

**Detecting Change in Professional Conduct Using Information
from the Web: A Differential Effect for Different Business Entity
Types, Implications for Privacy, and “the Right to be Forgotten”**

Roozmehr Safi

University of Missouri–Kansas City
safir@umkc.edu

Detecting Change in Professional Conduct Using Information from the Web: A Differential Effect for Different Business Entity Types, Implications for Privacy, and “the Right to be Forgotten”

Abstract:

Increasing reliance on the Internet’s perpetual memory has raised concerns regarding how dated information that would otherwise be forgotten or inaccessible can unduly or disproportionately influence current assessments and decisions. I investigate aspects of this topic for two major business entity types: one-person businesses (i.e., sole proprietors) and firms. Results show that one-person businesses tend to be more severely impacted than firms by past adverse information, and furthermore their improvement trends over time are more likely to be dismissed as noise than recognized as true signals of change. While firms can offset old unfavorable conduct by engaging in new favorable behaviors, a sole proprietor’s current favorable operations can remain dominated by decades-old actions. Results also indicate that decision makers perceive firms as more capable of truly changing. Also, while only decision makers with certain personality characteristics recognize signs of positive change from a sole proprietor, all decision makers detect and appreciate such changes in a firm’s conduct. This study finds that limiting access to adverse past information is likely to be more helpful (or necessary) for one-person businesses or more generally for individuals than for firms.

Keywords: right to be forgotten, privacy, reputation management, time-value of information, implicit theories.

Detecting Change in Professional Conduct Using Information from the Web: A Differential Effect for Different Business Entity Types, Implications for Privacy, and “the Right to be Forgotten”

1- INTRODUCTION

Massive volumes of information are placed on the Web every day. This information is often retained indefinitely as part of the Internet’s omnipresent “memory” that can be effortlessly accessed with a simple search. Although this is justly considered a breakthrough of the information age, like any other technological development it can also create unintended societal, economic, or ethical dilemmas. One case in point is the way past information that might otherwise be forgotten or inaccessible remains “fresh” on the internet, possibly improperly influencing present judgements. This concern has resulted in demands and legislation limiting access to certain past information, or “a right to be forgotten” (Jones, 2018; Mayer-Schönberger, 2009). This right would allow individuals and businesses to request that search engines such as Google exclude specified past information from search results.

Information available on the Internet now has consequential effects on nearly all decisions, particularly in professional and business contexts. Statistics show that a majority of clients now routinely use the Internet to research service providers before hiring them (Hayes, 2018). Given the prevalence of such “cyber-vetting” practices, it is no wonder that following the passage of the right to be forgotten as law in Europe the largest number of delisting requests have involved professional information (Bertram et al., 2018). The concept of limiting access to past information has become the center of many debates: while some see it as a virtuous privacy right allowing entities to become “different from their pasts,” others see it as threat to the free flow of information and citizens’ right to know. Although the idea has received substantial

attention from philosophical, legal, and implementation perspectives, certain fundamental questions have remained unexplored. At the core of a desire to restrict access to certain past information is the implicit assumption that decision makers' opinions are disproportionately affected by dated information, even when the focal entity has changed for the better over time and therefore necessitating an external protective mechanism (i.e., the right to be forgotten). At its face value this assumption may appear sensible, but there is surprisingly little research empirically investigating either how this assumption holds or its boundary conditions. The goal of this study is to address this knowledge gap with respect to decisions made within business contexts.

In this study I investigate how favorable and unfavorable information from either the near or distant past as obtained from a typical cyber-screening activity can shape decision makers' judgments concerning a business. I also investigate the role of decision makers' personality characteristics during this process. In answering these questions I make a distinction between two major businesses entity types: one-person businesses (i.e., sole proprietors, or unincorporated businesses such as independent contractors and freelancers), and firms. From a theoretical perspective this distinction is important because sole proprietors are perceived to be more "human" as extensions of their individual owners. From a practical perspective these entities constitute the majority of businesses (nearly 75% in the United States).

The remainder of this paper is structured as follows. First, I provide a background on the topic. Next, I establish a set of research hypotheses drawing on past research and observations. I investigate these hypotheses using two studies, present the results, and discuss their implications for the theory and practice of both Information Systems and privacy rights.

2- BACKGROUND

From language learning and skill development to managing social interactions, humans depend on memory. It is therefore unsurprising that forgetting is generally considered undesirable. Yet forgetting is not necessarily a “flaw” of the human mind; it has important evolutionary advantages. For example, forgetting old or forsaken beliefs in favor of new ones helps avoid cognitive distress. Forgetting the past helps us focus on the present and plan for the future, as well as reestablish broken social ties. Forgetting is not only a beneficial coping mechanism at the individual or dyadic levels (Weiner 1968), but can also be advantageous at a societal scale. Over time the collective memory of society “forgets” its members’ misdeeds, providing a second chance to be included and contribute.

Forgetting can be beneficial since time is a major determinant of information’s value (Mayer-Schönberger, 2009). Information - even if originally true and accurate - tends to render out of context, inaccurate, or otherwise less diagnostic over time. However, the omnipresent and shared memory of the Internet has disrupted this natural role of time (Ambrose 2012; Korenhof et al. 2015), making recovery from past errors increasingly difficult (Ambrose et al., 2012; Jones, 2018). Removing information from the Internet is not easy since the stakeholders of this information are many and diverse, each with potentially competing interests in the retention or erasure of data. This context has given rise to demands toward provisioning a formal mechanism helping those adversely affected by past information with a right to limit access, or a “right to be forgotten.”

Limiting access to certain past information is not a new concept but has roots in many legal philosophies¹. The major event bringing this concept to the fore was a 2014 landmark ruling by the court of the European Union in favor of a citizen requesting that Google remove links to a digitized newspaper article concerning prior delinquent debt that he had subsequently paid back. The grounds for this ruling was that even when information was initially lawful and accurate, it should later be removed at the request of those affected when it could, “appear to be inadequate, irrelevant or no longer relevant, or excessive ... in the light of the time that has elapsed”². During the first three years following this ruling Google received approximately 2.4 million delisting requests in Europe and approved the removal of nearly one million links, most involving professional information or wrongdoings (Bertram et al. 2018). This has clearly significant economic and societal implications. As public support for the idea increases (approximately 75% approval in the US; Auxier, 2020) and many countries consider adopting comparable laws (Frosio, 2016), greater understanding of this topic is a pressing need.

Prior Research

The right to be forgotten is closely related to privacy, a topic widely studied in Information Systems and other disciplines (e.g., Aivazpour & Rao, 2020; Bélanger & Crossler, 2011; Lowry et al., 2017; Smith et al., 2011). Nonetheless, while privacy amounts to limiting the *release* of certain information, a right to be forgotten concerns limiting access to information *already released* to the public. Research on the right to be forgotten has grown rapidly since the early 2010s, although discussion has mostly remained limited to its philosophical, ethical, legal,

¹ For example, in England certain criminal convictions are considered “spent” after a set time and the ex-offender has no obligation to disclose related information for purposes such as employment.

² “Judgment of the Court in Case C-131/12 (Costeja),” available at <http://curia.europa.eu/juris/document>.

and implementation aspects (Shim 2016). Proponents from the ethical and philosophical perspectives argue that this right is necessary for the normal and just functioning of society. Legal researchers similarly argue that the right is a logical extension of existing privacy laws. On the other hand, the right has been seen as a form of censorship as well as a threat to freedom of the press and speech. Some of the biggest controversies regarding the right involve its implementation since the conditions required to qualify information for delisting (e.g., irrelevancy and excessiveness) are highly open to interpretation. The technical aspects of implementing the right to be forgotten are equally challenging; some have gone as far as to say that the right is “empty talk” by arguing that reversing the information dissemination process is simply not technically possible (Xiao & Lin, 2019). Overall, the literature shows little consensus on how to either properly interpret the concept or reliably enforce it as a law.

