Taming Uncertainty

Ralph Hertwig, Timothy J. Pleskac, Thorsten Pachur,
and The Center for Adaptive Rationality

The MIT Press
Cambridge, Massachusetts
London, England

9 Tomorrow Never Knows: Why and How Uncertainty Matters in Intertemporal Choice

Junyi Dai, Thorsten Pachur, Timothy J. Pleskac, and Ralph Hertwig

9.1 The Pervasiveness of Uncertainty in Intertemporal Choice

In the 1960s, Walter Mischel and his graduate students looked for a way to measure the development of self-control among preschoolers. In the now famous Marshmallow Test, they offered children a choice of either eating a single treat straight away or waiting until the experimenter returned from a brief errand and then being rewarded with a second treat (Mischel, 2014). This is perhaps the best-known example of an *intertemporal choice*—a decision between options whose outcomes materialize at different times. What is not widely appreciated is that some of Mischel's experiments contained an element of uncertainty: children did not always know how long they would have to wait. The experimenter simply said, "you know, sometimes, I'm gone a long time" (Mischel & Ebbesen, 1970, p. 332).

Uncertainty is an inherent property of intertemporal choice in many real-world situations (Frederick, Loewenstein, & O'Donoghue, 2002). Imagine you have a lump sum of money that you can either spend immediately or invest in the stock market in the hope of reaping benefits later. You face *outcome uncertainty* about whether an investment will turn out to be profitable and about how much profit will be made, and *temporal uncertainty* regarding the length of time it will take to reap the benefits. Or imagine you are considering changing your eating habits. You will have to make the decision without knowing when you can expect to notice results, and without being sure that the results will be worth the effort. Despite the important role that uncertainty plays in many intertemporal choices, most experimental studies of these decisions try to factor it out. Often this is done by asking people to assume that the outcomes of any chosen option would be sure to materialize (for a review, see Frederick et al., 2002; Urminsky & Zauberman, 2015). For instance, a person might be asked to choose between receiving €100 now and receiving €200 in a year, with both outcomes guaranteed. Relative to the extent of the work on intertemporal choice in both economics and psychology over the past century, discussions of the possible role of uncertainty in these decisions have been quite rare. There are, of course, exceptions (see, e.g., Benzion, Rapoport, & Yagil, 1989; Kagel, Green, & Caraco, 1986; Keren & Roelofsma, 1995; Mischel & Grusec, 1967; Stevenson, 1986; B. J. Weber & Chapman, 2007). Yet, of the 10 most cited articles on intertemporal choice on Google Scholar as of July 2017, the majority do not mention the issue of uncertainty at all, and the others either treat it only tangentially or list it as one of many possible factors that might influence people's decisions (e.g., Frederick et al., 2002).

Our goal in this chapter is to zoom in on the role that uncertainty might play in intertemporal choice and to show that uncertainty could be a key factor shaping people's behavior in these choices. To do so, we highlight the various ways in which uncertainty is relevant in intertemporal choice. Table 9.1 provides an overview of the different types of uncertainty we consider; the first three types are specific cases of outcome uncertainty, whereas the last relates to temporal uncertainty. One important insight is that key regularities in intertemporal choice could be understood as adaptive responses to uncertainty. Taking uncertainty into consideration could thus impact how many well-known findings are interpreted. Moreover, we suggest that

Table 9.1
Types of uncertainty in intertemporal choice with examples.

Type of uncertainty	Examples for the choice "Should I start exercising now to be healthier in the future?"
Uncertainty in the materialization of the future outcome	Will the decision definitely lead to better health? If not, how likely is it that I will end up healthier?
Uncertainty in the size of the future outcome	How much healthier will I be if I start exercising?
Uncertainty in the subjective value of the future outcome (i.e., utility uncertainty)	How much benefit will I have from better health?
Uncertainty in the delay until the future outcome materializes	When will I actually attain better health?

decisions from experience, which have been employed to study the impact of uncertainty in risky choice (see chapters 7 and 8), can provide a helpful methodological framework for understanding intertemporal choice under uncertainty. We conclude the chapter with such an illustration.

