



2016

Deciding to Curtail Persistence

Joseph T. McGuire

University of Pennsylvania, mcguirej@psych.upenn.edu

Joseph W. Kable

University of Pennsylvania, kable@psych.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/psychology_papers



Part of the [Psychology Commons](#)

Recommended Citation (OVERRIDE)

McGuire, J. T. & Kable, J. W. (2016). Deciding to Curtail Persistence. In K.D. Vohs & R.F. Baumeister (Eds.), *Handbook of Self-Regulation: Research, Theory, and Applications* (pp. 533-546). New York: Guilford Press.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/psychology_papers/33
For more information, please contact repository@pobox.upenn.edu.

Deciding to Curtail Persistence

Abstract

Imagine that a few seconds ago you called a restaurant to book a reservation and were placed on hold. How soon do you expect to be helped? Are you having any difficulty waiting?

Now imagine 5 minutes have gone by and you are still hearing hold music. Is it getting more difficult? Have your expectations changed? How much longer will you give them?

Voluntary persistence toward delayed rewards has often been framed, in the psychological literature, as a self-control problem. This view presumes that it is generally beneficial to direct one's behavior toward valuable prospects in the future, but that the fallible nature of self-control makes people sometimes succumb to immediate temptations instead. In laboratory studies, individuals who wait longer for delayed rewards have been deemed to possess greater self-control capacity.

In real life, though, how long it is worth holding out for future rewards can be a more vexed question. Not all long-run rewards is complicated by the fact that future events are uncertain in both their substance and their timing. When it comes to choosing how long to wait for everything from city buses to customer service representatives, decision makers can as easily err by waiting too long--chasing sunk costs-- as by waiting too little. In this chapter we review research suggesting that the challenge of delaying gratification does not emerge merely from psychological limitations but instead reflects the genuine complexity of the environments in which real-world decisions take place.

Disciplines

Psychology

CHAPTER 29

Deciding to Curtail Persistence

JOSEPH T. MCGUIRE

JOSEPH W. KABLE

Imagine that a few seconds ago you called a restaurant to book a reservation and were placed on hold. How soon do you expect to be helped? Are you having any difficulty waiting?

Now imagine 5 minutes have gone by and you are still hearing hold music. Is it getting more difficult? Have your expectations changed? How much longer will you give them?

Voluntary persistence toward delayed rewards has often been framed, in the psychological literature, as a self-control problem. This view presumes that it is generally beneficial to direct one's behavior toward valuable prospects in the future, but that the fallible nature of self-control makes people sometimes succumb to immediate temptations instead. In laboratory studies, individuals who wait longer for delayed rewards have been deemed to possess greater self-control capacity.

In real life, though, how long it is worth holding out for future rewards can be a more vexed question. Not all long-run prospects are worthwhile goals, and not all immediate rewards are wayward temptations. Evaluating the relative merits of immediate and long-run rewards is complicated by the fact that future events are uncertain in both their substance and their timing. When it comes to choosing how long to wait for everything from city buses to customer service representatives, decision makers can as easily err by waiting too long—chasing sunk costs—as by waiting too little. In this chapter we review research suggesting that the challenge of delaying gratification does not emerge merely from psychological limitations but instead reflects the genuine complexity of the environments in which real-world decisions take place.

SEEMINGLY INCONSISTENT DECISION BEHAVIOR

When it comes to diagnosing self-control failure, one of the most incriminating phenomena occurs when people reverse their own previous choices about the future. Take, for

example, decisions about exercise. If you encounter someone who does not regularly exercise, you might simply assume he or she has other priorities. If, however, you encounter someone who began an exercise routine and dropped it within a week, you might be more inclined to suspect he or she is having difficulty following through on goals.

One of the best-known laboratory paradigms for eliciting reversals of future-oriented behavior was developed to study delay of gratification in children (Mischel & Ebbesen, 1970; Mischel, Ebbesen, & Zeiss, 1972; Mischel, Shoda, & Rodriguez, 1989). In the delay-of-gratification task, an experimenter offers a 4-year-old child a choice between a smaller and a larger food reward. The child can obtain the larger and more preferred reward by waiting alone in a room until the experimenter returns, or can obtain the smaller reward by summoning the experimenter at any time.