Unanswered Questions

While prior research has investigated various philosophical, legal, and technical aspects of the right to be forgotten, key empirical questions concerning the very need for such a right have remained underexplored. The *raison d’être* for limiting access to past adverse information is allowing individuals to live, “without being perpetually or periodically stigmatized as a consequence of a specific action performed in the past” (Mantelero, 2013, p. 231). This concern is based on the implicit assumption that dated information does in fact substantially affect decision makers’ opinions even when the focal entity has changed for the better over time. Nonetheless, there is little empirical research either on the degree to which this assumption holds or the conditions under which it holds (e.g., the type of entity under consideration or decision makers’ characteristics). The goal of this study is answering these questions.

3- THEORY AND HYPOTHESES

Humans form impressions about others based on small pieces of information and “thin slices” of observed behavior (Ambady et al. 2000). These impressions are updated over time as new information moderates or “overwrites” old information in memory (O’Brien and Klein 2017). A negative impression can therefore later change as the focal entity engages in newer positive behaviors. However, both old and new information are equally and vividly accessible when evaluating others based on information from the Internet. The right to be forgotten therefore concerns fears that searchers may not intuitively discount the decision weight of older and less relevant information, or may fail to “see” trends of positive change.

In investigating this concern one notes that researching a subject on the Internet typically results in a collection of disparate mixed (i.e., both positive and negative) information. Research shows that when presented with a collection of evidence people tend to spontaneously use these pieces of information to devise a holistic causal story (Pennington & Hastie, 1988). Extant work also shows that people often rely on the temporal order of evidence in order to infer causality (Block & Zakay, 2001). People are therefore likely to pay attention to information time tags in order to uncover trends of change (Favere-Marchesi, 2006). It is therefore hypothesized in the context of evaluating businesses based on Internet search results that:

Hypothesis 1 (H1): Decision makers use information timestamps and successfully detect trends of positive or negative conduct changes when evaluating a business based on a collection of mixed information from the Web.

The next hypothesis concerns the possible effects of the order in which information items are displayed. Although a higher Internet search result ranking is supposed to reflect greater importance, Internet search results appear for the most part at random and certainly without

respect to time. In general the order of encountering information can influence decision outcomes via a “primacy effect,” an opposite “recency effect,” or more generally an “order effect” (see Hogarth & Einhorn, 1992). Nonetheless, when information items are timestamped they can also reveal temporal trends (Favere-Marchesi, 2006). The lack of any temporal order, a common characteristic of information found via Web searches, has important implications for evaluations since it can conceal patterns of change that may otherwise be readily noticeable. It is hypothesized that:

Hypothesis 2 (H2): The lack of a chronological order in search results hinders decision makers from identifying conduct change trends when evaluating a business based on a collection of mixed information from the Web.

The next question is whether there is a differential effect in detecting positive versus negative conduct changes when evaluating a business based on evidence from the Web. Prior research indicates that individuals tend to hold asymmetric views concerning change for better versus for worse by perceiving the former to be less plausible (O’Brien & Klein, 2017). Social scientists explain this phenomenon using the conceptually parallel Law of Entropy from physics holding that all systems have a natural tendency to progress toward disorderliness and chaos. Such asymmetrical views regarding the likelihood of changing for better versus for worse are predicted to lead decision makers into taking signs of gradual decline in behavior as true signals of deterioration, but dismiss comparable signs of improvement as random noise. This perspective accordingly places those who improve over time at a disadvantage. It is hypothesized that:

Hypothesis 3 (H3): Decision makers identify a negative change trend more readily than they identify a comparable positive change trend when evaluating a business entity based on a collection of mixed evidence from the Web.

While previous hypotheses concern decision makers' overall evaluations of an entity based on a *collection* of information, the next hypothesis concerns the perceived decision value for individual pieces of positive or negative information from various points of time. Prior research shows that unfavorable information is generally perceived as more important than favorable information (i.e., people have a “negativity bias”; Baumeister et al. 2001; Rozin and Royzman 2001). Apart from the potency of negative versus positive information per se regardless of time, research has shown that negative information from the far past is minimally discounted while positive information from the far past is significantly discounted. This effect is termed the, “differential discounting of past information” (Brandimarte et al. 2018). It can therefore be predicted that information concerning old unfavorable business conduct is only discounted marginally to the extent that such information continues to overshadow current favorable conduct. It is therefore hypothesized that:

Hypothesis 4 (H4): Decision makers perceive unfavorable information about business conduct as more important than positive information regardless of the age of information.

The next hypothesis concerns the role of decision makers' personality characteristics in the ability to detect changes in others when using a collection of past information. People hold stable subjective views (a.k.a. “lay theories”) regarding the degree to which others can change and grow. While some individuals believe in the fixed nature of personalities, others believe in the malleability of character. The former group, referred to as “entity theorists,” attributes behavior to stable and unchanging traits. They accordingly tend to predict consistent behaviors by others across time and circumstances. The latter group, referred to as “incremental theorists,” tend to believe that conduct quality is dynamic and changing with respect to time and conditions.

These general views are referred to as, “implicit theories” (Chiu, Hong, et al., 1997; Dweck et al., 1995). Such views are extendible from individuals to various aspects of the social and material world (Kruglanski, 1989), including firms. These views are perhaps the most important personality characteristics relating to assessing others’ present based on their history of past behaviors, and accordingly the right to be forgotten. Although the potential implications of implicit theories for the right to be forgotten have been probed to some degree (see for example Berkelaar, 2017; Brandimarte, Vosgerau, & Acquisti, 2018; S. L. Williams, 2015), they have not yet been thoroughly investigated. However, it is predicted that implicit theories play a significant role in detecting positive changes in business conduct over time. It is hypothesized that:

Hypothesis 5 (H5): Incremental theorists detect and appreciate signs of improvement in business conduct more than entity theorists.

The above hypotheses are postulated with respect to businesses in general. However, in testing it is useful to make a distinction between two major types of business entities: owner-operated businesses and firms. Owner-operated businesses function as extensions of a single person and are therefore viewed by decision makers as possessing human-like traits, capabilities, and limitations. Although firms are also frequently viewed as having persistent qualities (i.e., some form of predisposed “corporate DNA”; Baskin 1998), they are typically perceived as more liable to change over time. Furthermore, while firms may be seen as capable of changing or firing managers, owner-operated businesses are not. In addition to investigating the above hypotheses for both entity types I postulate the following hypothesis:

Hypothesis 6 (H6): Decision makers perceive firms as more liable to change and grow versus owner-operated businesses.

4. EXPERIMENT

I investigated the above hypotheses by conducting an experiment. Participants in this experiment were presented with a cover story concerning a public institution in the process of selecting a contractor that would renovate its facilities. Participants were told that the institution's selection policy requires that shortlisted candidates be evaluated by various stakeholders, and were told that their help was needed to use personal judgment in anonymously evaluating a shortlisted candidate. This scenario was used in two studies: in Study 1 the candidate was described as a person, while in Study 2 it was described as a firm.

Participants made their evaluations based on a collection of information they were told had been collected about the candidate from the Internet. This collection was comprised of three favorable, three unfavorable, and two neutral information items (see Appendix A-1 for the procedure involved in developing these items as well their analysis in terms of potency and valance). All participants were presented with the same collection of information. However, under different experimental conditions the information items were paired with different timestamps. Some of the timestamps were two to three decades old and some were almost current. The experiment had a two-by-two factorial design. First, the collection reflected either an overall improvement or overall deterioration in candidate's conduct over time depending on how timestamps were paired with information items (i.e., negative items paired with old timestamps and positive items paired with more current timestamps or vice-versa). The second manipulation was regarding the order of information items' listing: in some conditions the items were arranged in (reverse) chronological order, while in others the items were presented out of temporal order. This manipulation made the pattern of change more versus less readily apparent, with the latter condition usually being the case with results from Web searches. A control

condition was also included in order to capture the baseline. The time tags in this condition were assigned such that the collection conveyed no trend and the items were presented in (reverse) temporal order. Table 1 summarizes the five conditions. Appendix A-2 provides details of timestamps and presentation orders under each condition in Studies 1 and 2.

