not so straightforward, and certainly do not proceed in a vacuum. The hot cognition thesis (Abelson, 1963) argues that affect and cognition are inexorably linked; for every concept or piece of information in memory there is an associated affective evaluation that is activated whenever the concept is accessed (Lodge & Taber, 2000, 2005; Taber & Lodge, 2006). Whether positive or negative, affect cannot be separated from the underlying information, so theories that focus only on cognitive updating can only tell part of the story. Indeed, Zajonc (1984) convincingly argues for the primacy of affect; that affective responses occur before conscious processing. There are both cognitive and affective systems wired in the human brain, and these may well work in parallel and independently.

Marcus and his colleagues (Marcus & MacKuen, 1993; Marcus et al., 2000) propose a theory of affective intelligence to describe the direct effects of the affective systems on cognitive processes. Affective responses are the result of a dual process emotional system: a behavioral inhibition system and a behavioral approach system (Marcus & MacKuen, 1993). The first system compares a new stimulus to existing expectations, and if a stimulus is found to be incongruent with those expectations, attention is shifted to it. The new stimulus is thus a potential threat, and this perceived threat generates negative affect like anxiety, interrupting normal (essentially below consciousness) processing. This interruption leads to active processing, where both attention to the new information and the time taken to process it increases. Thus, negative affect directly motivates the individual to learn more about the stimulus and the environment in general. If affective intelligence is right, we would expect that during an election campaign anxious voters would be more attentive, more informed, and more likely to make good choices (Lau & Redlawsk, 1997; Lau, Anderson, & Redlawsk, 2008). Marcus and his colleagues provide evidence that anxious voters show more learning and attention to campaigns, while calm voters pay less attention (Marcus et al., 2000). And while they do not test it directly, the logical conclusion of their theory is that affectively intelligent voters should look more like rational updaters.

But other research on affect and its effects on cognition somewhat muddy these waters. Holbrook (2005) used political ads to generate positive, negative, and neutral affect in subjects. Highly anxious individuals were more responsive to new information, but overall they were less able to accurately recall information after the fact. On the other hand, Brader (2005), also using political ads to generate affect, found that anxious subjects were more likely to recall information related to the issue in the ads, but failed to seek out more information. Isen (2000) argues that positive affect improves cognitive processing; contrary to Marcus et al., affectively positive individuals are more likely to ignore incongruent information, rather than pay special attention to it. Finally, other evidence supports the notion that anxious voters pay more attention and accurately process more information, but this effect may be limited to conditions where there is a great deal of incongruent information in the environment (Redlawsk, Civettini, & Lau, 2007). Feeling just a little bit anxious might not be enough to trigger careful attention.

This latter point may be crucial to understanding the apparent contradictions in the research on anxiety and updating. The affective intelligence thesis was tested by Marcus in the aggregate, using American National Election Studies data, collected during the particularly rich political environment of a presidential campaign. Most of the other studies have examined negative affect in a more limited form in a laboratory with subjects exposed to relatively little information. If negative affect has its greatest effect when there is a great deal to be anxious about—for example, when there are significant amounts of affectively incongruent

information which increases the threat to existing evaluations—we might not find consistent effects unless the information environment becomes especially threatening to existing beliefs. Why not? Because other processes, such as motivating reasoning, may be working at cross purposes with affective intelligence.

Where affective intelligence argues that negative affect may produce better decisions, motivated reasoning suggests that this is not quite the whole story (Kunda, 1990; Taber & Lodge, 2006; Redlawsk, 2002, 2006). Motivated reasoners make an immediate evaluation of new information and use it to update an online tally that summarizes their evaluative affect (Hastie & Park, 1986; Lodge, McGraw, & Stroh, 1989; Lodge, Steenbergen, & Brau, 1995; Redlawsk, 2001). Newly encountered information carries with it an affective value. Given an existing evaluation (represented by the online tally), these affective components interact so that the online tally directly influences how the new information is evaluated *before* it is used to update the tally. This is the key insight missing from both the cognitive approaches and affective intelligence. Even anxious voters presumably motivated to learn more and make more accurate assessments may well be subject to the processing biases of motivated reasoning as they affectively evaluate *before* they begin to cognitively process new information.

While a negative emotional response may be generated by incongruence between expectations (existing affect as summarized by the online tally) and new information, motivated reasoning suggests that this incongruence does not necessarily lead to greater accuracy in evaluation or greater information search. Instead voters committed to a candidate may be motivated to discount incongruent information; they may mentally argue against it, bolstering their existing evaluation by recalling all the good things about a liked candidate even in the face of something negative. Motivated reasoning describes an interaction between existing affective evaluations and new information, but unlike affective intelligence, the effect of affect may lead to *less* accurate updating, rather than more.

Inaccurate updating might be of different kinds. One possibility is that a voter updates her beliefs in the correct direction, but not to the appropriate magnitude. That is, instead of becoming less positive by a factor of "X" in the face of negative information, the evaluation might become less positive by something less than "X," in effect a conservatism bias (Steenbergen, 2001). Such updating failures might have relatively limited consequences, as long as they are directionally accurate. On the other hand, attitude strengthening effects (also called polarization effects, Lord, Ross, & Lepper, 1979) where updating is in the wrong *direction* have been demonstrated. Redlawsk (2002) finds that voters with an existing positive evaluation of a candidate become more positive about that candidate when encountering negative information. Lodge and Taber (2000, 2005; Taber & Lodge, 2006) show a similar effect in studying issue preferences, and Edwards and Smith (1996) find that individuals confronted with an argument in conflict with their prior beliefs judge that argument to be weak, spend longer scrutinizing it, and generate a list of relevant thoughts and arguments that tend to refute the argument rather

than support it, a process consistent with attitude strengthening. In the context of a campaign, a voter learning something negative about a favorite candidate might first doubt the validity of the information, spend time reviewing and trying to comprehend it, and in the process create a list of relevant thoughts, most of which argue that the information is either false or unimportant. This thought listing, in refuting the new piece of information, could call to mind many of the reasons for the initial support of the candidate and leave a better feeling about the candidate even after encountering negative information.

In the end we consider that there may be two processes with contradictory effects. One, motivated reasoning, suggests that small amounts of incongruent information can be countered in the service of maintaining existing evaluations. The other, affective intelligence, might not come into play until there is a significant threat to expectations, driving up anxiety, and overcoming the motivation to maintain evaluations. It is the existence of this hybrid process that we test for in the remainder of this paper.

Hypotheses

We know that voters update their evaluations based on new information encountered during the campaign. We also know that affect influences this process in many ways. Voters with positive affect towards a candidate may want to hold on to that evaluation and resist changing their opinion.[2] The question here is whether there is a point at which the positive affect motivated reasoners try to maintain is overwhelmed by a growing threat to the existing evaluation as more incongruent information is encountered, and thus leading to more accurate updating.[3] And if there is such a point, what drives it? We know that updating is not a linear process, but it may be that it is also a hybrid process, with voters updating one way when encountering just a bit of unexpected (incongruent) information, but another way when the threat represented by the new information grows large. To test this empirically we must do several things. First, an initial candidate evaluation must be established—that is, the voter must have time to learn something about the candidates. Second, we must challenge the voter's positive evaluation of the

---

[2] It is also possible that a similar effect occurs with a disliked candidate—that is, that voters who develop negative affect towards a candidate may be unwilling to positively update their evaluations, at least at first. While the design of our study allows us to examine this possibility, for reasons of both space and theoretical clarity our focus in this paper will be on positively evaluated candidates.
[3] Another way to think of this in more Bayesian terms is to consider whether as more negative information is encountered, a voter's positive "priors" become less and less important in the calculation of the posterior evaluation. Gill (2007) notes in a different context that as the N of new information points goes to infinity, the "data" eventually win. In other words, is there a point at which the priors no longer exert influence on the calculation of a revised evaluation given new information? Of course, we would argue in the present case that the N need not go to infinity at all, but that there is a finite tipping point at which evaluation begins to update more accurately.
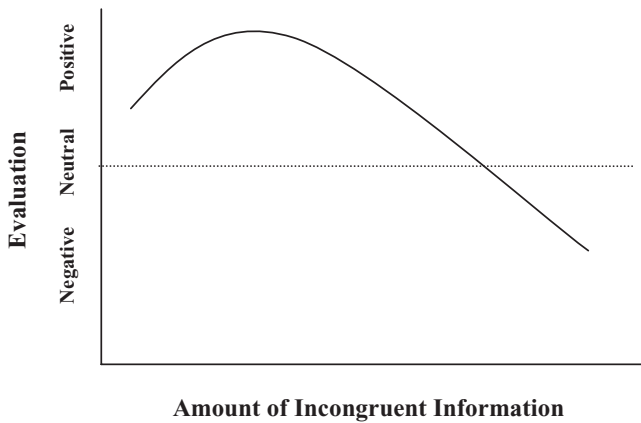
Figure 1. Expected Effects of the Amount of Incongruent Information on Evaluation of a Preferred Candidate.

preferred candidate by providing negative information about that candidate. Third, we must assess whether the updated evaluation shows evidence of motivated reasoning or some more accurate form of updating, and finally, if both processes are at work, we must find the point at which the impact of new negative information on evaluation changes.

