

Temporal Horizons and Decision-Making: A Big Data Approach

Robert Thorstad (rthorst@emory.edu)

Phillip Wolff (pwolff@emory.edu)

Emory University Department of Psychology

Abstract

Human behavior is plagued by shortsightedness. When faced with two options, smaller rewards are often chosen over larger rewards, even when such choices are potentially costly. In three experiments, we use big data techniques to examine how such choices might be driven by people's temporal horizons. In Experiment 1, we determine the average distance into the future people talk about in their tweets in order to determine the temporal horizon of each U.S. state. States with further future horizons had lower rates of risk taking behavior (smoking, binge drinking) and higher rates of investment (e.g., education, infrastructure). In Experiment 2, we used an individual's tweets to establish their temporal horizon and found that those with longer temporal horizons were more willing to wait for larger rewards. In Experiment 3, we were once again able to predict the choice behaviors of individuals from their tweets, this time showing that those with longer future horizons were less likely to take risks. The findings help establish a powerful relationship between people's thoughts about the future and their decisions.

Keywords: prospection; future thinking; big data.

Introduction

People often act impulsively. They eat unhealthy foods, gamble, overspend, and engage in behaviors that are likely to have a negative impact on their wellbeing. Such shortsighted behaviors demonstrate the phenomenon of *delay discounting*, a type of reasoning in which long-term benefits or risks are minimized in favor of smaller short-term rewards. Delay discounting is seemingly non-normative and maladaptive. Immediate and distant futures are equal parts of one's life as a whole, so it would appear to be in people's best interest to wait for larger rewards. And yet, people continue to make poor decisions, despite the long-term risks and in spite of the potential benefits of waiting for larger rewards. In this paper, we use big data techniques to address the question of why delay discounting might occur.

We hypothesize that delay discounting may be due, in part, to the way individuals think about the future. Specifically, we hypothesize that individuals with a long *temporal horizon* – who tend to think far into the future – will be less likely to discount future rewards.

There is mixed evidence in the literature for the role of future thinking in delay discounting. There is some evidence that individuals who tend to think about the future more in general (e.g. have a future time perspective) are less likely to discount future rewards (Steinberg et al, 2009; Daugherty & Brase, 2010), although effect sizes tend to be small. Additionally, there are reports that encouraging individuals to think about the future more vividly (e.g. episodic future

thinking) reduces delay discounting (Peters & Buchel, 2010; but see also Kwan, Craver, Green, Myerson, & Rosenbaum, 2013).

Findings for temporal horizon, however, are more mixed. Heerey, Matveeva, & Gold (2011) report that schizophrenic patients with a longer time horizon are less likely to discount future rewards. However, the authors find no relationship between time horizon and discounting for healthy adult controls. Additionally, Fellows & Farah (2005) report that patients with lesions to ventromedial prefrontal cortex have a shortened time horizon but show normal delay discounting. These results provide mixed evidence for the role of time horizon in delay discounting.

One reason for these mixed results may be the problem of how to measure temporal horizon. Temporal horizon is usually measured by asking participants to imagine future events and then estimate their distance in the future. However, such explicit tasks may not accurately reflect participants' time horizon. Additionally, time horizon may be subject to state effects that are not captured in a single laboratory measurement. The current research addresses these limitation through an analysis of people's naturally occurring language on twitter about future events. By measuring future horizon in this way, we can more accurately assess its potential impact on people's choices about immediate and future rewards. Additionally, because twitter data is longitudinal in nature, we can address potential state effects by measuring a participant's time horizon over many different time points.