Behavior in this paradigm varies greatly across individuals. Data from a large replication study sponsored by the National Institutes of Health are shown in Figure 29.1 (McGuire & Kable, 2013; National Institute of Child Health and Human Development [NICHD] Early Child Care Research Network, 2002). The results are plotted in the form of a survival curve, showing what fraction of the children were still waiting at each time from the beginning of the delay until the 7-minute mark, when the large reward was delivered. At one extreme, about 53% of the participants waited the entire time. At the other extreme, about 17% quit within the first 15 seconds, consistent with having rapidly decided that it was not worth waiting for the larger reward. In between, the remaining 30% of the participants showed a more puzzling and seemingly inconsistent pattern: They waited at first, but at some point gave up and received the same small reward they could have obtained without waiting at all. It seems reasonable to construe this reversing sequence of choices—first trying to wait but then quitting—as a *delay of gratification failure*.

Various potential explanations for delay of gratification failure have been discussed in the psychological literature, and most appeal to the idea of limited self-control capacity. The child's initial choice is thought to reflect a belief that the benefits of waiting outweighed the costs, but the subsequent reversal is taken as evidence that the child lacked the ability to follow through (Mischel, Ayduk, & Mendoza-Denton, 2003). The specific cognitive skills that make the difference between successful and unsuccessful waiting are




FIGURE 29.1. Behavior of 4-year-olds in the delay-of-gratification paradigm. The experimenter returned and delivered the large reward after 7 minutes if the child was still waiting. Data are from the Study of Early Child Care and Youth Development (NICHD Early Child Care Research Network, 2002). From McGuire and Kable (2013). Copyright 2013 by the American Psychological Association. Adapted by permission.

thought to include cognitive tactics, such as the strategic control of attention, that help to manage the balance between hot and cool motivational systems (Duckworth, Gendler, & Gross, 2014; Eigsti et al., 2006; Metcalfe & Mischel, 1999; Mischel et al., 1972; Rodriguez, Mischel, & Shoda, 1989).

Less-often appreciated, however, is that decision makers faced with delay might have good reason to reverse their original patient choice even if their self-control abilities are fully intact. When the timing of future events is uncertain, it can be appropriate to change one's subjective valuation of future rewards merely as a function of how much time has passed already.

TEMPORAL UNCERTAINTY AND TIME-VARYING REASSESSMENT

Most future rewards are associated with some degree of uncertainty, and this uncertainty can take at least two forms. Sometimes one is unsure whether the rewarding event will occur at all. In laboratory experiments, children quite reasonably become less interested in waiting for future rewards if they have reason to mistrust whether the experimenter will deliver the rewards as promised (Kidd, Palmeri, & Aslin, 2013; Mahrer, 1956). In other cases, one might believe an event is fairly sure to occur eventually but still might face substantial uncertainty about the duration of the delay. Consider real-life situations such as waiting for a city bus, waiting in a long checkout line, or waiting for the end of a boring speech. In these situations there is no real risk of the awaited event never arriving, but one might be anxious about the possibility that it might take an unreasonable amount of time.

The original delay-of-gratification laboratory experiments included deliberate measures to remove the first kind of mistrust as a factor. The protocol ensured that the children knew the large reward existed and the experimenter would eventually deliver it (Mischel & Ebbesen, 1970; Mischel et al., 1972). However, the paradigm was designed to create a high degree of temporal uncertainty. Participants received no specific information in advance about how long they would be expected to wait: The experimenter told them the delay would last "a while," and added, "Sometimes, I'm gone a long time" (Mischel & Ebbesen, 1970, p. 332). Even if a specific time interval had been mentioned, it is far from guaranteed that 4-year-old participants would have been able to make full use of the information. This is an important and face-valid feature of the paradigm, since temporal uncertainty is the rule rather than the exception for many kinds of real-life future rewards.

When people are unsure how long something will take, they nonetheless are able to express an intuitive sense of what the possibilities are. People's explicit predictions about the timing of various kinds of future events turn out to be reasonably well calibrated to the true statistics of the relevant domain (Griffiths & Tenenbaum, 2006). People understand, based on their background knowledge, that different forms of uncertainty characterize different situations.

A decision maker's expectations about an uncertain future event can be thought of in terms of a subjective probability distribution. For example, if you were wondering when the next subway train would arrive during a typical evening commute, your expectations might be well characterized by the normal distribution in the top-left of Figure 29.2.