Condition	Details	Notes
1	Favorable information items were assigned timestamps from recent past and unfavorable ones from distant past; items presented in (reverse) chronological order.	Reflects a pattern of improving conduct over time (<i>improving, ordered</i>).
2	Unfavorable information items were assigned timestamps from recent past and favorable ones from distant past; items presented in (reverse) chronological order.	Reflects a pattern of deteriorating conduct over time (<i>deteriorating, ordered</i>).
3	Similar to Condition 1, other than items were presented in random chronological order.	Reflects a pattern of improving conduct over time, but the pattern may not be readily noticeable (<i>improving, unordered</i>).
4	Similar to Condition 2, other than items were presented in random chronological order.	Reflects a pattern of deteriorating conduct over time, but the pattern may not be readily noticeable (<i>deteriorating, unordered</i>).
5 (Baseline)	Some favorable information was assigned timestamps from recent past and some from distant past, <i>and</i> some unfavorable information was assigned timestamps from recent past and some from distant past; presented in (reverse) chronological order.	Reflects neither an improving nor a deteriorating pattern of change in conduct over time (<i>as control</i>).

Table 1. Experimental conditions with regards to recent versus past information valance and presentation order.

Outcome Variables:

The first outcome variable was the overall rating of the candidate's appropriateness for the job based on the collection of information. Participants were also asked about the reasoning for their evaluations in an open-ended question. Next, they were asked to rate each of the eight timestamped information items in the collection separately in terms of importance. Participants were then asked to self-report - in retrospect - the degree to which they thought the age of information items (timestamps) affected evaluations. Participants also completed a scale capturing their implicit theory orientations, as well as answered demographic questions and other general questions concerning Internet use.

5-PARTICIPANTS AND RESULTS

Data was collected from subjects on Amazon's Mechanical Turk platform (MTurk). This population is considerably diverse in terms of age, education, and other demographic characteristics and provides rich life experiences with respect to this study. No specific participation criteria was set other than being from the United States. In terms of familiarity with the task, nearly 90% of participants said that they had hired a service provider (e.g., a repair person or a tax preparer) for their personal or professional lives, and 85% of these noted that they researched service providers on the Internet before hiring them at least half the time. Data was screened on collection in order to exclude either incomplete responses or responses from subjects who had clearly not engaged with the task (determined by answers to two check questions and the overall relevance of responses to open-ended questions - such screening procedures are recommended by prior research using MTurk; Sheehan, 2018). This resulted in 338 responses for Study 1 in which the candidate was described as a person (mean age: 38.5 years, median age: 36

years, females 51%), and 333 responses for Study 2 in which the candidate was described as a firm (mean age: 38 years, median age: 35 years, females 55%).

5.1. Overall Candidate Evaluation

The first outcome variable was participants' overall evaluations of the candidate based on the collection of eight information items measured on a seven-point Likert scale (1: Extremely Negative, 7: Extremely Positive)¹. Table 2 provides summary results.

Condition Entity Type	1 (improving, ordered)	2 (deteriorating, ordered)	3 (improving, unordered)	4 (deteriorating, unordered)	Baseline (no trend)
Owner-Operated Business (a person) – Study 1	3.61 (sd = 1.83) (n=69)	1.93 (sd = 1.31) (n=69)	2.88 (sd = 1.44) (n=64)	2.11 (sd = 1.24) (n=71)	2.49 (sd = 1.21) (n=65)
Firm – Study 2	3.93 (sd= 1.66) (n=70)	2.05 (sd= 1.40) (n=63)	3.71 (sd= 1.67) (n=66)	2.29 (sd= 1.40) (n=69)	2.89 (sd=1.56) (n=65)

Table 2. Mean evaluation (along with standard deviations and subject count) for owner-operated business (Study 1) versus a firm (Study 2).

Figure 1 depicts the results and features 95% confidence intervals.

¹ In addition to specifying their overall *evaluation* of the candidate, participants specified the degree to which they *recommended* that the applicant be awarded the job (1: Strongly Disagree, 7: Strongly Agree). In the following analysis using either response variable led to the same results, denoting the robustness of these findings. For space considerations details related to analysis of the second variable are not reported.

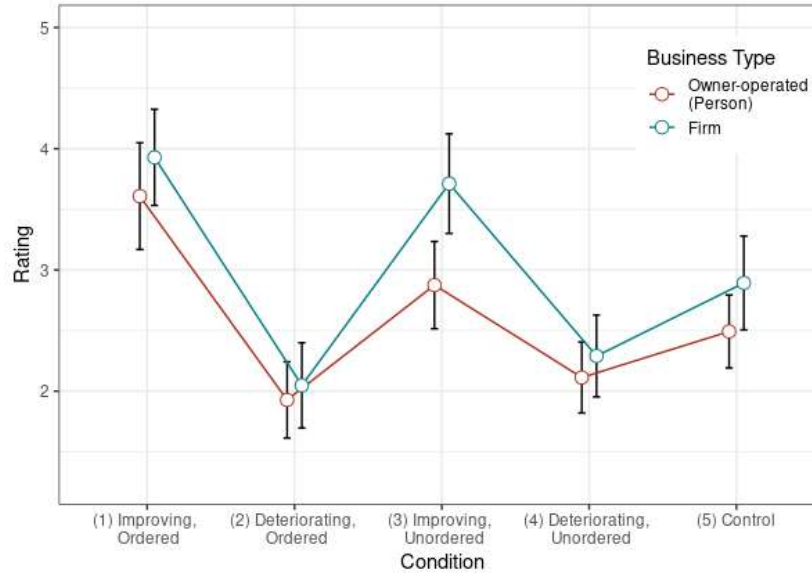


Figure 1. Mean evaluation of the candidate with 95% confidence intervals for Studies 1 and 2.

A one-way ANOVA comparing mean evaluation scores across conditions was statistically significant in Study 1 in which the candidate was described as a person ($F(4,333)=15.22$, $p < .001$), and so was the corresponding analysis for Study 2 in which the candidate was described as a firm ($F(4,328)=19.52$, $p < .001$). In Study 1 the mean rating in Condition 5 (baseline) was statistically different from mean rating in the treatment Condition 4. In Study 2 the mean rating in Condition 5 was statistically different from mean ratings in treatment Conditions 1, 2, and 3 from that study. This demonstrates a more pronounced separation between ratings in experimental conditions and the baseline when the candidate was described as a firm.

Participants in Groups 1 through 4 were subject to two simultaneous manipulations: (1) the direction of conduct change (change for better in Conditions 1 and 3 versus for worse in Conditions 2 and 4), and (2) information presentation order (chronologically ordered in Conditions 1 and 2 versus unordered in Conditions 3 and 4). I conducted a two-way ANOVA investigating the main effect for these two variables and their interactions on evaluations.

For Study 1 in which the candidate was introduced as a person the main effect from the direction of conduct change ($F(1, 269) = 47.73, p < 0.001$) as well as the interaction between the direction of conduct change and information presentation order ($F(1, 269) = 6.64, p = 0.01$) were statistically significant. The main effect of information presentation order ($F(1, 269) = 2.72, p = 0.10$) was not significant in this study¹. Panel A in Figure 2 depicts the results. A post-hoc analysis using Tukey correction for multiple comparisons showed that all pair-wise comparisons between the means are statically significant, except for the two conditions in which conduct has deteriorated over time (bottom left in Panel A). A similar analysis on results when the business is introduced as a firm (Study 2) shows that only the main effect of change direction ($F(1, 264) = 76.62, p < 0.001$) was significant. Neither the main effect of information presentation order ($F(1, 264) = 0.08, p = 0.78$) nor the interaction between these two terms ($F(1, 264) = 1.48, p = 0.23$) was significant. Figure 2, Panel B visualizes the results.