9.2 The Many Shades of Uncertainty in Intertemporal Choice

Three of the most prominent findings in the literature on intertemporal choice are as follows: (a) people choose between payoffs that materialize at different times as though the values of delayed payoffs were discounted (Fisher, 1930); (b) the degree of discounting per unit of time tends to decrease as delays get longer (Benzion et al., 1989; Chapman, 1996; Thaler, 1981); and (c) over time, people's preference between one option with a smaller-but-sooner reward and another option with a larger-but-later reward may reverse, from preferring the latter to preferring the former (e.g., Ainslie, 1975; Green, Fristoe, & Myerson, 1994).

A common approach to accounting for these three findings is to assume a mathematical function that describes how the utility of an outcome decreases as the delay to its realization increases. These functions are often referred to as *discount functions*. One discount function that is particularly effective at capturing people's intertemporal choices is the hyperbolic function (Mazur, 1987). Figure 9.1 shows such a function and how it produces the preference reversal mentioned above. As we will show, this and other behavioral regularities captured by a hyperbolic discount function may be related to the uncertainty inherent in the prospect of future outcomes.

9.2.1 Will the Future Outcome Materialize? And How Likely Is It?

As indicated in table 9.1, one important source of uncertainty in intertemporal choice is that it is unclear whether the anticipated consequence will actually materialize. In 1965, when French lawyer André-François Raffrey was 47 years old, he offered a 90-year-old widow a deal: he would pay her 2,500 francs every month until her death, at which point he would inherit her beautiful apartment. She accepted—and went on to live another 32 years. She outlasted her unlucky beneficiary, who died after 30 years, having paid approximately 920,000 francs for an apartment he never got to live in ("A 120-year lease", 1995; Coatney, 1997). In the same vein, when the dot-com bubble burst in 2000, washing out 52% of dot-com companies

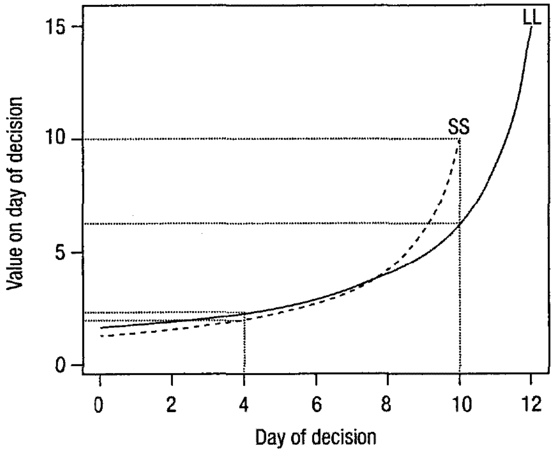

Figure 9.1
The value of two options on different days according to a hyperbolic discount function. In this example, a choice is offered between a smaller-but-sooner (SS) option that is obtained on day 10, and a larger-but-later (LL) option, where a larger payoff is obtained on day 12. Note that the SS option is not available after day 10, so it has no value on subsequent days. The figure shows how a preference reversal can occur. For example, if the choice between the SS and LL options is made on day 4, then the LL option has a greater value, but if the choice is on day 10, then the value of the SS option is greater.

by 2004 and severely affecting many of those that remained (Berlin, 2008), many stockholders lost almost all of what initially seemed to be a profitable investment.

These examples raise the question of what role uncertainty about the materialization of future outcomes might play in the intertemporal choices that people make. One answer is that outcome uncertainty might be an important reason for *delay discounting*—the apparent discounting of the value of future payoffs. In work dealing with animals, researchers have highlighted the possibility that a delayed reward may become unavailable before it is ready for collection (collection risk; Houston, Kacelnik, & McNamara, 1982). Food, for instance, may have already been consumed by a competitor by the time an animal returns to it (see also chapters 12 and 15). Similarly, the anticipated outcome of harvesting and consuming a reward might not actually be realized (e.g., because a planned retrieval is interrupted by a dangerous predator, or because the animal dies before it can return). These and other risks associated with a delayed reward might make it reasonable to find a delayed reward less attractive than an immediate reward of the same—or even smaller—magnitude (e.g., Green & Myerson, 1996; Stephens, 2002; Stevens, 2010). In other words, delay discounting could represent an adaptive response to the inherent uncertainty about whether or not a delayed option will actually materialize (see also Wendt & Czaczkes, 2017).