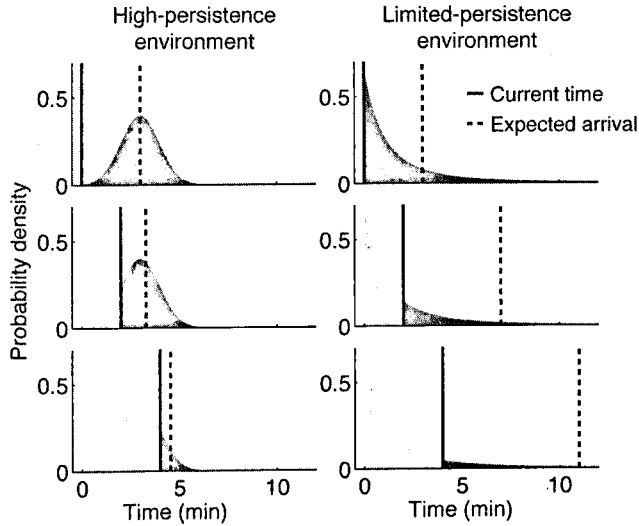


FIGURE 29.2. Two environments that entail opposite implications for behavioral persistence. The filled regions represent a decision maker's subjective probability distribution for a reward's arrival time. Faded areas represent times that have been ruled out because they are shorter than the time one has already waited; darker areas represent times that are still possible. The dashed line represents a prediction of the reward's arrival time, obtained by averaging over the times that are still possible. This prediction evolves as a function of elapsed time (solid line), and does so differently in the two environments. From McGuire and Kable (2012). Copyright 2012 by Elsevier B. V. Adapted by permission.

This reflects an expectation that the train will probably take somewhere between 1 and 5 minutes. If pressed for a prediction, you might use the average of all the possibilities (or perhaps the median) and predict a waiting time of 3 minutes.

Once the subjective probabilities are written down in this way, it becomes straightforward to think about how one's prediction would evolve with the passage of time. As time elapses (with no train in sight), short delay lengths are ruled out. This is represented in the left-hand column of Figure 29.2, with rows representing the situation in which zero, 2, or 4 minutes have already elapsed. As long as the train does not appear, one can adjust one's expectations by excluding all the delays shorter than the amount of time that has already elapsed. At any given moment, one might offer an updated prediction based on the average of the times that are still possible. The result of this, in this scenario, is that time is felt to be bringing the event progressively nearer (Griffiths & Tenenbaum, 2006; McGuire & Kable, 2013). Even when the delay has exceeded the original prediction, one infers that there must be only a short time still remaining (bottom-left panel of Figure 29.2).

However, not all events are as reliable as the rush-hour subway train. Other situations involve a different form of temporal uncertainty, which leads to a qualitatively different pattern of dynamic change in one's predictions. Suppose that instead of waiting for the subway one is waiting for a reply to an e-mail. Normally distributed expectations would no longer be appropriate, since e-mail reply times follow something more like the power-law distribution shown on the right-hand side of Figure 29.2 (Barabási,

2005). This distribution represents a mixed and open-ended set of possibilities: Many replies come quickly, but others can take a very long time. As in the previous scenario, one might initially predict a delay of 3 minutes based on the average of the distribution. Here, however, one's prediction would increase rather than decrease as time elapsed. As shorter delays were ruled out, longer and longer delays would remain possible, and the awaited event would be felt to be receding into the future (Griffiths & Tenenbaum, 2006; McGuire & Kable, 2013; Rachlin, 2000). This would be like inferring that someone who has not written back immediately is probably away from e-mail and will take hours rather than minutes to respond.

It might seem counterintuitive that decision makers would change their predictions in the absence of an obvious change in their situation. The previous example illustrates, however, that time passage itself can constitute a source of novel and relevant information. If one initially believed there was some chance that the event could happen quickly, its nonoccurrence should lead to revised expectations in light of what possibilities remain.

HIGH-PERSISTENCE AND LIMITED-PERSISTENCE ENVIRONMENTS

When a reward's predicted arrival time changes, its subjective value changes as well. Decision makers discount future rewards according to their delay. Although the exact shape that best describes the discounting function is a matter of debate, it is well established that when rewards are at stake, delay counts as a cost. More temporally remote rewards are worth less in terms of their present subjective value to the decision maker than otherwise equivalent rewards that are available sooner (Frederick, Loewenstein, & O'Donoghue, 2002). There is nothing irrational about treating delay as a cost, especially in situations such as the delay-of-gratification task, in which waiting prevents the decision maker from doing anything else. A cost-benefit assessment of the future reward should pit the expected gain (the margin by which one prefers the large reward to the small reward) against the expected opportunity cost of time (the value of whatever else one could be doing during the anticipated delay).

The large reward's attractiveness should therefore vary inversely with the predicted remaining wait time. When rewards have normally distributed timing and are felt to be approaching, like the subway train, their subjective value to the decision maker should increase as the wait progresses (Figure 29.2, left). If one preferred to wait when the predicted delay was 3 minutes, this preference should be even stronger when the predicted delay is only 1 minute. Environments with normally distributed timing statistics are conducive to high persistence. A decision maker who chose to wait at the outset would not be expected to want to reverse this decision later.