¹ I tested ANOVA assumptions by conducting Levin's test of homogeneity of variance and the result was significant, indicating that the assumption of variance equality in the dependent variable for different independent variable levels may be violated. This is not conclusive since the result from this hypothesis-based test is influenced by sample size (the larger the sample, the greater the power to reject the equal variance null hypothesis and vice-versa). For robustness I re-conducted the analysis by correcting for variance heteroscedasticity (i.e., "White-correcting") using the package "car" (Fox & Weisberg, 2019) in R (R Core Team, 2019). The same results in terms of the statistical significance for the predictor effects as those reported above were obtained.

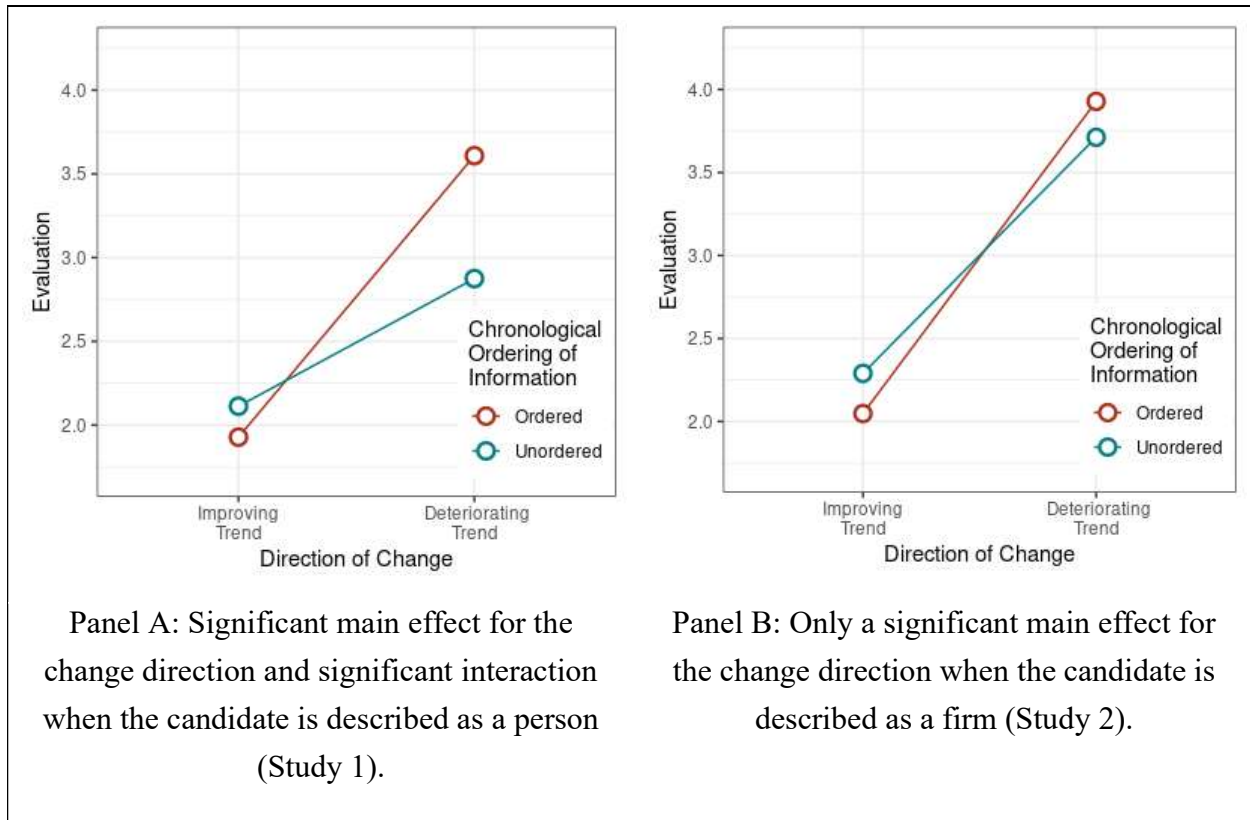


Figure 2. Effects of change direction and information presentation order.

The significant main effect for the change direction in Studies 1 and 2 indicates that overall subjects took information timestamps into account and rated the candidate higher when there was an improving conduct trend, supporting H1 for both entity types. However, this main effect was qualified by an interaction with presentation order in Study 1. Results show that the candidate was evaluated significantly lower following conduct improvement when the information items were presented out of chronological order versus when conduct similarly improved but the information items were shown in order. Results therefore support H2 for Study 1 by showing that a lack of chronological ordering in information items makes detecting conduct changes difficult. Unlike in Study 1, here the business was rated equally high when it had improved and equally low when it had deteriorated regardless of the order in which information was presented. This finds that H2 is not supported for Study 2.

For Study 1 in which the candidate was introduced as a person results also indicate that a positive trend of conduct change is not recognized as easily as a negative change trend (recognizing positive change trend is contingent on information being presented in order to make the change trend readily visible), supporting H3 in that Study. However, a positive change trend in conduct over time is recognized as easily as a negative trend regardless of whether or not information is presented in chronological order, failing to support H3 in Study 2.

A comparison of results for H2 and H3 across Studies 1 and 2 reveals a trend: while noticing improvement in persons is contingent on information items being presented in chronological order, improvement in firms is noticed regardless of how information is presented. I conjecture that this is due to the difference in the degree to which persons versus firms are deemed likely to and capable of changing, as probed further in the following sections.

5.2. Age Effect of Positive vs. Negative Information on Perceived Importance

After rating the candidate based on the collection of eight information items as discussed above, participants in Conditions 1 through 4 in both studies were asked to specify - in retrospect - how important they thought each individual piece of information was in forming overall opinions concerning the candidate. Participants rated each of the eight timestamped information items on a seven-point scale from “Extremely Unimportant” to “Extremely Important.” I then generated an aggregate measure of importance for favorable and unfavorable information respectively by averaging the scores of the three positive and three negative information items. It is notable that in Conditions 1 and 3 the favorable items had timestamps from the recent past and unfavorable ones had timestamps from the distant past, while for Conditions 2 and 4 the opposite was true. For this analysis responses from Conditions 1 and 3 and those from Conditions 2 and 4 were therefore aggregated. Table 3 provides the mean ratings for these two studies.

		Temporal Distance of Information Items			
		Study 1		Study 2	
		Recent Past	Distant Past	Recent Past	Distant Past
Information Item Valance	Favorable (+)	5.27 (1.16)	4.86 (1.26)	5.62 (1.00)	4.49 (1.35)
	Unfavorable (-)	6.49 (0.66)	5.92 (1.15)	6.28 (0.83)	5.88 (0.99)

Table 3. Mean (and SD) self-reported importance of items by information valance and temporal distance in rating the candidate described as a person (Study 1) versus a firm (Study 2).

I used a two-way repeated measure ANOVA in analyzing responses. In this analysis temporal distance served as a between-subject factor since every responder saw either favorable facts from the recent past and unfavorable facts from the distant past (an improving trend), or vice-versa (a deteriorating trend). The valance of the information items was a within-subject factor since every responder rated both favorable and unfavorable information. The dependent variable was responders' self-reported rating of information importance. For Study 1 in which the business was portrayed as a person there was a significant main effect for information valance ($F(1, 271) = 122.31, p < 0.001$) as well as a significant interaction effect between valance and temporal distance ($F(1, 271) = 21.87, p < 0.001$). The main effect of temporal distance was not significant ($F(1, 271) = .85, p = 0.36$); Figure 3, Panel A visualizes the results. A post-hoc analysis using Tukey correction showed that the observed mean differences among all four conditions are statistically significant.

For Study 2 in which the business was portrayed as a firm results showed a significant main effect for both information valance ($F(1, 266) = 102.9, p < 0.001$) and temporal distance

($F(1, 266) = 20.07, p < 0.001$). The interaction term was also significant ($F(1, 266) = 57.6, p < 0.001$). A post-hoc analysis using Tukey correction shows that the mean differences among all four conditions are statistically significant, except for favorable information that is recent and unfavorable information from the distant past (two points connected by the blue line in Figure 3, Panel B).




Figure 3. Simultaneous effect of information items' valance and temporal distance on self-reported importance.