In summary, we believe the updating process works something like this: (1) A voter develops an initial positive evaluation of a candidate through the early information that is learned; (2) if a small amount of negative information is encountered, rather than adjusting downward this initial evaluation becomes more positive, showing the motivated reasoning attitude strengthening effect; (3) if enough negative information is encountered to heighten the voter's anxiety about the preferred candidate, affective intelligence suggests that he or she will become more careful in processing additional new information; (4) This increased anxiety and careful processing may lead to an affective tipping point where additional negative information begins to generate downward adjustments to the evaluation. We suggest that a voter's evaluation of a positively evaluated candidate should follow a pattern such as the one seen in Figure 1 as greater amounts of negative information are encountered.[4]

Putting this into a more structured hypothesis, we expect that:

[4] It is certainly possible that a voter might never actually encounter negative information about a liked candidate. There is evidence that given a choice, voters may well work to confirm their evaluations by avoiding information that might challenge them (Taber & Lodge, 2006). This particular study does not address the issue of avoiding incongruent information; instead most of the subjects in the experiment to be described did in fact encounter such information, though a small number did not and thus provided a useful control group.

Hypothesis 1: Updating for positively evaluated candidates in the face of incongruent information will begin as a motivated reasoning process, showing attitude strengthening effects. Given enough incongruent information, an affective tipping point exists where increasingly anxious voters will begin to update candidate evaluations more accurately in the face of increasing incongruent information.

Hypothesis 1 suggests that if we find attitude strengthening effects for small amounts of incongruent information then these effects are the result of motivated reasoning. Of course, we cannot see motivated reasoning as it happens, only its results. But one means by which motivated reasoners might support their existing evaluations is by bolstering, that is, bringing to mind positive information already known to offset the new negative information that has been encountered. The results of such a process may then be visible in the memories people report about candidates after the election. In particular:

Hypothesis 2: Voters encountering small amounts of negative information about a liked candidate will report more *positive* memories about that candidate than will those encountering no negative information or those encountering large amounts of negative information.

Finally, increasing anxiety on the part of the voter who is learning increasingly negative information about a positively evaluated candidate may be the mechanism by which motivated reasoning is overcome and attitude strengthening ends. In following Marcus et al. (2000) on this point, we argue that at a high enough level of incongruency—unexpectedly negative information about a liked candidate—the environment in which these voters are operating will become more threatening, increasing anxiety resulting in more accurate updating. Another way of thinking about this is that voters should also become less certain that they made the "right" choice as their anxiety grows. Thus:

Hypothesis 3a: As increasing incongruency drives up anxiety about a positively evaluated candidate, voters will become less certain that they have made the right choice when called upon to cast a vote.

Direct evidence of greater anxiety as voters encounter more incongruency will also provide some evidence of this otherwise unseen process:

Hypothesis 3b: Encountering incongruent information generates anxiety that grows as more incongruency is encountered. Anxiety will increase until the voter adjusts to the new information and begins to consider other candidates.

Methodology

While a decision may be a single choice made at one *point* in time, evaluation is a process that occurs over some *period* of time. To understand a process, we should observe it as it occurs. Process-tracing experiments have been employed outside of political science by using information boards that allow subjects to choose exactly what they would like to learn about a set of alternatives presented to them (Ford, Schmitt, Schechtman, Hults, & Doherty, 1989; Jacoby, Jaccard, Kuss, Troutman, & Mazursky, 1987). Within political science similar information boards have been used to examine voting (Herstein, 1981), political decision making (Riggle & Johnson, 1996), and information search in political environments (Huang, 2000; Huang & Price, 1998) among other subjects. However, they have rarely been used to study candidate evaluation, though it would seem that process tracing could yield great insights into this subject.

The problem is that the traditional information board is *static* and allows constant access to all attributes for all alternatives under consideration. In the context of an election, this would be as if a voter had access to any piece of information about a candidate at any time, allowing easy comparison between candidates across all attributes. In a real election, however, information is much less organized, somewhat more chaotic, and the time allowed for learning and information gathering is limited by Election Day. Information comes and goes, and candidates do not always make it easy for voters to make a comparison or even get a clear understanding of where they stand on issues. Lau and Redlawsk's (Lau, 1995; Lau & Redlawsk, 2001, 2006) computer-based *dynamic* process-tracing methodology offers a way to model the vagaries of a political campaign in a controlled experimental environment. The system generates an ever-changing information environment that mimics the flow of information throughout a campaign and makes only limited amounts and types of information available at any point in time, much like in an election campaign in which many issues are "here today, gone tomorrow." It can also overwhelm voters with potentially unmanageable amounts of unorganized information in a way that resembles the media maelstrom in a real political environment. Yet, the dynamic information board retains the essential characteristic of process-tracing experiments in that it tracks the evaluation and decision-making process as it happens and as information is acquired.

We use this dynamic process-tracing environment to present a simulated presidential primary election campaign to subjects who learn about four candidates from within their party, evaluate them, and make a vote choice.[5] The

---

[5] A primary election was chosen to limit the direct effects of partisanship in this particular study. Obviously partisanship is of great import during a general election and undoubtedly plays a critical role in establishing candidate preference and possibly in resisting change to that preference. However, it adds a layer of individual difference that negatively impacts experimental control. Thus we settled on offering a primary election where partisanship would not be a factor for this study. If we can

campaign consists of a wide range of information about each candidate, including 27 issue positions, plus group endorsements, personality traits, and background information, along with preelection polls. As voters learn about the candidates, the system collects data on the information they access, how long they spend on each item, how they feel about each item, and their vote choices and evaluations of each candidate. These last two measures are obtained multiple times throughout the campaign in the form of "polls" in which voters are asked to choose a favorite candidate and rate all candidates. Given all the measures we are able to monitor, this methodology clearly provides an excellent way of tracking the evaluation and decision-making process during a campaign and the role that existing evaluative affect plays in the processing of new information and updating of candidate evaluations.

Experimental Design

A total of 207 nonstudent subjects were recruited from the Eastern Iowa area to participate in a mock presidential primary featuring four candidates from one party.[6] Candidates in the primary were fictional but designed to realistically represent the range of ideologies within their parties. Since the candidates were not real, subjects clearly had no prior knowledge about any of them, requiring evaluations to be determined only by the information accessed and inferences made during the campaign. Subjects registered as either Democrat or Republican before being exposed to information for the candidates from their party. Subjects were only allowed to vote in the party for which they had registered. Once the campaign began, subjects chose what they wished to learn about the candidates from an ever-changing set of candidate attributes presented over a 25-minute time period.

When subjects arrived they were seated at a computer and given an oral introduction to the experiment. They then completed an online questionnaire measuring their political preferences, knowledge, and interests. These questions allow us to gauge each subject's placement on the issues that were used in the simulation, which was necessary to be able to manipulate incongruency through subject-candidate agreement. Subjects were given a chance to practice with the dynamic process-tracing environment. They then began the primary campaign where they had 25 minutes to learn about the four candidates in their party, following which they voted for one of them. The campaign was interrupted after about seven minutes by a poll where subjects were asked to report their vote

---

establish the existence of a tipping point in the first place in this simpler environment, we can move on in future work to examine the conditions under which the tipping point itself varies. The strength of partisanship would clearly be one of those conditions worth closer examination.

[6] Subjects were recruited in a variety of ways to ensure some level of diversity, specifically in age and income. We do not claim that the subject pool is representative of any particular population. Subjects ranged in age from 18 to 88 years and had household incomes ranging from 7.5 to 100 thousand dollars per year. Fifty-six (56) percent of the subjects were female. Subjects who completed the study received $20 for their time.

preference and feeling thermometer evaluations of all four candidates. This poll was repeated two more times, once at about 13 minutes and again at about 20 minutes into the campaign. At the end of the campaign subjects voted and did one more set of evaluations. They were then asked a number of follow-up questions, including a memory-listing task asking subjects to record "everything you can remember" about each candidate, using only the candidate name as a prompt. Subjects were then prompted to indicate whether each memory made them feel enthusiastic, anxious, or angry about the candidate. Finally, subjects completed a cued recall task where they indicated whether they recalled examining each piece of information they had seen and if so, what they recalled about their affective response to it. They were then debriefed and dismissed.