Experiment 1: Discounting in Populations

The hypothesis that delay discounting depends on future horizon should extend beyond the choices of individuals to the collective choices of an entire population. Because various statistics of impulsive / long-range thinking are already available for populations, we began our investigation by analyzing the temporal horizon and collective choices of entire states. To conduct this kind of analysis we first collected a large number of tweets from each state. The tweets were then automatically analyzed for whether they referred to the future or past, as well as their distance into the future or past. Using these analyses, we could determine each state's *future orientation*, that is, the degree to which the tended to talk about the future or the past, as well as each state's *temporal horizon*, that is, distance into the future and past. Once each state's temporal horizon was determined, we could examine the extent to which each state's future horizon correlated with various kinds of collective choices. In particular, we examined the association between future horizon and various kinds of




Figure 1: Mean future orientation of each U.S. state. Colors range from yellow (least future) to dark red (most future). Note that the 4 states excluded from analysis are dummy-coded with the mean value.




Figure 2: Mean temporal horizon of each U.S. state. Colors range from yellow (shortest horizon) to dark red (longest horizon). Note that the 4 states excluded from analysis are dummy-coded with the mean value.

impulsive behaviors, such as binge drinking rates, cigarette smoking rates, and drunk driving rates. We also examined the association between future horizons and healthy choice behaviors, such as seatbelt usage, education spending, and highway spending. Because of the theoretical novelty of temporal horizon, we also compared our measure of future horizon with several demographic variables including political orientation and economic activity.

Methods

Tweet Collection. We collected 8,550,131 tweets from April to July 2015 using the Twitter Developer API. Tweets were restricted to English language only. Tweets were collected for equal durations from each U.S. state.

Future Orientation Classification. We defined the future orientation of a state as the number of tweets about the future, divided by the number of tweets that could be classified as either past or future ($\text{future} / (\text{future} + \text{past})$). We classified tweets as about the past, future, or neither using a custom-built classifier. For each tweet, we performed part-of-speech parsing using the Stanford Parser (Chen & Manning, 2014). The result is a sentence represented as a combination of part-of-speech-tags (e.g. “NP” for noun phrase) and words. We then developed a set of 112 lexical and syntactic rules to classify the tweet as past, future, or neither. For example, the tweet “Maybe ill stay home w mom tomorrow” was classified as future because it contains a singular noun “home” dominating the lexical item “tomorrow” (the relevant rule is $\text{NN|RB} > \text{tomorrow}$). In total, about 31% of tweets were classified as future and about 23% of tweets were classified as past. The future classifier was validated against human ratings of 1,000 sentences, showing 76.61% agreement with human raters (precision=0.75, recall=0.90). This agreement with human raters compares favorably to other attempts in the literature (e.g. Nakajima et al, 2014), and approaches human accuracy (human-to-human percent agreement = 86.44%). This

classifier is available for use on our website: <http://mindandlanguagelab.com/futureAnalysis>.

Temporal Horizon Classification. For future tweets only, we classified their temporal horizon using a keyword approach. We created a list of 58 twitter-appropriate keywords marking explicit periods of time: for example, “tomorrow”, “tmrw,” etc. For each keyword, we estimated the average number of minutes in the future it occurs. For example, “tomorrow” occurs on average 1440 minutes (24 hours) in the future. For each future tweet, we used regular expression matching to identify temporal keywords in the tweet. For example, “maybe ill stay home with mom tomorrow” matched the temporal noun “tomorrow.” 258,281 tweets had identifiable time horizons under this criterion. For each tweet that matched one or more temporal keywords, we averaged the resulting number of minutes to estimate its time horizon. Because time horizon had a large range (180-259,200 minutes) we took the natural logarithm of the result. We then averaged temporal horizon at the state level, creating a single score representing the average temporal horizon of each U.S. state. To ensure accurate estimation, we excluded states with fewer than 100 tweets with identifiable time horizons (3 States excluded: AK, ME, MT). We also excluded 1 state (HI) because its time horizon was more than 5 SD above the mean; this may be due to a large number of tweets from non-residents on vacation.

Demographic Measures. To explore the characteristics of states with different time horizons and future orientation, we collected state-level economic, political, religiosity & well-being indices from Gallup inc. We separately correlated time horizon and future orientation at the state level with each measure.