The opposite is true for events with power-law timing statistics, like e-mail responses (Figure 29.2, right). As time passes and the predicted delay grows longer, the subjective value of the delayed reward progressively deteriorates. Eventually it might be worth less to the decision maker than the immediate alternative, and at that point the value-maximizing course of action would be to give up. The fact that one was willing to wait 3 minutes for a particular delayed reward does not necessarily imply one is also willing to wait 10 minutes. Environments with power-law timing favor limited persistence: A decision maker who was willing to wait at first would have some giving-up time, beyond which the expected costs of waiting outweigh the expected benefits. In these

circumstances, a reversing pattern of choices (first waiting, then quitting) can be consistent with rational, value-guided behavior.

Effective self-regulation in delay-of-gratification scenarios is therefore considerably more complex than merely finding a way to resist the temptation to quit. Decision makers must continually gauge whether the reward is still worth waiting for, given their internal model of the statistics of the environment, and calibrate persistence accordingly. Because we all find ourselves sometimes in high-persistence environments and sometimes in limited-persistence environments, it is important that the mental processes supporting this behavior be flexible and context-sensitive.

Actual waiting behavior might therefore depend not merely on how much self-control one has, but more importantly on what kind of environment one thinks one is in. Because adaptive behavior varies across situations, it is possible to go wrong either by waiting too long or too little. For example, a decision maker might actually be in a high-persistence environment but incorrectly believe that a more limited degree of persistence is appropriate (Figure 29.3, top). To an observer who knows more about the environment it might look like the decision maker has self-control problems and lacks the capacity to delay gratification. Conversely, a decision maker might actually be in a limited-persistence environment but incorrectly think that high persistence is warranted (Figure 29.3, bottom). A better-informed observer might notice that the person is waiting too long and assume he or she is irrationally unwilling to abandon sunk costs (Arkes & Blumer, 1985; Navarro & Fantino, 2009).

THE UBIQUITY OF LIMITED-PERSISTENCE ENVIRONMENTS

Limited-persistence environments are ubiquitous. We routinely encounter delayed outcomes that are worth waiting for up to a point, but not indefinitely. One reason for this is that the ecological time intervals in many domains of human activities are well




FIGURE 29.3. Two ways a decision maker's beliefs can be miscalibrated to the environment.

characterized by power-law distributions (Anderson & Schooler, 1991; Barabási, 2005). Power-law timing represents situations in which there is a mixture of possibilities, with events able to happen on either a fast or slow time scale. Because this type of distribution allows for multiple time scales, it also makes a reasonable default assumption in situations where one initially does not know what time scale is applicable (Caves, 2000; Griffiths & Tenenbaum, 2006; Sozou, 1998). As a delay drags on, a decision maker can begin to infer that a slower time scale is more probable.

Data on people's explicit judgments lend support to the idea that limits on persistence could have their basis in temporal expectations. In several relevant domains people express temporal beliefs that resemble a power-law distribution. This can be assessed by asking people to predict how long something will take given that it has taken a certain amount of time already (Griffiths & Tenenbaum, 2006). For open-ended situations such as waiting on hold, people's predictions imply that the longer they have waited already, the more time they think is left. A similar pattern also holds for situations that are usually perceived as presenting self-control challenges, such as waiting for a diet to show positive results. When a diet had shown no results after 5 days, people predicted it would take about 10 more days, but if it had already been 15 days, people predicted that another 15 days would be required (see Figure 29.4). A similar pattern occurred when people were asked to make a prediction about a scenario modeled after the original delay-of-gratification experimental paradigm. At first they predicted that the experimenter would only take a few minutes to return and deliver the large reward, but by the time 10 minutes had elapsed, they predicted an additional 10-minute wait (McGuire & Kable, 2013). These findings support the idea that rewards can grow subjectively more remote as a function of elapsed time, implying that at some point they would cease to be worth waiting for.