The key experimental manipulation embedded in the election simulation varied the probability of encountering incongruent information during the campaign, and thus varied the information environment in which subjects operated. Incongruent information is defined as any candidate attribute at odds with the subject's preferences. For example, if a subject was pro-choice, an incongruent piece of information about a positively evaluated candidate would be that the candidate was pro-life. In this way a pro-choice subject learning her preferred candidate was pro-life would clearly have her expectations violated.[7] We varied the probability that subjects would encounter a certain amount of information like this; unexpectedly bad positions taken by a liked candidate.

As noted above, about seven minutes into the campaign subjects were polled and asked to indicate which candidate they would vote for if the election were held

---

[7] A second manipulation was also included, but is not central to the analyses here which focus on how individual pieces of information generate a sense of threat or anxiety as they accumulate. An "embedded instructions" manipulation was intended to alter the overall emotional state of the subjects just before the simulation began. Approximately half of the subjects were given special instructions about the experiment designed to heighten their overall sense of anxiety about their performance in the study. The instructions told subjects that their performance in the experiment was critical to the continuation of our research funding and that they were expected to do a good job. The other half of the subjects did not receive these instructions. An analysis of this intended manipulation showed that it was not strong enough to generate the expected reaction; as a consequence we do not consider it further here.

A third manipulation involved asking subjects how they felt about individual pieces of information that they accessed. One-half of subjects (immediate affect group) were asked immediately after viewing each item whether or not it made them feel enthusiastic, anxious, and/or angry toward the candidate. They were also asked to recall this affective response at the end of the study during a cued recall process. The other half (post affect group) were only asked to recall their affect at the end of the study, well after the simulation had been completed. This manipulation was designed to test whether or not subjects could accurately recall the affect attached to information they learned during the election when they are asked about it later. An initial examination of the data (Civettini & Redlawsk, 2005) suggests that affect recall is problematic at best. This manipulation is not directly relevant to our discussion of the affective tipping point since we do not rely here on subjects' own assessment of affect. However, because asking immediate affective responses impacted the number of items that could be examined by taking up time, we use this manipulation as an instrumental variable to predict the amount of information examined for each liked candidate in Table 2 below, in order to control for its effects.

at that point.[8] Subjects also evaluated each candidate on a 0–100 feeling thermometer, providing a ranking of candidate preferences. Following this first poll, subjects were randomly assigned to one of five levels of incongruent information. Those in Group 0 viewed candidates who were assigned issue positions that remained ideologically consistent throughout the campaign (i.e., the most liberal Democrat always taking the most liberal position or the moderate Democrat always taking more conservative positions held by his party). For this group candidate attributes were not manipulated in any way. In Group 1, a random 10% of information made available to the subject was manipulated to be incongruent with the subject's own preferences, while 90% of the information remained consistent with the candidate's established ideology. Group 2 subjects were assigned a 20% probability that available information would be incongruent with the subject's preferences with 80% remaining ideologically consistent, and Groups 3 and 4 were assigned 40% and 80% incongruency, respectively.

The assignment of incongruency occurred without the subjects' knowledge. In choosing what information to view about candidates, subjects could not know before choosing an item whether it would be congruent or incongruent with their initial candidate evaluation and therefore they could not control their information environment or the amount of incongruency they encountered.[9]

Data

Process-tracing methodologies provide data that are both extensive and complex. Before we could begin analyzing our results several steps were needed to clean the dataset. To begin with, we dropped eight subjects who either failed to complete the study or who did not take the study seriously. We then removed 10 additional subjects who looked at fewer than 50 pieces of information (less than two per minute over the course of the election) or more than 200 pieces of information (more than eight per minute). While retaining these 10 subjects does not change either the significance or substance of our results, these subjects were clear outliers when we examined the distribution of the number of items accessed across all subjects. We were left with 189 subjects whose data were suitable for analysis.[10]

It is important to make clear that subjects themselves decided when to click on information headlines to learn the details about any given candidate. Thus the

[8] By the first poll, the average subject had looked at 15–20 pieces of information, which were generally evenly spread across the four candidates in his or her party.

[9] Subjects clicked on boxes that contained "headlines" stating what information was available. These headlines were generally unsourced and valence-free. Examples include "Singer's position in Iraq" and "Rodgers' political philosophy."

[10] After debriefing the experimenter coded the degree of seriousness with which the subject approached the study, so this measure is based on observation of the subject's demeanor at the time of the study. Subjects examining fewer than 50 or more than 200 items were clear outliers. The mean number of items examined was 111.43, with a standard deviation of 35.07. Removing these subjects from the dataset does not substantively change the results and maintains consistency with Civettini and Redlawsk (2009).

actual effect of the manipulation depends to some extent on the headlines that subjects chose to examine. Within the experimental groups the actual amount of incongruent information that was encountered was distributed around the assigned probability based on two factors. First, because subjects chose from the available headlines what they wanted to know, and because we were randomly assigning the proportion of available information that could be incongruent, we could not control exactly how much incongruent information subjects actually learned. Second, because unmanipulated information was held ideologically consistent for the candidate based on his initial position on the liberal-conservative spectrum, subjects who themselves were not as ideologically consistent as the candidates would have encountered information that was incongruent relative to their own preferences but which was not specifically manipulated.

Ultimately for each subject there is a probability of x that any single item of information about a liked candidate was incongruent with the subject's preference and a probably of 1-x that the item was at a set point along the liberal-conservative spectrum based on the candidate's assigned ideology, where x is the manipulated probability: 0, 10, 20, 40, or 80%. If a subject's initially preferred candidate was the most liberal, we would expect that the subject would herself generally prefer the most liberal positions. However, if the subject held some more moderate preferences in her mix then she would still have the chance of encountering candidate positions incongruent with her own preferences from among the items we did not manipulate. The more internally inconsistent a subject's own issue preferences the more likely that the subject would encounter incongruent information. In the end, though, our interest in this study is to examine the extent to which encountering a more or less incongruent information environment results in more or less or accurate updating of candidate evaluations; we are not examining effects of the specific items themselves.

As the dynamic process-tracing experiment progressed, subjects clicked on valence and source neutral headlines that appeared and disappeared from the screen; each named an attribute that could be learned for a candidate. By clicking on the headline, subjects could learn detailed issue positions, candidate traits, and group endorsements. For each of these items that a subject accessed we recorded whether it was congruent with the subject's evaluation of the candidate. For traits this was fairly simple, since the information hidden behind the headline was clearly positive or clearly negative—for example, "Martin is considered egotistical and difficult to work with" would certainly be seen as negative, while "Even Martin's opponents consider him an honorable man" would be viewed as a positive trait.[11] Endorsements could also be readily determined to carry positive or negative valence, based on the subject's own expressed group preferences (as

---

[11] An initial set of trait statements was independently evaluated by three research assistants, who coded them as positive or negative. Only statements that were agreed to be positive or negative by all three coders were actually used in the study.

measured in the preexperiment questionnaire). If a disliked group endorsed a liked candidate, this was incongruent, while a liked group endorsing the candidate would be clearly congruent.

Coding positions as congruent or incongruent was a little more complicated. For each issue in the campaign, eight different positions were created which could then be assigned to any of the candidates on the fly as the simulation progressed. Once an issue position was assigned to a candidate that candidate consistently took that position and the position was not available to any other candidate. We were then able to calculate the distance between the subject and the candidate on a standard 7-point liberal-conservative scale for each individual issue. Where a subject was randomly assigned to receive an incongruent issue item for a liked candidate, the computer attempted to assign the most distant available issue position to that candidate.[12]

As described earlier, subjects in Group 0 where candidate issue positions were not manipulated could still encounter incongruent information since congruency was determined by subject-candidate issue agreement. Many voters have inconsistent ideologies, more liberal or conservative on some issues than others, so these subjects were bound to disagree with even their favorite candidate at times. And of course in the four experimental groups where congruency was manipulated, some unmanipulated items might still have been incongruent with a subject's own preferences for the same reason. Thus our calculation of the amount of incongruency subjects encountered is a combination of the proportion of *manipulated* items that were incongruent *plus* incongruence encountered because of a subject's own ideological inconsistency. No matter what the cause of incongruence, it operated in the same way—subjects would learn unexpectedly negative information about a liked candidate. To establish the congruency of unmanipulated items, we manually coded items as congruent or incongruent based on the actual candidate-subject issue distance. Issue positions closer than 3.5 points on the liberal-conservative scale were coded as congruent and issues 3.5 or more points away were coded as incongruent, though tests of other cut points make no significance difference in the results. Finally, to calculate the percentage of incongruent information each subject encountered over the entire campaign, we divided the total number of incongruent items the subject examined by the overall total number of items that the subject examined, resulting in a measure of total incongruency for each

---

[12] The issue position ratings were obtained by giving the list of positions to several graduate students and faculty members in the University of Iowa's Department of Political Science. Each person coded all the items and the final rating was the average rating across all coders. The subject's placement was self-reported during the questionnaire at the beginning of the study.