Risky Decision-Making. We hypothesized that time horizon might relate to risk, such that states with longer time horizon would take fewer risks. To evaluate this hypothesis, we collected 6 state-level indices of risky behavior from publicly available government data. We collected 3 indices of risky behavior (binge drinking,

cigarette smoking, and drunk driving) and 3 indices of behavior we considered the opposite of risky (seatbelt usage, education spending, and highway spending). Education and highway spending data were collected from the US Census; the other indices were collected from the Center for Disease Control. We then created a *risk composite* by separately normalizing the risk indices, correcting sign such that positive numbers represent risk, and averaging the resulting z-scores for each state.




Figure 3: Temporal horizon and risky decisions by state. Labels represent state abbreviations. Note that the horizontal axis is plotted in logarithmic scale.

Results

Temporal Horizon by State. We first present the average temporal horizon for each U.S. state. The mean temporal horizon was 1.27 days, see Figure 1. The U.S. census region with the longest time horizon was the Northeast; the Midwest had the shortest time horizon.

Future Orientation by State. We next present the average future orientation for each state. The mean future orientation was 31.55% ($SD = 1.02\%$). Figure 2 displays the mean future orientation for each U.S. state. To evaluate the relationship between time horizon and future orientation, we calculated the Pearson correlation between states' temporal horizon and future orientation. There was a negative relationship between future orientation and temporal horizon, $r(46) = -0.411$, $p < 0.01$. One potential explanation for this finding is that thoughts about the close future may be more frequent than thoughts about the far future.

Demographic Measures. To begin to investigate the characteristics of states with different time horizon and future orientation, we conducted an exploratory correlation analysis with 29 economic, political, religiosity, and well-being measures from Gallup, inc. States with a long time horizon were more politically liberal than states with a short time horizon, $r(45) = 0.533$, $p < 0.001$. By contrast, states with greater future orientation, were less politically liberal than less future-oriented states, $r(45) = -0.398$, $p < 0.01$.

Table 1: Results of a linear regression predicting risky decision-making at the state level. Future time horizon was still significantly negatively related to risk when including demographic and computational controls.

	β	p
Future Horizon	-0.368	0.023 *
Population	0.575	0.482
Gender	0.175	0.175
Median Age	0.015	0.916
GDP	-0.930	0.260
Past Horizon	-0.057	0.688

Risky Decision-Making. The main hypothesis evaluated was that states with long time horizon will take fewer risks than states with short time horizon. To evaluate this hypothesis, we created a composite of 6 risky decision-making indices at the state level, and separately correlated the composite with time horizon and future orientation. States with a long time horizon took significantly fewer risks than states with a short time horizon, $r(45) = -0.467$, $p < 0.01$ (Figure 3). Because New Jersey had a risk composite score more than 3 SD less than the mean, we verified that the relationship still holds excluding New Jersey, $r(45) = -0.352$, $p < 0.05$.

To verify that this relationship was not due to a simple tendency to think more about the future, we correlated states' future orientation with risk composite score. There was no relationship between future orientation and risk, $r(46) = 0.271$, $p = 0.069$.

In order to rule out alternative explanations for the relationship between time horizon and risk, we conducted a linear regression analysis controlling for several additional variables (Table 1). To control for demographic differences between states, we controlled for gender (percent male), population, median age, and GDP at the state level. To validate our computational method, we included an additional control by calculating a state's past horizon. To do this, we performed every computational step used to calculate a state's future horizon, except using tweets classified as past, not future. To evaluate whether future time horizon is negatively related to risk, we separately fit linear regression models with all control variables, with and without future time horizon. Adding future time horizon significantly improved model fit: $R^2 = 0.401$ (with time horizon) versus 0.315 (without time horizon), F-test: $F(1,39) = 5.564$, $p < 0.05$.

Discussion

In Experiment 1, we evaluated the hypothesis that having a long temporal horizon increases sensitivity to potential future costs. To do this we built a classifier to identify how near or far U.S. states tend to think in the future, and correlated the results with an index of risky decision-making.