ADAPTIVE CALIBRATION OF PERSISTENCE

If persistence is governed in a rational manner, this predicts that people ought to be able to adapt their behavior to contexts in which either high or low persistence is more beneficial. The hypothesis of adaptive calibration is consistent with previous demonstrations

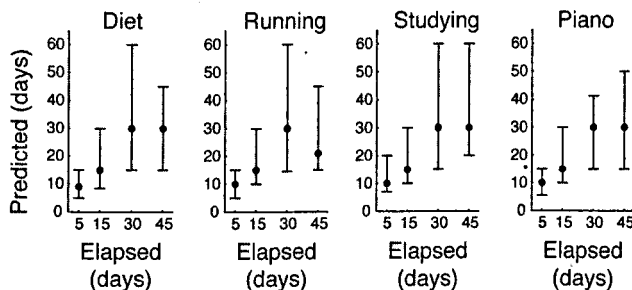


FIGURE 29.4. Explicit temporal predictions for everyday self-control dilemmas. Points show median predictions, and bars show interquartile range ($n = 256$). From McGuire and Kable (2013). Copyright 2013 by the American Psychological Association. Adapted by permission.

that environments differ in how much they reward future-oriented behavior (Fawcett, McNamara, & Houston, 2012; Otto, Markman, & Love, 2012), and that decision makers can flexibly adapt to environmental structure (Schweighofer et al., 2006; Simen et al., 2009). It stands in contrast to the alternative possibility that persistence is primarily dictated by a more rigid factor, such as a limitation on self-control capacity.

We tested the adaptive calibration hypothesis by giving participants direct experience with high-persistence and limited-persistence environments in a laboratory task (McGuire & Kable, 2012). The basic structure of the task was the same across both environments: Participants had repeated opportunities to wait for a monetary reward (e.g., \$0.15), which would be delivered at the end of an unspecified and randomly determined delay. Their alternative option was to take a smaller monetary reward immediately instead of waiting (e.g., \$0.01). Either the large or small reward was followed, a moment later, by a new trial, with a new opportunity to wait. Participants had a fixed period of 10 minutes to play and could obtain as many new trials as time allowed. Their goal was to make as much money as possible in the available time.

The most beneficial strategy depended on the probability distribution from which the random delays were sampled. In the high-persistence environment, the delay had a uniform probability of taking on any duration between 0 and 12 seconds. Like a normal distribution, this implied that the expected remaining delay steadily decreased throughout the trial. The reward-maximizing strategy was to wait as long as necessary. In the limited-persistence environment, conversely, the delay duration was sampled from a power-law distribution that encompassed a more heterogeneous set of possibilities. Many of the delays were very short (the fastest quartile of trials took less than 1 second), but others were very long (the slowest quartile took longer than 15 seconds, and up to 90 seconds). The reward-maximizing strategy was to wait for each trial at first but to quit if the reward were not delivered in about 2 seconds, at which point one could infer that the current delay would probably fall in the long tail of the distribution. The most advantageous behavior in the limited-persistence environment entailed the same reversing sequence of choices that is sometimes viewed as emblematic of delay-of-gratification failure: waiting at first but then quitting partway through the delay.

Participants adapted their behavior successfully to the two timing contexts. In a between-subject design, two groups received identical instructions but then encountered different distributions of delay intervals. Waiting times in the two groups gradually diverged as a function of experience (see Figure 29.5). Individuals who encountered high-persistence timing statistics increased their willingness to wait, and those who experienced limited-persistence timing statistics did the opposite. Performance in both conditions still fell short of optimality on average, and individuals in the limited-persistence environment tended to be less successful (as indexed by their total monetary earnings) than those in the high-persistence condition.

These empirical results make clear that high persistence is counterproductive in certain situations, and that people recognize and adapt to such situations on the basis of direct experience. At the same time, the results also imply that waiting can be easy in contexts in which one has reason to expect that time brings rewards steadily nearer. The outcomes of delay-of-gratification decisions may be governed not just by how long the decision maker is capable of waiting but by the way in which the decision maker perceives the structure of the environment. Accurately representing environmental structure would presumably be even more difficult in situations—unlike this laboratory task—in which




FIGURE 29.5. Modulation of persistence in response to direct experience with different time-interval statistics. Arrows mark ideal performance. From McGuire and Kable (2012). Copyright 2012 by Elsevier B. V. Adapted by permission.

the decision maker lacked the opportunity for experience-based learning in the relevant domain.

ERRING THROUGH OVERPERSISTENCE

The empirical finding that people performed less successfully in an environment that required less waiting (McGuire & Kable, 2012) is counterintuitive according to the classical view that waiting is difficult. There are many other examples of situations in which decision makers err on the side of persisting excessively. Previous work has examined factors that determine how much time people spend unproductively trying to solve difficult or impossible puzzles (McFarlin, Baumeister, & Blascovich, 1984), or how long people persist at an unpleasant activity after compensating incentives have been removed (Halkjelsvik & Rise, 2015). Manipulations designed to induce self-regulatory depletion can sometimes cause people to stick with a pointless task longer, especially if quitting requires a proactive response (Baumeister, Bratslavsky, Muraven, & Tice, 1998). Overpersistence in all these cases is seen as signaling a failure to act in line with well-calibrated cost–benefit assessments.