During the experiment before the first poll was administered and the manipulation began, positions were assigned to candidates as the subject selected pieces of information to read. Since these positions could not be changed after the subject had viewed them, the fewer options were available for the computer to choose from when executing the manipulation. In effect, the strength of the manipulation was attenuated when only a less distant position was available to assign to the candidate.

subject.[13] This measure is a continuous variable, ranging from 0 to 100 percent, with a mean of 37.9% and a standard deviation of 30.5.

Results

We have two goals in these analyses. First, we examine our process-tracing data and candidate evaluations to determine whether the patterns of evaluation match our theorized process. That is, as more negative information is encountered about a liked candidate do we see a pattern of initial attitude strengthening, which reaches some tipping point after which evaluations update in a negative direction? Second, if the patterns our subjects exhibit match our theory, we must then examine whether we have evidence to support our argument that both motivated reasoning and affective intelligence processes are at work.

The Interaction of Affect and Candidate Evaluation

We turn first to candidate evaluation.[14] The first part of Hypothesis 1 suggests an attitude-strengthening effect in the face of a small amount of negative information about a liked candidate. Voters who process information according to normative standards should become less positive about a liked candidate with each piece of negative information they learn. But motivated reasoners will try to counter the negativity to maintain their existing evaluation and may in the course of doing so become even more positive about a liked candidate. We begin by simply looking graphically at the actual evaluations for each of our randomly assigned groups. Figure 2 presents this information graphing the mean evaluation of the initially preferred candidate from the first poll (before any manipulation of information) through the second and third preelection polls, and the final postvote evaluation (poll 4). The evaluations were made on a 0–100 point feeling thermometer scale.

It is immediately clear that updating of evaluations does not proceed linearly across the groups. The group assigned to no incongruency (that is, where all information remained ideologically consistent for the candidate) shows a small uptick in the evaluation of the liked candidate by the end of the campaign. This is what we would expect since these subjects learned little or nothing that should

---

[13] Nearly every item that was available to subjects could be assessed for congruency with a few exceptions. In addition to issues, endorsements, and candidate personality traits, subjects could learn poll results throughout the campaign and could also learn about candidate background and experience. These items were not assessed as either congruent or incongruent and are ignored in the analysis.

[14] While we manipulated both the highest and lower rated candidates, we focus in this analysis only on the candidate most preferred (highest rated) in the first poll, after which the manipulation of congruency began. We would expect some similar effects (though reversed) for a disliked candidate; however, it may also simply be that voters ignore a disliked candidate once they establish the evaluation. Available space precludes us from this examination here.
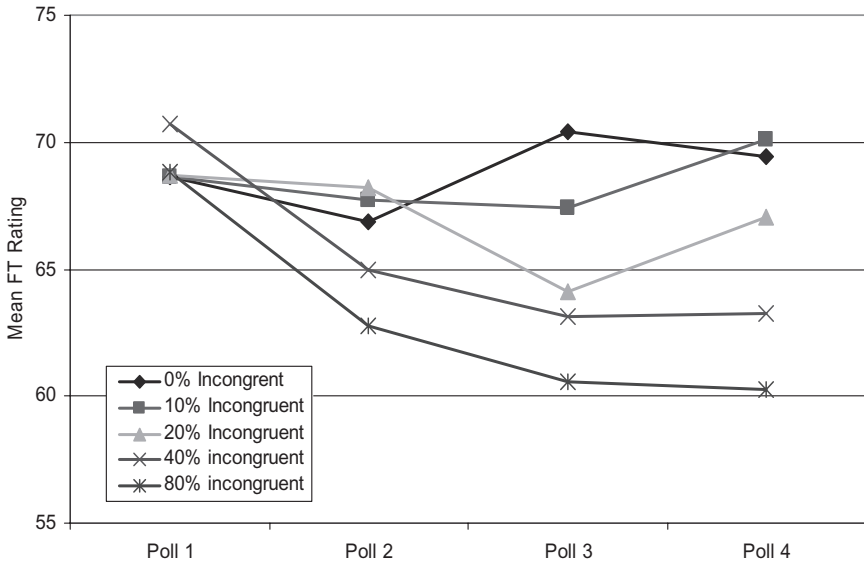
Figure 2. Evaluation of Preferred Candidate by Assigned Incongruency Groups Over the Course of the Campaign.

have made them feel negatively about the candidate. But interestingly, so does the group that was assigned to 10% incongruency. Even those assigned to 20% negative information about the initially most preferred candidate do not, in the end, show significant updating of their evaluations of that candidate. It is only within the groups assigned to encounter 40% and 80% incongruent information that we see the expected decline in the evaluation over time.

But this does not give us the full picture of the evaluation as a factor of the *actual* level of incongruent information encountered, since as detailed above, the randomly assigned groups provided a probability of encountering incongruent information, rather than a certainty. The exact level of incongruency was controlled by both the manipulation and the subject's own level of ideological constraint. The more constrained the subject the more likely that unmanipulated information would be perceived as congruent with expectations and manipulated information would be perceived as incongruent. We recoded our subjects into incongruency quartiles, after first removing those subjects who demonstratively never encountered any incongruency (based on our calculation of actual incongruency as described earlier). This gives us five groups with differing incongruency means, encountering from zero to 74.6% incongruent (negative) information about the initially most preferred candidate. Table 1 describes these observed incongruency groups and compares them to the original randomly assigned experimental groups.

Table 1. Incongruency Encountered by Observed and Randomly Assigned Groups

Target		Observed			Randomly Assigned		
		Mean	SD	N	Mean	SD	N
Group 0	0	0.0	0.0	18	15.5	24.9	37
Group 1	10	20.2	18.3	51	24.2	24.4	37
Group 2	20	30.2	17.7	35	34.3	23.8	34
Group 3	40	40.5	19.3	43	39.3	18.2	43
Group 4	80	79.2	19.7	42	74.6	23.8	38

Note. Means in percentages of Incongruent Information.

When we graph the mean evaluations over time for each of these observed groups, the differences seen in Figure 2 become even more pronounced. Now the group that never actually encountered any incongruent information (Group 0) actually ends up somewhat less positive about their favorite candidate at the end than either of the first two quartiles of incongruency (Groups 1 and 2). And those in the first quartile—averaging about 20% incongruent information—become consistently more positive about their candidate, even in the face of a nonnegligible amount of negative information. This finding clearly fits patterns previously shown by Taber and Lodge (2006) as well as Redlawsk (2002) with attitude strengthening effects evident in the face of incongruent information.

But it is also clear in Figure 3 that this strengthening effect does not occur for all levels of incongruent information. While subjects who encountered smaller amounts of incongruency seem to have their positive impression strengthened over time, those encountering above the median amount of incongruency (Groups 3 and 4) show a very different pattern. These subjects begin lowering their evaluation of the preferred candidate immediately, and that evaluation continues to decline over time, though leveling off towards the end of the campaign. These subjects, then, do show evidence of more accurate updating of their priors compared to those encountering less incongruency.