The main result of Experiment 1 was that states with longer time horizons took fewer risks than did states with shorter time horizons. This result was specific to temporal

horizon, rather than future orientation in general. This result remained significant when including several demographic and computational controls.

One limitation of Experiment 1 was that analysis was conducted at the level of states, not individuals. As such, it was impossible to obtain an experimental measure of risky decision-making. The aim of Experiment 2 was to address this limitation by linking individual participants' temporal horizon to an experimental measure of future decision-making. Additionally, we aimed to extend our results beyond future costs to potential future benefits. If having a long time horizon increases sensitivity to potential future costs, does it also increase sensitivity to potential future benefits?




Figure 4: A sample delay discounting trial. In this trial, the participant is asked to choose between \$60 today and \$100 after a delay of 6 months.

Experiment 2: Discounting in Individuals

The aim of Experiment 2 was to test the hypothesis that individuals with a long temporal horizon are less likely to discount future rewards. Participants completed a delay discounting task where they made a series of 60 choices between smaller present and larger future rewards (Figure 4). Participants also provided their twitter handle, allowing identification of their individual temporal horizon. If temporal horizon affects delay discounting, participants with a long temporal horizon should be more likely to choose to wait for future rewards.

An additional aim of Experiment 2 was to investigate the stability of temporal horizon over time. To do this, we separately classified participants' temporal horizon using only recent tweets from the past 2 and 4 weeks, as well as using tweets from all dates.

Methods

Participants. 198 participants were recruited via Amazon Mechanical Turk and received \$1 for participation. 29 participants were excluded for failing to provide a valid twitter handle ($N=7$), providing a "protected" twitter account that could not be processed ($N=14$), or failing to complete the delay discounting task ($N=8$).

Delay Discounting. Delay discounting questions were composed by fulling crossing 6 delay lengths (1 week, 6 months, 1 year, 5 years, 10 years, 20 years) with 10 immediate reward amounts (\$1, \$5, \$10, \$20, \$40, \$60, \$80, \$90, \$95, and \$99). The delayed reward was always \$100. For example, on one trial, participants chose between \$60 today and \$100 in 5 years. Participants completed 60 delay

discounting trials in random order. We excluded data from 13 participants who chose the future reward less than 5 times.

Temporal Horizon Classification. Participants were asked to provide their valid twitter handle with at least 50 tweets. Temporal horizon was classified using the same methods as in Experiment 1, with 2 additions. First, we added an additional preprocessing step to remove URLs, hashtags (#), and twitter user mentions (@) using regular expression matching. Second, to investigate the stability of temporal horizon over time, we separately classified participants' temporal horizon using only tweets from within 7 days, 14 days, 30 days, 60 days, 180 days, and 360 days from the experiment, as well as data from all time periods. To ensure accurate estimation of time horizon, we excluded participants with fewer than 10 tweets (at any date) for which time horizon could be calculated.




Figure 5: Participants with long time horizon were more willing to wait for future rewards. The horizontal axis is plotted in logarithmic scale.

Results

First, we evaluated whether participants' temporal horizon affected delay discounting. For each participant, we calculated a *reward index*, e.g. the total money earned on the delay discounting task, divided by the total possible earnings. A large reward index indicates a willingness to wait for future rewards. Participants with a long time horizon were more likely to wait for future rewards, $r(100) = 0.204$, $p < 0.05$ (Figure 5).

Second, we asked whether participants' decisions were better predicted by their recent time horizon, or by their time horizon based on all tweets. We did not find that the temporal horizon associated with more recent tweets was more predictive of delay discounting than the temporal horizon associated with all tweets (all $ps > 0.2$). This could imply that temporal horizon is a trait-like characteristic that is relatively stable in an individual. However, it is also the case that the number of recent tweets was relatively low, which could reduce the ability to find an effect for recent tweets.

Discussion

The main finding of Experiment 2 was that participants' temporal horizon affected their delay discounting. Using all tweets to classify temporal horizon, participants with a long temporal horizon were more likely to wait for future rewards. This result supports the hypothesis that one reason participants discount future rewards is their cognitions about the future.