Overpersistence is also a risk in ecological foraging scenarios. Consider the situation of a predator, such as an owl that hunts by sitting and waiting for prey to come into view. A sit-and-wait predator needs to be able to wait patiently but also should not wait forever. After some period of time with no results, the best thing to do is give up and go forage in a new place. Similarly, an animal whose food is distributed in patches needs to decide when to stop harvesting from the current patch and move on in search of a fresh patch. Research in behavioral ecology has examined in detail how the ideal giving-up time depends on the structure and statistics of the foraging environment (Charnov, 1976; McNamara, 1982; Nishimura, 1991).

Animals in behavioral foraging experiments typically appear well calibrated to their environment, but sometimes they wait too long. In one example, the behavior of macaque monkeys in a patch-foraging task broadly corresponded to the predictions of optimal foraging theory (Blanchard & Hayden, 2015; Hayden, Pearson, & Platt, 2011). The monkeys would repeatedly harvest from the same patch for a period of time amid diminishing

returns, and would eventually depart for a new patch. However, there was a systematic tendency to overharvest; that is, to persist with diminishing returns from the current patch slightly past the point at which it would have been more beneficial to move on (Blanchard & Hayden, 2015; Hayden et al., 2011). Similar phenomena have been documented in species ranging from rats to humans: Compared to an optimal model, decision makers show excessive willingness to tolerate long delays (Constantino & Daw, 2015; Wikenheiser, Stephens, & Redish, 2013).

The specific explanation for overpersistence in each of these cases is up for debate. One possibility is that decision makers experience an extra subjective cost related to abandoning sunk costs or deviating from the status quo. Overpersistence might also arise from misestimating the environmental statistics. For example, if one underestimated how rich the environment was, one might fail to realize how large an opportunity cost was entailed by long delays. Conversely, if one underestimated or neglected the possibility that the delay could turn out to be far longer than originally predicted, this too could make persistence seem unduly beneficial from the decision maker's perspective.

Regardless of which specific explanation fits best, the larger message from these findings is that overpersistence is a commonplace maladaptive feature of behavior. The general challenge in delay-of-gratification situations is to strike a balance, pursuing delayed rewards as long as they remain worth waiting for, but not longer. Because effective self-regulation requires curtailing persistence when necessary, both insufficient and excessive persistence deserve to be considered as important and ecologically relevant forms of self-control failure (Baumeister & Scher, 1988).

ADAPTIVE CALIBRATION AND RULE-GOVERNED SELF-CONTROL

The need to adapt flexibly to incoming information is fundamentally at odds with certain self-control strategies that involve attempting to lock in or constrain one's own future behavior. One of the most frequently discussed forms of self-control is *precommitment* (Ainslie, 1975; Schelling, 1984; Strotz, 1955), which refers to deliberately constraining one's own options in the future, for instance, by handing one's car keys to a friend or freezing one's credit cards in blocks of ice. The reason to take such measures would be to frustrate one's own future preferences in the event these disagree with the preferences one holds now.

Although some precommitment devices involve setting up physical constraints in the external environment, such devices can also be purely internal. Mental tricks such as reframing incentives or shifting the focus of attention can help ensure that one's current decision will be adhered to in the future (Ainslie, 1975; Duckworth et al., 2014). Perhaps the most archetypal form of self-control involves setting up a mental rule—for instance, “keep waiting until the delayed reward arrives”—and rigidly adhering to it even at moments when an alternative course of action begins to seem more valuable. This is one possible way of thinking about dual-system self-control models in which a “cool” system overrides a “hot” system (Metcalf & Mischel, 1999).

Rule-governed behavior is not without its costs. Ainslie (1975, p. 476) made this observation in an early discussion of mental precommitment: “Unpredictability of reward reduces the differential reward for rigid behavior and would be expected to reduce preference for precommitting devices.” Precommitment would be inappropriate in temporally

uncertain delay-of-gratification scenarios (McGuire & Kable, 2012, 2013). These are situations in which the passage of time furnishes new and valid information about the relative value of waiting. Decision makers in these kinds of situations might choose to wait at first, but should not commit to waiting indefinitely. Rather than locking in their initial preferences, they would be best served by reserving the ability to revisit those preferences later in light of updated information.