The results from Study 1 show that for persons unfavorable information has a larger effect than favorable information regardless of how old or new the information is, supporting H4. Specifically, decision makers consider unfavorable conduct from as distant as two to three decades ago more important than nearly current favorable conduct. On the other hand, for firms (Study 2) negative information does not indiscriminately dominate positive information, failing

to support for H4. For firms positive information from the recent past is considered on par with negative information from the distant past, indicating that firms may be able to make up for past negative deeds by engaging in positive ones.

5.3. Comparing Results from Studies 1 and 2; Implications for Perceived Changeability

Certain observations stand out when comparing results for H1 through H4 concerning persons versus firms. First, while the order of presenting information items (chronological versus unordered) does not significantly affect judgements regarding a firm, it *does* affect judgments regarding a person. Specifically, an improving person is judged significantly less favorably when information items about him are presented out of order. In order to understand this difference one notes that presenting information out of temporal order - as with almost all Web search results - can render trends of change so as to be less readily noticeable. Nonetheless, this effect should apply to both persons and firms. The differential appreciation of improvement over time for persons versus firms when information is out of order may be explained in terms of the degree to which these two entity types are considered *liable to change*. Decision makers who perceive a high possibility for change may exert extra effort to consider timestamps closely and “connect the dots” in order to find any trend that may exist. Decision makers apparently perceive persons as less likely to change, leading to a higher likelihood of “missing” signs of their improvement. Observations later in this study with regard to implicit theories concerning persons versus firms support this conjecture.

5.4. Further Investigation of the Effect of Information Age on Perceived Importance

After first rating the candidate based on the entire collection of timestamped information and next rating the importance of each information item in that collection, participants were asked to specify - in retrospect - the degree to which they thought the age of information items

(timestamps) affected their decisions (1: far too little, 7: far too much). When collapsing across conditions the mean reported importance of timestamps was significantly lower when judging a person than when judging a firm (3.95 versus 4.19, $t(662.05) = 2.20$, $p = 0.03$), suggesting that decision makers do not temporally weight information concerning persons as much as they do for firms. Figure 4 depicts these results.





Figure 4. Self-reported importance of information item timestamps in assessing owner-operated contractor (Study 1) versus firm (Study 2).

Figure 5 presents a more fine-grained account of responses. It suggests that the perceived importance of timestamps is a function of: (1) the presence of a trend, (2) the visibility of that trend, and (3) the direction of that trend. In both studies the reported importance of timestamps is lowest when there is no trend (Condition 5). It is highest when there is a readily visible negative trend (Condition 2), followed by when there is a readily visible positive trend (Condition 1). Consistent with findings in Section 5.1, results in Figure 5 suggest that under Condition 3 (in which the entity has improved, but the change trend is not as readily apparent - a particularly important condition with respect to the right to be forgotten) - subjects are mostly “blind” to change in *persons* as indicated by the lower rating of timestamp importance in that condition

relative to Condition 1 in that Study 1 and the corresponding Condition 3 in Study 2. As also discussed in Section 5.1, I postulate that this is due to decision makers' differential perceptions concerning the likelihood of change in persons versus firms, as discussed further below.



Figures 5. Self-reported importance of timestamps across conditions and candidate types.

5.5. Effect of Changeability Perceptions on Noticing Behavior Change

I investigated the possible role of decision makers' perceptions concerning changeability of character on interpreting historical information regarding others and its implications for the right to be forgotten by administering the scale known as the "kind of person" (KOP) scale (Dweck, 1999) toward the end of the Study 1. This scale captures responders' general beliefs concerning the malleability or fixedness of character. An example item on this scale is: "people can do things differently, but the important parts of who they are can't really be changed." A lower score on this scale indicates holding an entity theory (i.e., people are not likely to change) while a high score indicates holding an incremental theory (i.e., people can change). I adapted the KOP scale in order to develop a new scale measuring perceptions concerning how much

firms can change and grow (see scale items and validation in Appendix 2). I call the new scale the “kind of company” scale (KOC) and administered this assessment toward the end of Study 2.

The panel on the left in Figure 6 shows the mean KOP and KOC along with 95% confidence intervals. The mean KOP score is 3.95, whereas the mean KOC score is 4.43 on a six-point scale. This difference is statically significant ($t = 5.80, p < 0.001$), supporting H6 predicting that firms are considered more liable (or capable) to change relative to persons. The panel on the right in Figure 6 shows a kernel density plot for KOP and KOC.




Figure 6. Left: mean “kind of person” (KOP) score (Study 1) and mean “kind of company” (KOC) score (Study 2). Right: distribution of KOP and KOC scores.

In order to investigate H5 regarding the role of implicit theories on candidate evaluations I median-split responders in Study 1 (Study 2) based on their KOP (KOC) scores, labeling those with above-median scores as “Incremental Theorists” and those with below-median score as “Entity Theorists.” I then investigated the effect of the change direction in candidates over time and responders’ implicit theories on candidate evaluations.

For Study 1 the results from ANOVA showed a significant main effect for the direction of change ($F(1, 227) = 41.88, p < 0.001$). There was additionally a significant interaction effect

between that variable and implicit theory ($F(1, 227) = 8.12, p = 0.005$), but no significant main effect for implicit theory ($F(1, 227) = .83, p = 0.36$) (see Panel A in Figure 7). As predicted, responders with an incremental implicit theory rated the candidate lowest when there was a negative change trend and rated highest when there was a positive trend. Nonetheless, as also predicted, changes in conduct over time did not significantly affect evaluations by entity theorists. A post-hoc analysis using Tukey correction showed that the latter group's evaluation did not significantly differ for candidate improvement or deterioration ($p = .10$). The other non-significant comparison was between the two conditions where the candidate had deteriorated ($p = .70$). Results accordingly show that implicit theories play a significant role for detecting change in *persons*, supporting H5 for this entity types.




Figure 7. Effect of responders' implicit theories and focal entity's change direction on the candidate's evaluation

A similar analysis on responses from Study 2 showed a significant main effect for the direction of change in candidate conduct ($F(1, 237) = 64.72, p < 0.001$). There was also a significant interaction effect between this variable and responder implicit theory ($F(1, 237) = 4.44, p = 0.04$), but no significant main effect for implicit theory ($F(1, 237) = 0.001, p = 0.97$) (see Panel B in Figure 7). A post-hoc analysis using Tukey correction showed that the differences in the candidate mean evaluation depicted in Figure 11 are significant except for two conditions at the bottom left ($p = .22$) and two conditions at the top right ($p = .70$). Contrary to the prediction by H5, these results suggest that there is no significant difference in *firm* rating by decision makers classified as incremental versus entity theorist. From the perspective of implicit theories, while only some decision makers “see” and appreciate signs of positive change in persons and the rest dismiss such signs as noise, all decision makers take signs of improvement in firms as signals of change; evidently, virtually all decision makers believe that firms can change.

6- GENERAL DISCUSSION AND CONCLUSIONS

We increasingly rely on information from the virtually eternal memory of the Internet in order to learn about others. Although increased access to information generally improves decision quality, concerns have also risen about how easy access to older information that is likely less diagnostic and would otherwise not be available can disproportionately influence our judgements (Ambrose, 2012; Jones, 2018; Mayer-Schönberger, 2009). In this paper I investigated several questions related to how decision makers use old versus new information in evaluating two major businesses entity types: owner-operated businesses and firms.

A major concern with using information from the Internet is that decision makers may not intuitively discern between old and new information since both are easily and vividly

available via a simple search. If true, such a phenomenon can hinder decision makers' abilities to identify change trends in others' behavior. Results from this study show that users do detect and appreciate signs of gradual improvement in others under certain but not all conditions.

Specifically, they detect change trends in firm conduct regardless of the order in which facts are presented, which is consistent with prior literature on trend effects (Favere-Marchesi, 2006).

However, decision makers fail to "see" and appreciate change for the better in owner-operated businesses (or more generally in persons) unless search results are chronologically ordered so that positive change is readily visible. This is seldom the case since Internet search results nearly always appear in random temporal order. I theorized that this observed differential effect can be explained in terms of the degree to which decision makers perceive persons versus firms to be capable of changing.