We can calculate the change in evaluation from the first poll to the vote for each of the randomly assigned and observed groups and test whether the changes are statistically different from zero and from each other. We use a one-way ANOVA to predict the mean evaluation for each incongruency group.[15] We present the results graphically in Figure 4, which includes both the random and observed groups. The evaluation updating process appears to correspond with the theorized

---

[15] A *post hoc* LSD analysis of the significance of the difference between each group shows in both cases (manipulated and observed groups) that the change in evaluation from first poll to the vote for the two groups with the most incongruent information is statistically different from the other three though not from each other. Likewise, the change in evaluations for the three groups with the least incongruency, while different from the other two, are not statistically different from each other.
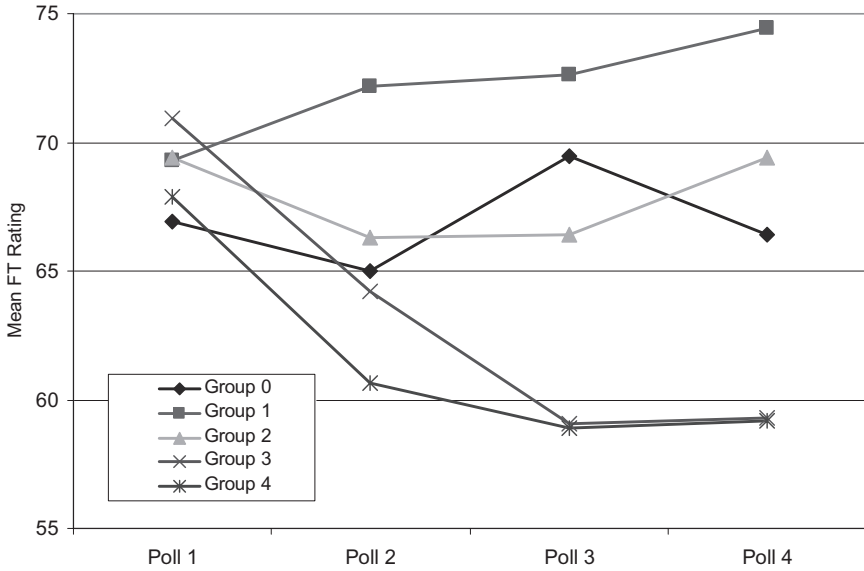
Figure 3. Evaluation of Preferred Candidate by Observed Incongruency Groups Over the Course of the Campaign.
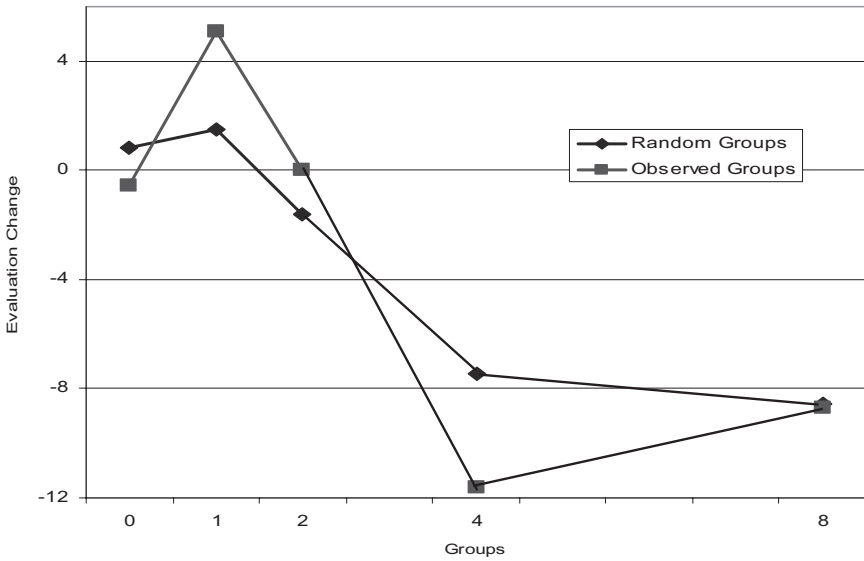

Figure 4. Mean Evaluation Change from First Poll to the Vote.

curve presented in Figure 1.[16] While the differences between the three groups with the lower levels of incongruency are attenuated in the curve for the randomly assigned groups, when we account for the information actually encountered through our observed groups, the pattern is clearer. But in neither case does the evaluation updating curve come close to approximating a normatively correct process where evaluations consistently decline as more negative information is encountered.

Calculating the Updating Curve and Tipping Point

The graphical presentation of the random and observed group data shows clear differences in updating by those who encountered small amounts of incongruency compared to those who encountered greater amounts, suggesting our first hypothesis finds support. Low incongruency subjects appear to resist negatively updating the evaluation of a preferred candidate while those encountering larger amounts of incongruency update as we would expect. We now proceed to calculating the nature of the evaluation updating curve and examining whether in fact a tipping point can be identified. Note that Figure 4—change in evaluation over time—suggests something more complex than the quadratic function described in Figure 1.

Recall that our observed groups—designed to account for the actual amount of incongruent information encountered by subjects—deviate from random assignment, raising the serious question of whether the amount of incongruency is endogenous to our experimental design and limiting the causal claims that we can make. One way to address this problem is to use an instrumental variable approach. Instead of using the randomly assigned groups or our artificially constructed groups based on observed information search, we will employ a two-stage least-squares regression with instrumental variables, the first stage of which will use the randomly assigned groups as instruments for the observed incongruency levels.

An initial fit of the data without any control variables (not shown) suggests that a cubic function works better than a quadratic, which is not surprising given the pattern shown in Figure 4. Our complete model allows for this cubic function and controls for the initial rating of the liked candidate and amount of information encountered about the candidate. We control for the initial rating to address ceiling effects. Subjects who started with a very high rating of their preferred candidate have less room to move up than those who started lower. We control for the amount of information viewed because incongruency is measured in terms of the percentage of information encountered. Subjects varied in how much information they examined for their liked candidate; controlling for this is necessary to recognize

---

[16] Note that the x-axis scale in Figure 4 describes the relative amounts of incongruency for each group compared to the other groups,

Table 2. Two-Stage Least Squares Regression for Change in
Evaluation for Initially Preferred Candidate by Incongruency
Levels: Second Stage Regression Results

Change in Evaluation			
	B	SE	Sig.
Observed Incongruency Level			
Linear[1]	1.207	.712	.092
Quadratic[1]	−.055	.026	.036
Cubic[1]	.0005	.0002	.033
Total Items Examined for Candidate[1]	.156	.509	.760
Initial Evaluation of Candidate	−.064	.134	.635
Model significance: F-test = 2.668, p = .024.			

Table entries are un-standardized coefficients and Standard Errors.
Significance reported are z-tests, two-tailed. N = 189.
[1]Variable was estimated by instruments in first stage equation.
First stage instruments were 1) the randomly assigned target levels
for incongruency, the square of the target levels, and the cube of
the target levels and 2) a dummy variable indicating the presence
or absence of the immediate affect manipulation.

that some subjects simply learned more information than did others. We use the randomly assigned targets for information incongruency as instruments in the first stage to predict the actual amount of incongruent information encountered. Further, because one of our manipulations slowed half of the subjects down by asking their affective response after every item examined (see Note 7) those subjects systematically examined fewer items. Thus we employ a dummy for this manipulation as an instrument to account for the actual variation in total items examined for liked candidates.

Table 2 shows the results for the second-stage equation. The dependent variable is *change* in the initially preferred candidate's evaluation at the end of the election compared to the first poll. The coefficients on the linear, quadratic, and cubic terms for information incongruency are all significant and in the expected directions. The coefficients on initial candidate evaluation and total number of items examined for the liked candidate are not significant, though they are in the expected direction. Importantly, the results of this analysis confirm the results of our initial examination of the data using only the randomly assigned groups or our observed groups. The curve resulting from the equation shows an initial increase in the final evaluation for candidates when a small amount of incongruent information is encountered. At some tipping point, the evaluations begin to decline and updating proceeds in a more normatively correct direction, as more negative information is encountered about the preferred candidate. For subjects encountering relatively little incongruency, the predicted change in evaluations is positive over the course of the campaign, that is, encountering small amounts of negative

information does in fact result in a more positive evaluation. Likewise, at relatively large amounts of incongruency, evaluations no longer continue to decline, resulting in the cubic function as best fitting. This counterintuitive result most likely occurs because at high levels of incongruency subjects turn away from the initially preferred candidate and begin examining more information for other candidates.

Where does the tipping point occur? That is, how much incongruent information is necessary to force updating to begin to take "reality" into account? Calculating this tipping point is simple, since it is the local maximum of the cubic function within the bounds of 0–100% incongruency. This simple calculation yields a rounded value of 13.4, suggesting that in these data once about 13% of information about an initially preferred candidate is incongruent with a subject's own preferences, evaluations stop becoming *more* positive. But evaluations of candidates do not actually become more negative than for an ideal candidate (no incongruency) until about 28% incongruent information is encountered. Thus in our data there is a range of incongruency (I), $0 < I < 28$, between which evaluations of an initially liked candidate are on average higher than for the ideal candidate, that is, one who takes positions perfectly congruent with a subject's own preferences. But the exact tipping point itself is less important than the fact that our results strongly support both motivated reasoning effects and accurate updating, at different levels of incongruency.