Our main focus in this chapter has been on changing predictions of remaining delay time, but this is only one of many factors that can drive a dynamic reappraisal of future outcomes. Another such factor is risk: A long delay might alter one's prediction of whether an awaited outcome will occur at all, or of how good it will be. Shortly after a job interview, one might have high confidence in a favorable result, but this confidence might wane after 2 weeks without being contacted. If one's car fails to start on the first attempt, one might try a few more times, but a few more failures will suggest it may be better to stop trying and look for an alternative solution. It is also possible for the attractiveness of future rewards to be altered by tangible feedback beyond the mere passage of time. One might intend to carry a hot lasagna from the kitchen all the way to the dining room but justifiably change plans after the first few seconds. This would not be a self-control failure; it would reflect that in those few seconds, one has learned a great deal about the costs this course of action would entail. In a world that is both changeable and uncertain, well-regulated behavior requires allowing plans to be revised in response to incoming information.

The reason that precommitment strategies are sometimes valuable is that the uncertainty or changeability of the environment needs to be balanced against the potential for erratic variability in one's own future actions. The situations in which rule-governed self-control is most clearly beneficial are those in which the environment is stable and predictable but one's own behavior at the decisive moment is liable to be based on incomplete information. For example, one might place the alarm clock across the room because one does not trust oneself at 6:00 A.M. to give complete consideration to the costs and benefits of continuing to sleep.

A hypothesis worth investigating in the future, therefore, is that people's degree of inclination to precommit and rigidly adhere to predetermined plans should vary adaptively from one context to another. This perspective suggests that self-control is not as simple as overriding impulses by adhering to preestablished rules, but instead requires a more flexible set of cost-benefit assessments (Casey, 2015; Phelps, Lempert, & Sokol-Hessner, 2014). Key variables in these assessments would include both the instability/uncertainty of the environment and the potential for non-value-guided variability in one's own future behavior. Given the omnipresence of uncertainty in real-world decision environments, it is likely that the first of these variables is often the dominant factor, militating against rigidly precommitted decisions. The most significant challenge of self-control might not be to close off nonpreferred alternatives, but rather to leave all alternatives open while reevaluating them in a dynamic and context-sensitive manner.

CONCLUSIONS

Effective self-regulation requires dynamically adjusting one's subjective valuation of choice alternatives in response to new information. Even in situations such as delay of

gratification, which at first glance may look like a static self-control dilemma, value estimates can change over time due to changing assessments of temporal uncertainty. The passage of time can progressively degrade the subjective value of an awaited future reward because in some environments, a long delay so far supports an inference that the delay still remaining is longer than expected. This means one can err through either insufficient or excessive persistence, and either type of error can be explained in terms of a mismatch between one's expectations and the environment. The need to base decisions on dynamically updated inferences puts important limits on the adaptive usefulness of self-control strategies such as precommitment and rule adherence.

So what is the relevance of these ideas to domains in which people chronically exhibit self-control problems or delay-of-gratification failure, such as dieting or substance abuse? First, we should not overlook the possibility that people sometimes correctly infer, as time passes, that the future reward is more remote than they originally thought it was. This could very well be the case in dieting, as most diets are not effective in the long term (Mann et al., 2007). In situations where choosers are clearly miscalibrated (Figure 29.3, top), persistence could be promoted by providing information and/or allowing for learning from accurate feedback. The installation of waiting-time clocks in subway systems is one case in which clear temporal information can transform a situation in which waiting is difficult into one in which waiting is easy. In still other cases, reversals seem more clearly related to variability in the state of the chooser, for example, the smoker who relapses during withdrawal, and these might more clearly be openings for precommitment and rule-based strategies. Beneficial change in persistence might be best achieved through tailored, situation-specific interventions that acknowledge the complexity of calibrating behavior to diverse decision environments.

REFERENCES

- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, *82*, 463–496.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.
- Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, *35*, 124–140.
- Barabási, A.-L. (2005). The origin of bursts and heavy tails in human dynamics. *Nature*, *435*, 207–211.
- Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology*, *74*, 1252–1265.
- Baumeister, R. F., & Scher, S. J. (1988). Self-defeating behavior patterns among normal individuals: Review and analysis of common self-destructive tendencies. *Psychological Bulletin*, *104*, 3–22.
- Blanchard, T. C., & Hayden, B. Y. (2015). Monkeys are more patient in a foraging task than in a standard intertemporal choice task. *PLoS ONE*, *10*, e0117057.
- Casey, B. J. (2015). Beyond simple models of self-control to circuit-based accounts of adolescent behavior. *Annual Review of Psychology*, *66*, 295–319.
- Caves, C. M. (2000). Predicting future duration from present age: A critical assessment. *Contemporary Physics*, *41*, 143–153.
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, *9*, 129–136.