As hypothesized and consistent with prior literature (O'Brien & Klein, 2017), results also show that decision makers are less likely to detect a trend of change for the better versus a trend of change for the worse in individual conduct. This burden of persuasion and proof is apparently heavier in order to establish improvement than it is to infer deterioration. Contrary to the hypothesis, such an effect was not observed for firms since decision makers detected trends of improvement and deterioration equally effectively for that entity type. Again, a decision maker's inferior ability to detect improvement in a person versus a firm can be attributed to differential preconceptions regarding the degree to which each of these entity types can change and grow over time.

Apart from evaluations based on a collection of mixed information as discussed above, analyzing decision makers' reported weights for individual positive versus negative information items from various time points reflects a similar differential processing that places persons at a

disadvantage. As hypothesized and consistent with prior literature (Brandimarte et al., 2018), for persons bad conduct simply looms larger, regardless of time. However, for firms I find that decision makers demonstrate a higher appreciation for information concerning recent favorable conduct. Specifically, decision makers' valuations of importance for recent positive information were no different than their valuations for older adverse information. Prior research uses two argument streams explaining why "bad" is perceived as stronger than "good." The first stream attributes bad's strength to its higher affective load, while the other emphasizes the higher perceived informational value of bad relative to good (Baumeister et al., 2001). Results here lend credence to the latter argument: decision makers appear to consider bad conduct as more diagnostic of peoples' "true" and mostly fixed character, a conjecture supported by analysis of decision makers' implicit theories regarding change in persons versus firms.

These observed differential effects in evaluating owner-operated businesses (or more generally persons) versus firms can be attributable to the degree to which these two entity types are considered likely to change and grow over time. Analysis on decision makers' implicit theories revealed that persons are perceived as less likely to change than firms. Also, as predicted, when evaluating a person evaluators with an incremental implicit theory rated the candidate higher when there was a positive trend of conduct change while a positive conduct change did not affect decision makers with an entity implicit theory. However, firms are evaluated equally highly by all decision makers when there is an improving conduct trend, and equally poorly when there is a deteriorating trend. In sum, while only some decision makers "see" and appreciate signs of positive change in persons and the rest dismiss such signs as noise, all decision makers take signs of improvement in firms as true signals of change. These findings further highlight the importance of perceptions regarding change and changeability (i.e., implicit

theories) on how decision makers interpret historical information. Implicit theories have been investigated primarily in order to understand perceptions concerning change in persons, but more recently there have been calls for extending this idea toward firms (Wheeler & Omair, 2016). This study is one of the first to answer that call, and its results inform the theory and practice of both trust recovery (Haselhuhn et al., 2010) and firm recovery strategies (Yin et al., 2016).

6.1. Implications for Practice

The findings from this study have important implications for practice. Results indicate that owner-operated businesses, or more generally persons, should be more concerned with adverse information regarding past conduct since decision makers tend to attribute their errors primarily to fixed character while attributing firm behavior to time and circumstances.

Results also show that the “once a thief, always a thief” mentality investigated by prior research (Williams, 2015) is held more strongly for persons than firms. While firms’ older unfavorable conduct can be offset by newer favorable behavior in the eyes of the public, a person’s current favorable conduct tends to remain dominated by unfavorable behavior even from the distant past. In other words, unfavorable behaviors are more likely to “stick” to persons than to firms. Results also indicate that decision makers perceive that firms are more likely to change than persons. From a reputation management perspective this may provide a reason for owner-operated business to operate as firms rather than sole proprietors when possible, better positioning them to deal with the reputational consequences of errors.

6.2. Limitations and Directions for Future Research

As with any other empirical study this work has limitations that should be considered when interpreting the results. First, I investigated evaluations of businesses based on a collection

of mixed (favorable *and* unfavorable) information and found that given the same set of actions, those whose behavior shows a positive trend are evaluated higher than those whose behavior shows a negative trend or no trend. Nonetheless, this study did not investigate how decision makers evaluate those who improve relative to those who have consistently been good. Prior literature does provide interesting insights into this question. Specifically, research shows that those who change from bad to good tend to be considered *better* than those who have consistently been good. Prior research has explained this phenomenon using the contrast effect (Bless & Schwarz, 2010; Tversky & Griffin, 1991) through which the presence of “bad” provides a reference point for appreciating “good” more strongly, or using the general tenet that people prefer an improving sequence over an unchanging one even when all paths ultimately lead to the same absolute end point (Loewenstein & Prelec, 1993; Klein & O’Brien, 2017). If that is the case, then having some bad conduct to overcome in one’s distant past can in fact be advantageous rather than detrimental. However, such appreciation would be contingent on decision makers’ abilities to notice this trend of improvement, which per results from this study is more likely to be the case when evaluating a firm versus a person.

Future research can also investigate the nature of positive and negative past information being used to evaluate a business since the type of information (e.g., information on morality versus professional competency) can make a difference in these judgements (see Brandimarte et al., 2018; O’Brien & Klein, 2017; Reeder & Brewer, 1979). In addition, while this study focused primarily on the information items themselves, future research can investigate the possible effects of the information’s source (e.g., a generic Google search versus a search on Yelp, Bloomberg, etc.). Future research can also look into the possible effect of sole proprietors’ characteristics such as owner’s age or gender and firm characteristics such as size and type.

Finally, results in this study are based on responses from a US sample. Although this was necessary for control purposes, future research can gauge the degree to which these results generalize to other populations.

6.3. Conclusions

This study investigated various aspects of how past information from the Internet can affect decision makers' impressions concerning a business entity's present when the focal entity may have changed over time. Results from the study have important implications for the theory and practice of decision making and impression management in the information age. They additionally inform the debate concerning the merits and drawbacks of limiting access to adverse past information, commonly known as the right to be forgotten, as a topic with increasing appeal among researchers, practitioners, and the general public. Further investigating the effects identified here as well as other problems related to this emerging topic remain promising areas for impactful future research.

References

- Aivazpour, Z., & Rao, V. S. (Chino). (2020). Information Disclosure and Privacy Paradox: The Role of Impulsivity. *ACM SIGMIS Database: The DATABASE for Advances in Information Systems*, 51(1), 14–36. <https://doi.org/10.1145/3380799.3380803>
- Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgmental accuracy from thin slices of the behavioral stream. In *Advances in experimental social psychology* (Vol. 32, pp. 201–271). Elsevier.
- Ambrose, M. L. (2012). It's about time: Privacy, information life cycles, and the right to be forgotten. *Stan. Tech. L. Rev.*, 16, 369.
- Ambrose, M. L., Friess, N., & Van Matre, J. (2012). Seeking digital redemption: The future of forgiveness in the internet age. *Santa Clara Computer & High Tech. LJ*, 29, 99.