Tipping Point Mechanisms: Motivated Reasoning and Affective Intelligence

Having established the presence of a tipping point in candidate evaluation, we turn now to an explanation for this observed behavior. Recall that we are working with two different affective theories. Motivated reasoning suggests that evaluators will become more positive in the face of negative information about a liked person, as existing positive affect for that person interacts with negative affect for new information, triggering a (cognitive) effort to make sense of the new information while striving to maintain the existing (affective) feeling about the person. This process may lead to attitude strengthening. Somewhat in opposition to this, affective intelligence theory argues that a threatening environment generates increasing anxiety, which operates to cause a person to learn more about the environment in order to prepare a response. Thus increasing anxiety leads to learning, which normatively should lead to more, rather than less, accurate updating of evaluations. It is our contention that the updating pattern we have established can be explained by the operation of both motivated reasoning (at low levels of threat; that is, small amounts of incongruency) and affective intelligence at higher threat levels. We turn now to examining whether the pattern we see can be accounted for by these two affective processes.

Motivated Reasoning

It is well established that information that does not violate expectations is more readily processed than information that does. Encountering unexpected information piques our interest and forces us to concentrate more on it, compared to information that simply confirms expectations. Multiple explanations for this phenomenon have been advanced, including affective intelligence's dual affective systems (Marcus et al., 2000), Petty and Cacioppo's (1981, 1983) central versus peripheral routes, and hot cognition's interaction of existing affect and new information (Taber & Lodge, 2006; Redlawsk, 2002). Regardless of the mechanism that might explain information processing differences, the first step here is to show that such differences actually exist in our data. If they do not, then we can immediately discount any process—including motivated reasoning and affective intelligence—that would expect incongruency to take longer to process.

We examined the amount of time subjects took to read incongruent information for their preferred candidate and compared it to the time they took to read congruent information for the same candidate. Initially we simply compared mean processing time for each type of information across our subjects. The results show that processing time for incongruent information was more than 10% greater than for congruent information ($M_{incongruent} = 6.99$, $M_{congruent} = 6.35$ seconds, t = 2.187, 1797 df, $p < .03$.)[17] However, different information items in the dynamic process-tracing system may be of different lengths. Obviously the longer the item, the longer it will take to read. So the analysis must also control for the number of words in the item and the individual subject's reading ability, measured as the time it took to read a set of instructions. A regression analysis controlling for these factors is shown in Table 3. The result is clear. Even after controls are applied, congruency affects processing time. On average, incongruent items take more than half a second longer to process, all else equal. To put this into perspective, we calculated the length of time to read a 30-word item for someone at the mean reading speed. On average such an item took 8.56 seconds to read if it was congruent, but 9.20 seconds to read if it was incongruent, all else equal, an increase of 7.4% in processing time.

The candidate evaluation pattern supports our argument about the form of the updating process, with an initial increase in evaluations in the face of a small amount of negative information, followed by more accurate updating as incongruency builds. And our processing time analysis suggests that incongruent information takes longer to process, as we would expect. But are we seeing motivated reasoning as such? And when subjects do reach a tipping point do increasing levels of anxiety correspond to the downturn in evaluations as we would expect from affective intelligence?

---

[17] While we are not examining disliked candidates in this paper, we see the same effect of processing time for those candidates as well. Congruent information (negative information about a disliked candidate) is processed much faster (6.3 seconds compared to 7.1 seconds for incongruent). The t-test is significant at $p < .02$. These results replicate the findings in Redlawsk (2002) with a different dataset.

Table 3. Processing Time for Congruent and Incongruent
Information Preferred Candidate

| Predictor	B	SE
Item Incongruency (1 = Congruent)	−.637***	.228
# Words in Item	.154***	.006
Reading Speed	−.401***	.012
Constant	4.257***	.343
Adj r^2	.487	

*p < .1, **p < .05, ***p < .01
Table entries are un-standardized OLS coefficients and Standard
Errors. N = 1587 Items.

As Hypothesis 2 suggests, we can look to the memories our voters report as an indicator of the presence of motivated reasoning. Motivated reasoners attempt to maintain their affective evaluations in the face of unexpected information. One way in which they do this is to use bolstering—to recall into active memory factors which support the existing evaluation and which may then overwhelm the new incongruent information.[18] If this happens, memory for these attitude-supporting attributes may be enhanced, as repeated access to a concept increases the likelihood that the concept will be remembered (Fiske & Taylor, 1991). We can examine the likelihood that our subjects recalled positive (enthusiastic) memories as a function of the amount of incongruency encountered. If we find that small amounts of negative information increase positive memory, we will have evidence of a motivated reasoning process.

Figure 5 plots the mean number of reported memories for each of the instrumentally defined incongruency groups. The results almost perfectly support the expectations of Hypothesis 2. Subjects who saw a small amount of negative information about their most liked candidate (Group 1) report more positive memories and more overall memory than any other group, including those who encountered the least amount of incongruent information (Group 0). Further, those in Group 2 report about the same number of positive memories as Group 0. It is not until levels of incongruency beyond the tipping point that positive memories for the initially liked candidate begin to decline. Thus we have evidence of a motivated reasoning process that operates as would be expected if bolstering were under way at low levels of incongruency.

[18] Let us be clear about this process. We do not suggest that this is cognitively driven at the start. Bolstering itself is a by-product of the associative nature of memory (Anderson, 1983) and hot cognition (Lodge & Taber, 2005). In activating the construct stored for the candidate (and the affect associated with it) other connected memory nodes are also activated. Since we are dealing with positively evaluated candidates, most of these associations are also positive. As memory nodes are activated they are more likely to be recalled when later tested. The net result is that more positive memory traces should be evident when negative information stimulates processing since incongruent information is processed more carefully than congruent information (Redlawsk, 2002).
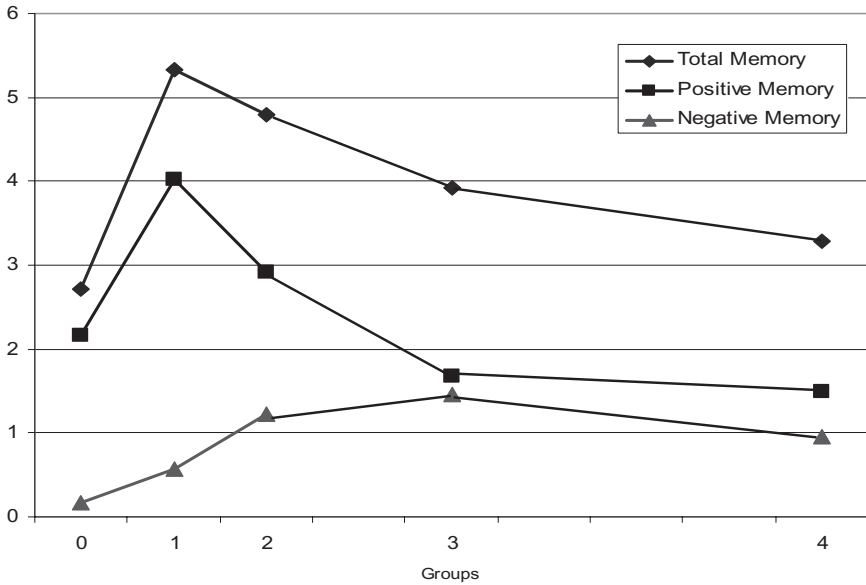
Figure 5. Incongruency and Memory for a Preferred Candidate.

Affective Intelligence

What about once evaluations begin to adjust downward in the face of incon-gruent information? What causes this tipping point? Hypotheses 3a and 3b test the claim of affective intelligence that more anxious voters are better voters (Marcus et al., 2000). If negative affect increases as incongruency increases, it may well be a mechanism that leads to the more accurate updating of evaluations past the tipping point. Recall that affective intelligence argues in favor of two emotional subsystems, one of which is characterized by positive affect and the other by negative. This negative affective subsystem—the behavioral inhibition system (BIS)—is responsible for focusing attention on negative stimuli so as to avoid (or address) potentially dangerous situations. Our expectation then is that as the information environment for our subjects becomes increasingly at odds with expectations about a liked candidate, feelings of negative affect—a sense of anxiety about the original evaluation—will grow.