- Constantino, S. M., & Daw, N. D. (2015). Learning the opportunity cost of time in a patch-foraging task. *Cognitive, Affective, and Behavioral Neuroscience*, *15*, 837–853.
- Duckworth, A. L., Gendler, T. S., & Gross, J. J. (2014). Self-control in school-age children. *Educational Psychologist*, *49*, 199–217.
- Eigsti, I.-M., Zayas, V., Mischel, W., Shoda, Y., Ayduk, O., Dadlani, M. B., et al. (2006). Predicting cognitive control from preschool to late adolescence and young adulthood. *Psychological Science*, *17*, 478–484.
- Fawcett, T. W., McNamara, J. M., & Houston, A. I. (2012). When is it adaptive to be patient?: A general framework for evaluating delayed rewards. *Behavioural Processes*, *89*, 128–136.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *40*, 351–401.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*, 767–773.
- Halkjelsvik, T., & Rise, J. (2015). Persistence motives in irrational decisions to complete a boring task. *Personality and Social Psychology Bulletin*, *41*, 90–102.
- Hayden, B. Y., Pearson, J. M., & Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, *14*, 933–939.
- Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition*, *126*, 109–114.
- Mahrer, A. R. (1956). The role of expectancy in delayed reinforcement. *Journal of Experimental Psychology*, *52*, 101–106.
- Mann, T., Tomiyama, A. J., Westling, E., Lew, A.-M., Samuels, B., & Chatman, J. (2007). Medicare's search for effective obesity treatments: Diets are not the answer. *American Psychologist*, *62*, 220–233.
- McFarlin, D. B., Baumeister, R. F., & Blascovich, J. (1984). On knowing when to quit: Task failure, self-esteem, advice, and nonproductive persistence. *Journal of Personality*, *52*, 138–155.
- McGuire, J. T., & Kable, J. W. (2012). Decision makers calibrate behavioral persistence on the basis of time-interval experience. *Cognition*, *124*, 216–226.
- McGuire, J. T., & Kable, J. W. (2013). Rational temporal predictions can underlie apparent failures to delay gratification. *Psychological Review*, *120*, 395–410.
- McNamara, J. (1982). Optimal patch use in a stochastic environment. *Theoretical Population Biology*, *21*, 269–288.
- Metcalf, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, *106*, 3–19.
- Mischel, W., Ayduk, O., & Mendoza-Denton, R. (2003). Sustaining delay of gratification over time: A hot-cool systems perspective. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 175–200). New York: Russell Sage Foundation.
- Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*, *16*, 329–337.
- Mischel, W., Ebbesen, E. B., & Zeiss, A. R. (1972). Cognitive and attentional mechanisms in delay of gratification. *Journal of Personality and Social Psychology*, *21*, 204–218.
- Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, *244*, 933–938.
- National Institute of Child Health and Human Development (NICHD) Early Child Care Research Network. (2002). Early child care and children's development prior to school entry: Results from the NICHD Study of Early Child Care. *American Educational Research Journal*, *39*, 133–164.
- Navarro, A. D., & Fantino, E. (2009). The sunk-time effect: An exploration. *Journal of Behavioral Decision Making*, *22*, 252–270.
- Nishimura, K. (1991). Optimal patch residence time of a sit-and-wait forager. *Behavioral Ecology*, *2*, 283–294.

- Otto, A. R., Markman, A. B., & Love, B. C. (2012). Taking more, now: The optimality of impulsive choice hinges on environment structure. *Social Psychological and Personality Science*, 3, 131–138.
- Phelps, E. A., Lempert, K. M., & Sokol-Hessner, P. (2014). Emotion and decision making: Multiple modulatory neural circuits. *Annual Review of Neuroscience*, 37, 263–287.
- Rachlin, H. (2000). *The science of self-control*. Cambridge, MA: Harvard University Press.
- Rodriguez, M. L., Mischel, W., & Shoda, Y. (1989). Cognitive person variables in the delay of gratification of older children at risk. *Journal of Personality and Social Psychology*, 57, 358–367.
- Schelling, T. C. (1984). Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review*, 74, 1–11.
- Schweighofer, N., Shishida, K., Han, C. E., Okamoto, Y., Tanaka, S. C., Yamawaki, S., et al. (2006). Humans can adopt optimal discounting strategy under real-time constraints. *PLoS Computational Biology*, 2, e152.
- Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 1865–1897.
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society B: Biological Sciences*, 265, 2015–2020.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23, 165–180.
- Wikenheiser, A. M., Stephens, D. W., & Redish, A. D. (2013). Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proceedings of the National Academy of Sciences*, 110, 8308–8313.