Experiment 2 demonstrated that temporal horizon affects decision-making in situations where a benefit occurs only in the long future. In Experiment 3, we wondered whether the same might be true for future costs.




Figure 6: The BART task. Participants earned points for inflating balloons, and at any point could bank points earned and proceed to the next trial.

Experiment 3: Risk Taking in Individuals

The aim of Experiment 3 was to explore the relationship between temporal horizon and decision-making where the future outcome is a cost, not a benefit. In risky decision-making, participants must trade off between a potential gain in the present, and a loss in the future. In Experiment 2, we found that a long temporal horizon increased participants' sensitivity to future gains. However, in Experiment 1, we also found at the level of states that a long temporal horizon increased states' sensitivity to future costs, e.g. risks. The aim of Experiment 3 was to test at the individual level whether a long temporal horizon might increase participants' sensitivity to future costs.

In Experiment 3, participants completed the Balloon Analogue Risk Task (BART; Lejuez et al, 2003). In this task, participants are presented with a series of 30 balloons (Figure 6). Participants earn points every time they inflate the balloon, but they take a risk by doing so: the balloon may pop, resulting in no points for the trial. Alternatively, at any time participants may bank points currently earned and proceed to the next trial.

To evaluate whether future horizon affects evaluation of future costs, we also asked participants to provide their twitter handle. For each individual, we identified their temporal horizon and future orientation on the basis of their tweets. If temporal horizon affects evaluation of future costs, participants with a long temporal horizon should take fewer risks in the BART task.

Methods

Participants. 83 participants were recruited via Amazon Mechanical Turk and received \$1 for participation. There were 3 exclusion criteria. Participants were excluded for failing to provide a valid twitter handle ($N=9$), providing a "protected" twitter account that could not be processed ($N=8$), or failing to complete the BART task ($N=1$).

Temporal Horizon and Future Orientation Classification. Participants were asked to provide their valid twitter handle with at least 50 tweets. Temporal horizon and future orientation classification were performed for each individual twitter account using the methods from Experiment 1, with 2 changes. First, we removed hashtags (#) and user mentions (@) using the same methods as Experiment 2. Second, to improve classification accuracy at the individual level, we simplified the temporal horizons classifier by restricting temporal keywords to those referencing tonight (e.g. "tonight," "2nite," etc.) and references to "tomorrow." The total number of tweets collected was 117,281. 14 participants were excluded because their twitter account did not provide any tweets with identifiable time horizon.

BART Task. Participants completed a 30-trial version of the BART task, using a javascript implementation by Timo Gnams (2013). Balloon explosion points were drawn from a uniform distribution from 1-128 clicks. There were 3 balloon colors (10 balloons each) with differing points gained per click: 0.5, 1, or 5 points. We calculated two measures of participants' risk taking. (1) *Adjusted pumps*: mean number of inflations per balloon, excluding trials where the balloon exploded. (2) *Number of explosions*: number of balloons inflated until explosion.




Figure 7: Participants with a long time horizon exploded fewer balloons in the BART task. The horizontal axis is plotted in logarithmic scale.

Results

The main hypothesis evaluated was that participants with a long time horizon will take fewer risks. Supporting this prediction, participants with a longer time horizon exploded

fewer balloons, $r(54) = -0.321$, $p < 0.05$ (Figure 7) and trended towards inflating balloons fewer times, adjusted pumps: $r(54) = -0.230$, $p = 0.094$. This supports the claim that participants with a long time horizon are more sensitive not only to future benefits, but also to future costs.

We also verified that our results were driven by temporal horizon, rather than a mere tendency to think about the future. There was no relationship between future orientation and BART performance, $r(54)=0.027$, $p=0.846$ (explosions); $r(54)=0.067$, $p=0.632$ (adjusted pumps). This lack of result supports the conclusion that temporal horizon, and not a mere tendency to think about the future, increases sensitivity to future costs.