In order to assess whether affective intelligence processes are at work, we have three measures of the impact of the primary election on our voters. The first of these measures is derived from asking subjects after they had voted to indicate how difficult it was to make a choice between the candidates. The second comes from asking how confident subjects were that they had chosen the "right" candidate. Both were asked on a simple scale from 1 to 5 coded so 1 was
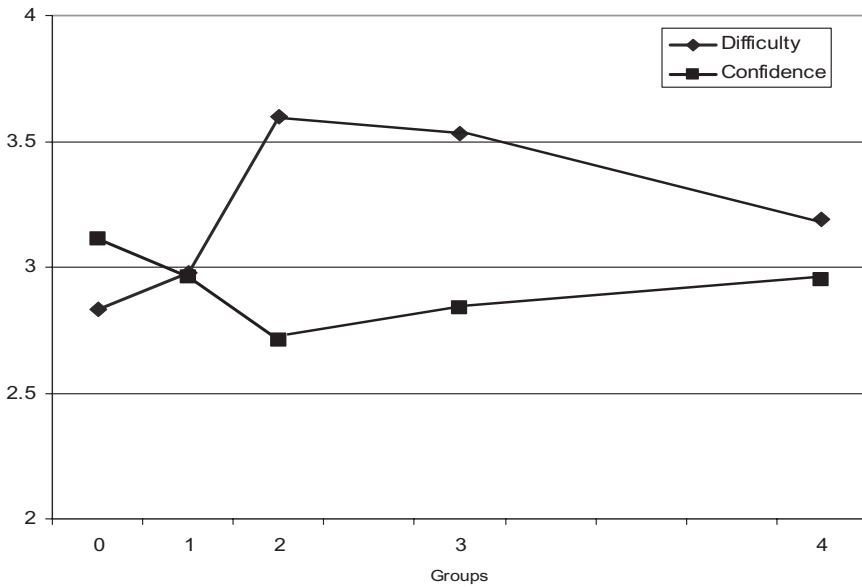
Figure 6. Reported Difficulty and Confidence in Primary Vote Decision by Incongruency Groups.

"Very Easy" or "Not at all Confident" and 5 was "Very Difficult" or "Very Confident."

Figure 6 presents a summary of both the difficulty and confidence measures by incongruency group. Both point in the same direction and support Hypothesis 3a. As levels of incongruent information increase so does reported difficulty while confidence decreases, but only up to a point. Subjects encountering incongruency about their most liked candidate that extends beyond the tipping point actually report an easier decision and greater confidence at the conclusion of the campaign than those at lower levels of incongruency. This occurs because once past the evaluative tipping point subjects make peace with the fact that their initially most liked candidate was just not what they thought he was and they moved on to another option. In fact, voters in the highest incongruency group were less than 20% likely to vote for the candidate they initially preferred, well below all other groups.

Our third and most direct measure of the impact of the information environment on subjects comes from the administration of Watson, Clark, and Tellegen's (1988) Positive and Negative Affect Scale (PANAS), which allows us to assess subjects' affective states. This measure is explicitly aimed at assessing how a subject feels at the time the scale is administered. A brief questionnaire was administered at two points in time—before the experiment began, and again immediately before the cued recall process. We can compare the level of negative
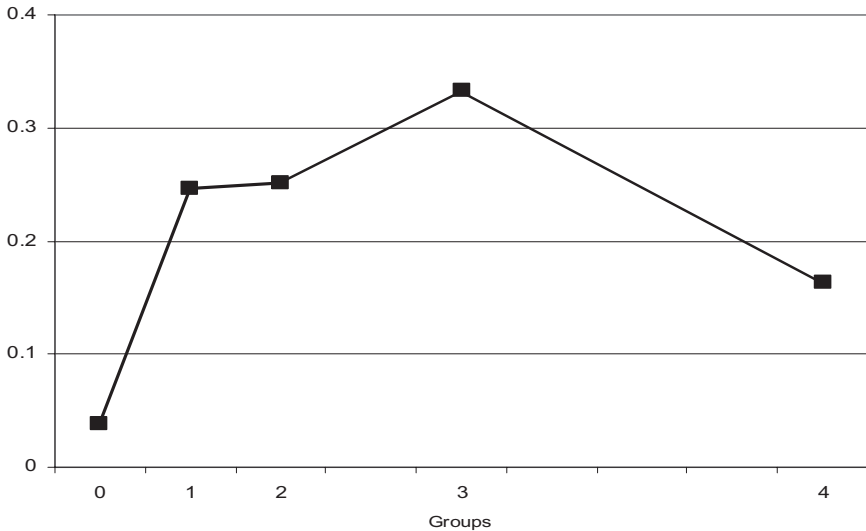
Figure 7. Change in Negative Affect at End of Campaign by Levels of Incongruency.

affect expressed by our subjects after they voted to the level they expressed before they started and examine the extent to which change in negative affect varies by the amount of incongruency subjects encountered. We specifically use the PANAS Negative Affect scale (NA), which consists of ten affect words.[19] For each, subjects were asked to indicate how strongly the affect word described their current feelings at the time of administration on a 1–5 scale, with 1 indicating that the word did not describe them at all and 5 indicating it described their current feelings very strongly. The NA scale in our data has a Cronbach's Alpha of .884, indicating high reliability. Previous studies have shown the NA scale itself to be a good measure of anxious states (Crawford & Henry, 2004; Mehrabian, 1997), and it is widely used in social psychology. We compare the postexperiment NA score with the preexperiment NA score and examine the net difference by the instrumentally defined levels of incongruency in the information environment. Results are in Figure 7.

Just as with reported difficulty and confidence, this more direct measure of anxiety responds as we predict in Hypothesis 3b. The group of subjects who encountered the least incongruency (Group 0) evidences no change at all in their negative affective state post experiment compared to preexperiment. Subjects encountering greater amounts of incongruent information about their liked candidate report a more negative affective state immediately following the election. As

---

[19] These ten words are: distressed, upset, guilty, scared, hostile, irritable, ashamed, nervous, jittery, and afraid.

levels of incongruency climb so too does negative affect. But at the highest levels of incongruency, the increase in negative affect, while still greater than those encountering no incongruency, drops substantially. Once past the affective tipping point subjects become less anxious than they do immediately before it, though the effect is not instantaneous. Given that subjects at the tipping point are significantly more positive about their initially liked candidate—despite his growing negatives—it takes more negative information to bring the rating down below the initial point before attitude strengthening began. Thus too it appears to take time for the sense of anxiety evidenced to dissipate.

As negative affect reached high enough levels, our subjects began to reassess their initial evaluations. While we cannot with these data establish without question the causal link we would like, we do see patterns we would expect if affective intelligence was at work. As more negative information about a liked candidate is encountered, the information environment becomes more threatening, leading to increasing negative feelings and a sense that the decision is more difficult. This increasing challenge is resolved at the tipping point, as subjects stop the process that leads to attitude strengthening and instead begin a process leading to accurate updating.

Discussion

Motivated reasoners strive to maintain existing evaluative affect, even in the face of countervailing information. This effect has been well established in the literature (Kunda, 1990) and replicated here. Thus where we would expect "rational" voters to approximate normatively correct updating, motivated reasoners show evidence of attitude strengthening, becoming even more positive about a liked candidate in the face of negative information about that candidate. And while the voters in our experiment show heightened negative affect and evidence that the choice becomes more difficult as incongruency grows, they also show attitude strengthening effects as motivated reasoning predicts. But we also show clear evidence that these effects do not necessary continue under all circumstances. At some point our voters appear to wise up, recognize that they are possibly wrong, and begin making adjustments. In short, they begin to act as rational updating processes would require. Why? Our results are consistent with the idea that as anxiety increases (leading to more difficulty in the decision and less confidence) voters pay closer attention to the environment (and processing time increases). They then begin to more carefully consider new information and potentially override existing affective expectations. Such a process would be consistent with affective intelligence overriding motivated reasoning.

It is worth keeping in mind the limitations of this study which represents only a first attempt to show that motivated reasoning does not continue ad infinitum and to propose a mechanism—affective intelligence—that might interfere with it. The study is experimental, the environment in which our voters operated, while having

many of the key features of a real world election environment, certainly was not a real campaign. Experimental control meant our candidates were made up, and while our subjects told us the candidates seemed quite real, no subject knew anything about them before the study. Further, we limited ourselves to a primary election, thus negating any effects of partisan identification. We would expect, for example, that strong partisans in a general election would have a very high tipping point, at least compared to non-partisans. Likewise, ideologues, as opposed to moderates, might also be harder to move off their initial support for a candidate very close to them. We do not test either of these possibilities here. Yet while our environment is not a real election, we believe that the psychological processes our subjects engage in their attempts to learn about and evaluate candidates are really no different whether in the laboratory or in the midst of a "real" campaign. Voters learn about candidates, assess whether the new information fits with their expectations, and in some fashion ultimately create and update evaluations in both environments.

We argue that the role of affect in candidate evaluation is not a matter of either/or. Voters are not either motivated reasoners or rational processors. Instead, as our experiment demonstrates, voters can be both depending on the information environment in which they are operating. When the amount of incongruency is relatively small, the heightened negative affect does not necessarily override the motivation to maintain support for a candidate in which the voter is already positively invested. But as V. O. Key (1966) noted four decades ago, voters are not always fools. An affective tipping point exists at which existing positive evaluations give way to a newly understood reality—the candidate is just not what he or she seemed to be at first.