- Auxier, B. (2020). *Most Americans support the right to be forgotten online* | Pew Research Center. Pew Research Center. <https://www.pewresearch.org/fact-tank/2020/01/27/most-americans-support-right-to-have-some-personal-info-removed-from-online-searches/>
- Baskin, K. (1998). *Corporate DNA: Learning from Life*. Routledge.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370.
- Bélanger, F., & Crossler, R. E. (2011). Privacy in the digital age: A review of information privacy research in information systems. *MIS Quarterly*, 35(4), 1017–1042.
- Berkelaar, B. L. (2017). How implicit theories help differentiate approaches to online impression management: A preliminary typology. *New Media & Society*, 19(12), 2039–2058.
- Bertram, T., Bursztein, E., Caro, S., Chao, H., Feman, R. C., Fleischer, P., Gustafsson, A., Hemerly, J., Hibbert, C., & Invernizzi, L. (2018). Three Years of the Right to be Forgotten. *Google, Inc.*, 1–17.
- Bless, H., & Schwarz, N. (2010). Mental Construal and the Emergence of Assimilation and Contrast Effects. In *Advances in Experimental Social Psychology* (Vol. 42, pp. 319–373). Elsevier. [https://doi.org/10.1016/S0065-2601\(10\)42006-7](https://doi.org/10.1016/S0065-2601(10)42006-7)
- Brandimarte, L., Vosgerau, J., & Acquisti, A. (2018). Differential discounting and present impact of past information. *Journal of Experimental Psychology: General*, 147(1), 74.
- Chiu, C., Dweck, C. S., Tong, J. Y., & Fu, J. H. (1997). Implicit theories and conceptions of morality. *Journal of Personality and Social Psychology*, 73(5), 923.
- Chiu, C., Hong, Y., & Dweck, C. S. (1997). Lay dispositionism and implicit theories of personality. *Journal of Personality and Social Psychology*, 73(1), 19.
- Dweck, C. S. (1999). *Self-theories: Their role in motivation, personality, and development*. Psychology Press.
- Dweck, C. S., Chiu, C., & Hong, Y. (1995). Implicit theories and their role in judgments and reactions: A word from two perspectives. *Psychological Inquiry*, 6(4), 267–285.
- Favere-Marchesi, M. (2006). “Order effects” revisited: The importance of chronology. *Auditing: A Journal of Practice & Theory*, 25(1), 69–83.
- Fox, J., & Weisberg, S. (2019). *An R Companion to Applied Regression* (Third). Sage. <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- Frosio, G. F. (2016). The Right to Be Forgotten: Much Ado About Nothing. *Colo. Tech. LJ*, 15, 307.

- Haselhuhn, M. P., Schweitzer, M. E., & Wood, A. M. (2010). How Implicit Beliefs Influence Trust Recovery. *Psychological Science*, 21(5), 645–648.
<https://doi.org/10.1177/0956797610367752>
- Hayes, L. (2018). *More Than Half of Employers Have Found Content on Social Media That Caused Them NOT to Hire a Candidate, According to Recent CareerBuilder Survey*. PR Newswire. <https://www.prnewswire.com/>
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55. [https://doi.org/10.1016/0010-0285\(92\)90002-J](https://doi.org/10.1016/0010-0285(92)90002-J)
- Jones, M. L. (2018). *Ctrl+ Z: The right to be forgotten*. NYU Press.
- Kline, P. (2013). *Handbook of Psychological Testing*. Routledge.
<https://doi.org/10.4324/9781315812274>
- Korenhof, P., Ausloos, J., Szekely, I., Ambrose, M., Sartor, G., & Leenes, R. (2015). Timing the right to be forgotten: A study into “time” as a factor in deciding about retention or erasure of data. In *Reforming European data protection law* (pp. 171–201). Springer.
- Kruglanski, A. W. (1989). The psychology of being "right": The problem of accuracy in social perception and cognition. *Psychological Bulletin*, 106(3), 395.
- Loewenstein, G. F., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review*, 100(1), 91–108. <https://doi.org/10.1037/0033-295X.100.1.91>
- Lowry, P. B., Dinev, T., & Willison, R. (2017). Why security and privacy research lies at the centre of the information systems (IS) artefact: Proposing a bold research agenda. *European Journal of Information Systems*, 26(6), 546–563.
- Mantelero, A. (2013). The EU Proposal for a General Data Protection Regulation and the roots of the ‘right to be forgotten.’ *Computer Law & Security Review*, 29(3), 229–235.
<https://doi.org/10.1016/j.clsr.2013.03.010>
- Mayer-Schönberger, V. (2009). *Delete: The virtue of forgetting in the digital age*. Princeton University Press.
- O’Brien, E., & Klein, N. (2017). The tipping point of perceived change: Asymmetric thresholds in diagnosing improvement versus decline. *Journal of Personality and Social Psychology*, 112(2), 161.
- R Core Team. (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>

- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review*, 86(1), 61–79.
<https://doi.org/10.1037/0033-295X.86.1.61>
- Revelle, W. (2018). *psych: Procedures for Psychological, Psychometric, and Personality Research*. Northwestern University. <https://CRAN.R-project.org/package=psych>
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5(4), 296–320.
- Sheehan, K. B. (2018). Crowdsourcing research: Data collection with Amazon’s Mechanical Turk. *Communication Monographs*, 85(1), 140–156.
<https://doi.org/10.1080/03637751.2017.1342043>
- Shim, M. (2016). A Study of the Research the Right to be Forgotten from 2010. *Journal of the Korea Institute of Information Security and Cryptology*, 26(4), 1073–1084.
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989–1016.
- Tversky, A., & Griffin, D. (1991). Endowment and contrast. In M. Strack & Argyle, Norbert Schwarz (Eds.), *Judgments of Well-Being* (pp. 101–118).
- Walsh, G., & Beatty, S. E. (2007). Customer-based corporate reputation of a service firm: Scale development and validation. *Journal of the Academy of Marketing Science*, 35(1), 127–143. <https://doi.org/10.1007/s11747-007-0015-7>
- Wheeler, S. C., & O’Hair, A. (2016). Potential growth areas for implicit theories research. *Journal of Consumer Psychology*, 26(1), 137–141.
<https://doi.org/10.1016/j.jcps.2015.06.008>
- Williams, S. L. (2015). *Once a thief, always a thief? How time, implicit theories, and race affect moral judgments*. University of Waterloo.
- Xiao, P., & Lin, H. (2019). The right to be forgotten debate: Pros and cons. *American Journal of Business, Economics and Management*, 7(1), 40.
- Yin, C.-Y., Yu, H.-Y., & Poon, P. (2016). Consumers’ attributions and brand evaluations in product-harm crises: The role of implicit theories of personality: Consumers’ attributions in product-harm crises. *Journal of Consumer Behaviour*, 15(1), 87–95.
<https://doi.org/10.1002/cb.1549>

Appendix A

1: Experimental Material Creation Procedures

Creating the information items for use in Studies 1 and 2 involved two phases: a fact building and a fact rating phase. The goal of the fact building phase was to generate a list of positive and negative facts of the types typically found on the Internet concerning businesses. A small group of subjects participating in this task were presented with a scenario involving an institution in the process of hiring a contractor to renovate parts of its buildings and campus. Participants were asked to consider that they were responsible for researching the Internet in order to determine if a candidate was fit and qualified for being awarded that job. They were told that their job was to think of and write: 1) examples of information or news that if found would encourage them to award the job to the candidate, 2) examples of information or news that if found would discourage them from awarding the job to the candidate, and 3) examples of information or news they would consider neutral as related to the decision. Each participant wrote four to five examples for each category. Prior literature identifies product and service quality, financial performance, and social and environmental responsibility among the key dimensions forming a typical service provider's reputation (Walsh and Beatty 2007); the information items generated by participants in this task touched on several of these dimensions. I selected six facts from each of the favorable, unfavorable, and neutral categories, shortening and editing them as necessary for use in the subsequent fact rating phase.

The goal of the fact rating phase was to select a subset of the information items generated during the previous step for use in the main experiment. For this another group of participants were presented with the same scenario involving hiring a contractor for a project. They were asked to presume that their Internet search had resulted in finding certain pieces of information

about the contractor. Each participant's job was to rank order that information in terms of how positively or negatively it would affect evaluation of the candidate. Each participant was presented with two separate lists: one contained favorable information from the previous task and the other the unfavorable information from that task. Both lists also included some of the neutral information items. Participants were asked to rank order either of the two lists from most favorable to neutral and from most unfavorable to neutral, respectively. Using responders' rankings I picked three highly favorable-rated facts, three highly unfavorable-rated facts, and two neutral facts. The reason for including neutral facts was to make the scenario more realistic. This formed a collection of eight information items for use in the main experiment. I then generated two separate versions of the eight facts: one version referred to the candidate as a person and was used in Study 1, while the other referred to the candidate as a company and was used in Study 2.

2. Information Items Used in Studies 1 and 2 and Their Order of Presentation

The eight information items regarding the candidate used in Studies 1 and 2 are presented in Table A1 below.

Fact ID	Valance	Fact Narrative about the Owner-operated [Firm] Candidate
F1	Favorable	Job completion times and price estimates by him [the company] were rated among the top 5%.
F2	Favorable	He [The company] obtained the Gold Medal of Quality Work by City officials.
F3	Favorable	He [The company] was rated as being highly responsive and easy-to-reach by clients according to the Consumers Watch.
N1	Neutral	He [The company] moved his business office [its headquarters] from Elm Street to Hopkins Street.