Discussion

The main result of Experiment 3 was that having a long temporal horizon increases sensitivity to potential future costs, not just benefits. Individual participants' temporal horizon, but not their future orientation in general, predicted risks taken in the BART risk task.

General Discussion

The results provide support for the hypothesis that delay discounting is determined, at least in part, by a person's future horizon. In Experiment 1, we showed that U.S. states with longer time horizons, but not greater future orientation, were more likely to engage in risky and impulsive decisions. In Experiment 2, we found that people with longer future horizons, as revealed by their tweets, were more likely to wait for future rewards. In Experiment 3, we found that people with longer horizons were more likely to avoid risks.

Although not directly examined in this paper, the experiments provide some insight into the length of people's future horizon. For example, in Experiment 1 we observed that the average temporal horizon was approximately 1.27 days. In contrast, in Experiment 2 we found a slightly longer temporal horizon: 2.01 days. Given that the temporal horizon in Experiment 2 was almost twice as long as in Experiment 1, the difference between these experiments may appear surprising. However, the difference between these experiments can probably be explained by differences in sampling. In Experiment 1, we randomly selected tweets from the twitter stream, and as a consequence no doubt sampled frequent twitter users more often than infrequent twitter users. It could be that frequent twitter users tend to have a shorter temporal horizon than infrequent users. The way we sampled in Experiment 2 was less susceptible to this potential bias because twitter users were selected based on their participation in a separate experiment. Arguably, then, the mean temporal horizon from Experiment 2 may be more representative of a typical temporal horizon.

One question not resolved by the current findings is whether temporal horizon is a relatively stable characteristic of individuals, i.e. a trait, or whether temporal may vary as a function of context and experience, making it more like a state. To the extent that temporal horizon is a state, it may be possible to have an impact on people's decision-making

processes by changing temporal horizon. On the other hand, it could be that the distance people look into the future is largely independent of context, and as a consequence these biases may be a relatively stable characteristic of individuals. Regardless, by looking at people's temporal horizon we can see how decisions may depend on people's temporal biases.

Acknowledgments

This research was made possible through the support of a grant from the John Templeton Foundation, "Prospective Psychology Stage 2: A Research Competition" to Martin Seligman. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation. The authors wish to thank Bridget Copley for assistance developing the future orientation classifier.

References

- Chen, D. & Manning, C. (2014). A fast and accurate dependency parser using neural networks. *Proceedings of EMNLP 2014*.
- Daugherty, J. & Brase, G. (2010). Taking time to be healthy: predicting health behaviors with delay discounting and time perspective. *Personality and Individual Differences*, 48(2), 202-207.
- Fellows, L. & Farah, M. (2005). Dissociable elements of human foresight: a role for the ventromedial frontal lobes in framing the future, but not in discounting future rewards. *Neuropsychologia*, 43(8), 1214-1221.
- Gnambs, T. (2013). Balloon Analogue Risk Task [Computer software]. Retrieved from timo.gnambs.at.
- Heerey, E., Matveeva, T., & Gold, F. (2012). Imagining the future: degraded representations of future rewards and events in schizophrenia. *Journal of Abnormal Psychology*, 120(2), 483-489.
- Kwan, D., Craver, C., Green, L., Myerson, J., & Rosenbaum, R. (2013). Dissociations in future thinking following hippocampal damage: evidence from discounting and time perspective in episodic amnesia. *JEP: General*, 142(4), 1335-1369.
- Lejuez, C. et al. (2003). The balloon analogue risk task (BART) differentiates smokers and nonsmokers. *Experimental and clinical psychopharmacology*, 11(1), 26.
- Nakajima et al. (2014). Investigation of future reference expressions in trend information. *Proceedings of the AAIL Spring Symposium Series*, 31-38.
- Peters, J. & Buchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron*, 66(1), 138-148.
- Steinberg, L., Graham, S., O'Brein, L., Woolard, J., Cauffman, E., & Banich, M. (2009). Age differences in future orientation and delay discounting. *Child Development*, 80(1), 28-4.