If voters are, in fact, somewhat immune to small amounts of negative information about their favored candidates, what are the implications in the real political world? Should candidates not worry about minor mess-ups, flip-flops, and fleeting scandals, or is the atmosphere of the modern campaign already so negative that most voters are pushed way past the tipping point months before Election Day? A closer look at the modern campaigning environment might help answer that question, but the fact remains that for a while at least, a candidate's early supporters will probably resist attempts to change their minds. Candidates who need to win new voters without alienating their bases should be able to lean to the middle, as long as they don't lean too far. However, in a real campaign where prior beliefs about candidates are long standing and based on much more information than subjects in our study were exposed to early on, reaching the affective tipping point will require substantially more negative information than our subjects encountered. It is easy to imagine a long-time fan of a presidential candidate rejecting virtually all new negative information about him or her and sticking to an early evaluation. Yet even such a fan might, in the face of overwhelming information counter to expectations, awake to the changed reality and revise her beliefs accordingly.

ACKNOWLEDGMENTS

REFERENCES

Abelson, R. P. (1963). Computer simulation of "hot cognitions." In S. Tomkins & S. Messick (Eds.), *Computer simulation and personality: Frontier of psychological theory* (pp. 277–298). New York: Wiley.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Brader, T. (2005). Striking a responsive chord: How political ads motivate and persuade voters by appealing to emotions. *American Journal of Political Science, 49*(2), 388–405.

Civettini, A. J. W., & Redlawsk, D. P. (2005). A feeling person's game: Affect and voter information processing and learning in a campaign. Paper presented at the Annual Meeting of the American Political Science Association, Washington, DC.

Civettini, A. J. W., & Redlawsk, D. P. (2009). Voters, emotions, and memory. *Political Psychology, 30*, 125–151.

Crawford, J. R., & Henry, J. D. (2004). The positive and negative affect schedule (PANAS): Construct validity, measurement properties and normative data in a large non-clinical sample. *British Journal of Clinical Psychology, 43*(Pt. 3), 245–265.

Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. New York: Harcourt.

Edwards, K., & Smith, E. (1996). A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology, 7*(1), 5–24.

Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd edn.). New York: McGraw Hill.

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University.

Ford, J. K., Schmitt, N., Schechtman, S. L., Hults, B. M., & Doherty, M. L. (1989). Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes, 43*, 75–117.

Gill, J. (2007). *Bayesian methods: A social and behavioral sciences approach* (2nd Ed). Boca Raton, FL: CRC Press.

Gladwell, M. (2000). *The tipping point: How little things can make a big difference*. New York: Little, Brown, and Company.

Green, D., & Gerber, A. (1999). Misperceptions about perceptual bias. *Annual Review of Political Science, 2*, 189–210.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the task is memory-based or on-line. *Psychological Review, 93*, 258–268.

Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.

Herstein, J. A. (1981). Keeping the voter's limits in mind: A cognitive process analysis of decision making in voting. *Journal of Personality and Social Psychology, 40*, 843–861.

Holbrook, R. A. (2005). Candidates for elected office and the causes of political anxiety: Differentiating between threat and novelty in candidate messages. Paper presented at the annual meeting of the Midwest Political Science Association.

Holbrook, A. L., Krosnick, J. A., Visser, P. S., Gardner, W. L., & Cacioppo, J. T. (2001). Attitudes toward presidential candidates and political parties: Initial optimism, inertial first impressions, and a focus on flaws. *American Journal of Political Science, 45*(4), 930–950.

Huang, L., & Price, V. (1998). Motivations, information search, and memory structure about political candidates. Paper presented at ICA '98.

Huang, L. (2000). Examining candidate information search processes: The impact of processing goals and sophistication. *Journal of Communication, 50*(Winter), 93–114.

Isen, A. M. (2000). Positive affect and decision making. In M. Lewis & J. Haviland-Jones (Eds.). *Handbook of emotions* (2nd ed., pp. 417–435). New York: Guilford.

Jacoby, J., Jaccard, J., Kuss, A., Troutman, T., & Mazursky, D. (1987). New directions in behavioral process research: Implications for social psychology. *Journal of Experimental Social Psychology, 23*, 146–175.

Key, V. O., Jr. (1966). *The responsible electorate*. Cambridge, MA: Belknap Press.

Kunda, Z. (1990). The case for motivated political reasoning. *Psychological Bulletin, 108*(3), 480–498.

Lau, R. R. (1995). Information search during an election campaign: Introducing a process tracing methodology for political scientists. In M. Lodge & K. McGraw (Eds.), *Political judgment: Structure and process* (pp. 179–206). Ann Arbor: University of Michigan Press.

Lau, R. R., Anderson, D., & Redlawsk, D. P. (2008). An exploration of correct voting in recent U.S. presidential elections. *American Journal of Political Science 52*(2), 395–411.

Lau, R. R., & Redlawsk, D. P. (1997). Voting correctly. *American Political Science Review, 91*(3), 585-598.

Lau, R. R., & Redlawsk, D. P. (2001). An experimental study of information search, memory, and decision making during a political campaign. In J. Kuklinski (Ed.), *Citizens and politics: Perspective from political psychology* (pp. 136–159). New York: Cambridge University Press.

Lau, R. R., & Redlawsk, D. P. (2006). *How voters decide: Information processing in an election campaign*. New York: Cambridge University Press.

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. 182–213). London: Cambridge University Press.

Lodge, M., & Taber, C. S. (2005). The primacy of affect for political candidates, groups, and issues: An experimental test of the hot cognition hypothesis. *Political Psychology, 26*, 455–482.

Lodge, M., McGraw, K., & Stroh, P. (1989). An impression-driven model of candidate evaluation. *American Political Science Review, 83*, 399–419.

Lodge, M., Steenbergen, M., & Brau, S. (1995). The responsive voter: Campaign information and the dynamics of candidate evaluation. *American Political Science Review, 89*, 309–326.

Lord, C. G., Ross, L., & Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*(11), 2098–2109.

Marcus, G. E., & MacKuen, M. (1993). Anxiety, enthusiasm, and the vote: The emotional underpinnings of learning and involvement during presidential campaigns. *American Political Science Review, 87*(September), 672–685.

Marcus, G. E., Newman, W. R., & MacKuen, M. (2000). *Affective intelligence and political judgment*. Chicago: University of Chicago Press.

Mehrabian, A. (1997). Comparison of the PAD and PANAS as models for describing emotions and for differentiating anxiety from depression. *Journal of Psychopathology and Behavioral Assessment, 19*(4), 331–357.

Petty, R. E., & Cacioppo, J. T. (1981). *Attitudes and persuasion: Classic and contemporary approaches*. Dubuque, IA: William C. Brown.

Petty, R. E., & Cacioppo, J. T. (1983). Central and peripheral routes to persuasion: Application to advertising. In L. Percy & A. Woodside (Eds.), *Advertising and consumer psychology* (pp. 3–23). Lexington, MA: Lexington Books.

Redlawsk, D. P. (2001). You must remember this: A test of the on-line voting model. *Journal of Politics, 63*, 29–58.

Redlawsk, D. P. (2002). Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making. *Journal of Politics, 64*, 1021–1044.

Redlawsk, D. P. (2006). Motivated reasoning, affect, and the role of memory in voter decision-making. In D. P. Redlawsk (Ed.), *Feeling politics: Emotion in information processing* (pp. 87–108). New York: Palgrave-Macmillan.

Redlawsk, D. P. (2007). Understanding vs. prediction in candidate evaluation. Paper presented at the annual meeting of the Midwest Political Science Association and at the annual meeting of the International Society of Political Psychology, Portland, OR.

Redlawsk, D. P., Civettini, A. J., & Lau, R. R. (2007). Affective intelligence and voting information processing and learning in a campaign. In A. Crigler, M. MacKuen, G. E. Marcus, & W. R. Neuman (Eds.), *The affect effect: Dynamics of emotion in political thinking and behavior*. Chicago: University of Chicago Press.

Riggle, E. D. B., & Johnson, M. M. S. (1996). Age differences in political decision making: Strategies for evaluating political candidates. *Political Behavior, 18*(1), 99–118.

Steenbergen, M. (2001). The Reverend Bayes meets J.Q. Public: Patterns of political belief updating in citizens. Paper presented at the annual meeting of the International Society of Political Psychology. Cuernavaca, Mexico.

Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science, 50*(3), 755–769

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect. *Journal of Personality and Social Psychology, 54*(6), 1063–1070.

Zajonc, R. B. (1984). On the primacy of affect. *American Psychologist, 39*(2), 117–123.