N2	Neutral	The local newspaper reported that the water pipeline was damaged in a digging one block south of his office [the company headquarters].
U1	Unfavorable	He [The company] was fined for dumping industrial waste in the protected forests north of his [its] workshop.
U2	Unfavorable	Government prosecuted him [the company] for fraud in bidding on a government project.
U3	Unfavorable	He [The company] was ordered by a court to redo approximately 15 client projects due to using substandard material.

Table A1. Facts concerning the owner-operated [firm] contractor used in Study 1 [Study 2].

3. Analyzing the Valance and Potency of Information Items

Responses from 74 new participants drawn from the same population as the main study were used to further analyze the potency (perceived importance) and valance (perceived positive versus negative load) of the information items. Participants in this task were presented with the same scenario as in the main task with the exception that the information items were not accompanied by timestamps. After assigning the candidate an overall rating, participants rated the importance of each individual piece of (un-timestamped) information item in forming opinions concerning the candidate (1: “Extremely Unimportant”, 7: “Extremely Important”). They were next asked to specify how negative (unfavorable) or positive (favorable) they thought each information item was in the context of the problem presented (1: “Extremely Negative”, 7: “Extremely Positive”). Table A2 presents summary results from averaging the scores associated with these three negative, two neutral, and three positive information items.

	Information Item Category		
	Unfavorable	Neutral	Favorable
Rated Valance	1.9 (1.4)	3.5 (0.9)	5.9 (0.8)
Rated Potency	6.2 (0.9)	3.5 (1.6)	5.0 (1.5)

Table A2. Average (and standard deviation) perceived valance and potency of the information item across three categories.

Figures A1 and A2 visualize the results in Table A2 and feature 95% confidence intervals.




Figure A1. Rated valance of information items by category.




Figure A2. Rated potency of information items by category.

The observed differences in the valance and potency of information items across three categories are all statistically significant (p -values < 0.05). In terms of valance Figure A1 shows that the information items in the unfavorable category and those in the favorable category are nearly equally-distant from the scale midpoint of 4, indicating that they are perceived as equally strongly negative or positive, respectively. Although on par in terms of valance strength, consistent with prior literature on negativity bias (Baumeister et al., 2001; Rozin & Royzman, 2001), Figure A2 shows that negative information is rated as more potent than positive information.

Presentation Order and Time Stamps Assigned to Information Items

The presentation order and time stamps associated with the eight information items used in Studies 1 and 2 are presented in Table A3 below. Note that Conditions 1 and 3 and Conditions 2 and 4 are identical except for the order of presentation.

Presentation Order	Time Stamp*, Fact ID				
	Condition 1	Condition 2	Condition 3	Condition 4	Condition 5
1	Jan. 2018, F1	Jan. 2018, U1	Jan. 2018, F1	Jan. 2018, U1	Jan. 2018, F1
2	Oct. 2017, F2	Oct. 2017, U2	Apr. 1991, U1	Apr. 1991, F1	Oct. 2017, U1
3	May. 2013, N1	May. 2013, N2	Apr. 2001, N2	Apr. 2001, N1	May. 2013, N2
4	Jan. 2012, F3	Jan. 2012, U3	Mar. 1998, U2	Mar. 1998, F2	Jan. 2012, U2
5	Jun. 2006, U1	Jun. 2006, F3	Oct. 2017, F2	Oct. 2017, U2	Jun. 2006, F2
6	Apr. 2001, N2	Apr. 2001, N1	Jun. 2006, U3	Jun. 2006, F3	Apr. 2001, U3
7	Mar. 1998, U2	Mar. 1998, F2	Jan. 2012, F3	Jan. 2012, U3	Mar. 1998, F3
8	Apr. 1991, U3	Apr. 1991, F1	May. 2013, N1	May. 2013, N2	Apr. 1991, N1

Table A3. Presentation order of information items in Studies 1 and 2.

* Data for this study was collected in mid-2019.

Example Data Presented

An example of the information provided under Condition 1 in Study 2 is presented below:

- **Jan. 2018** Job completion times and price estimates by the company were rated among the top 5%.
- **Oct. 2017** The company obtained the Gold Medal of Quality Work by City officials.
- **May. 2013** The company headquarters was moved from Elm Street to Hopkins Street.
- **Jan. 2012** The company was rated as being highly responsive and easy-to-reach by clients according to the Consumers Watch.
- **Jun. 2006** The company was fined for dumping industrial waste in the protected forests north of its workshop.
- **Apr. 2001** The local newspaper reported that the water pipeline was damaged in a digging one block south of the company headquarters.
- **Mar. 1998** Government prosecuted the company for fraud in bidding on a government project.
- **Apr. 1991** The company was ordered by a court to redo approximately 15 client projects due to using substandard material.

Appendix B

Scales for Measuring Implicit Theories

I used the eight-item scale developed by Dweck (1999), commonly referred to as the “kind of person” (KOP) scale in order measure implicit theories regarding persons. In order to measure implicit theories concerning firms I adapted that scale and developed a new one. I refer to the new scale as the “kind of company” (KOC) scale.

I conducted reliability and item analysis on both the KOP and the KOC scales in order to determine and compare their internal reliability in the context of this task and population. I calculated Cronbach’s alpha using the psych package (Revelle, 2018) in R (R Core Team, 2019). In general, scales with value of alpha larger than 0.70 or 0.80 are considered reliable (Kline, 2013). Prior research shows a high internal reliability for the KOP scale with alphas in the range of .90 to .96 (Chiu, Dweck, et al., 1997; Dweck et al., 1995). My results show a similar level of internal reliability for both scales with alphas equal to .94 for the KOP scale and .92 for the new KOC scale. It was further determined that dropping none of the eight items from the new scale would improve the alpha. Details for the two scales used are provided below.

“Kind of Person” (KOP) Scale

The scale by Dweck (1999) uses 6-point Likert measures with the two ends anchored at “Strongly Disagree” and “Strongly Agree.” The letter R indicates reverse coding.

Using the scale provided, specify the degree to which you agree or disagree to the following statements.

1. The kind of person someone is, is something basic about them, and it can’t be changed very much. (R)

2. People can do things differently, but the important parts of who they are can't really be changed. (R)
3. Everyone is a certain kind of person, and there is not much that they can do to really change that. (R)
4. As much as I hate to admit it, you can't teach an old dog new tricks. People can't really change their deepest attributes. (R)
5. Everyone, no matter who they are, can significantly change their basic characteristics.
6. People can substantially change the kind of person they are.
7. No matter what kind of a person someone is, they can always change very much.
8. People can change even their most basic qualities.

“Kind of Company” (KOC) Scale

The scale measures implicit theories regarding companies on a 6-point Likert measure with the two ends anchored at “Strongly Disagree” and “Strongly Agree.” The letter R indicates reverse coding.

Some companies are better than others; even within a certain industry, companies can differ substantially in terms of quality of work, performance, responsible behavior, and alike. Using the scale provided, specify the degree to which you agree or disagree to the following statements.

- 1- The kind of venture a company is in terms of work quality, performance, responsible behavior, and alike is something very basic about it and it can't be changed very much.
(R)

- 2- Firms can do things differently, but the important parts of what they are, such as their performance, work quality, and responsible behavior, can't really be changed. (R)
- 3- Every company, no matter what it is, can significantly change its basic characteristics, such as its performance, work quality, and responsible behavior. (R)
- 4- As much as I hate to admit it, you can't teach an old dog new tricks. Enterprises can't really change their deepest attributes. (R)
- 5- Companies can always substantially change the kind of ventures they are in terms of work quality, responsible behavior, performance and alike.
- 6- Every company is a certain kind of establishment, and there is not much that can be done to really change that.
- 7- No matter what kind of venture a company is, it can always change very much.
- 8- All companies can change even their most basic qualities.