The degree of delay discounting may be directly related to the degree of outcome uncertainty in the environment. M. Wilson and Daly (1997) compared 77 Chicago neighborhoods—including both poor, crime-ridden communities and wealthy, safe areas—in terms of life expectancy at birth and homicide rate. In estimating life expectancy in each community, the authors removed the contribution of deaths due to homicide. Although the direct effect of homicide on mortality was thus controlled for, the adjusted life expectancy was still strongly correlated with the homicide rate across communities. The authors proposed that low life expectancy in a neighborhood may create high uncertainty about surviving long enough to reap future benefits; this uncertainty might have led to stronger delay discounting and, in turn, more risk taking in social competition, sometimes resulting in loss of life.

Empirical support for the effect of outcome uncertainty experienced in the environment on intertemporal choices has been found in the lab. For instance, Kidd, Palmeri, and Aslin (2013) examined the role of outcome uncertainty in the Marshmallow Test. Before running the test, the experimenters presented their young participants with an art project task in which the promised art supplies or stickers were either provided or not, thus creating different degrees of uncertainty about the materialization of promised rewards. In the subsequent Marshmallow Test, children who had not obtained the promised objects in the art project (unreliable and thus uncertain environment) task were less willing to wait to get more marshmallows than were children who had obtained the promised objects (reliable and thus more certain environment).

A subtler form of outcome uncertainty is that decision makers might be uncertain not only about whether an anticipated outcome will materialize, but also about the likelihood that it will materialize. This type of uncertainty provides an interesting perspective on the enduring debate on the shape of the delay discount function in intertemporal choice (figure 9.1).

Recall that this function describes how the subjective valuation of an outcome changes as a function of the length of the delay. In figure 9.1, the discount rate per unit of time declines over time. Therefore, a person choosing between a smaller-but-sooner (SS) option and a larger-but-later (LL) option on day 10 would prefer the SS option, but the same person choosing on day 4 (and thus experiencing a longer delay before receiving either option) would prefer the LL option (see also box 9.1).

Yet the discount rate per unit of time is often assumed to be constant—that is, not dependent on the length of the delay. This assumption is formalized in the exponential discount function, which has been proposed in the context of the discounted utility model (Samuelson, 1937). Building on this model, Strotz (1955) showed that an exponential discount function can be normatively defensible if (a) a delayed option could become unavailable before the due date; (b) the probability (per unit of time) of this happening given that it has not occurred yet—also known as *hazard rate*—is constant over time; and (c) the hazard rate is known to the decision maker at the time a choice is made. Like other features of the environment, however, the hazard rate is typically unknown to the decision maker. For instance, there is usually no way for a foraging animal to assess the exact number of competitors present in the same area, or the chance of being devoured in the coming hours. The animal might therefore continuously update its belief about the hazard rate, depending on whether a food option remains available. If, as time passes, a delayed option remains available, the objective hazard rate might be rather low. Consequently, it is reasonable for the animal to update its belief such that a relatively low hazard rate is perceived as increasingly more likely. Sozou (1998) offered a formal explanation of this phenomenon, showing that when the initial (subjective) belief about the hazard rate can be expressed by an exponential distribution and this belief is updated using Bayes' rule, the uncertainty in the hazard rate naturally leads to a hyperbolic discount function, where the discount rate per unit of time declines with longer delays (e.g., Green, Fry, & Myerson, 1994; Mazur, 1987; Rachlin, 2006; see box 9.1).

Sozou (1998) showed that a hyperbolic discount function is also normatively defensible in an environment where the hazard rate is exponentially distributed and if one assumes that the overall discount function results from averaging across all possible hazard rates. Under these conditions, it may be appropriate to discount outcomes according to a hyperbolic discount

Box 9.1

Delay discounting: Exponential or hyperbolic?

Research on intertemporal choice has been strongly influenced by the concept of delay discounting, according to which the subjective valuation (or utility) of a payoff x declines as it is delayed into the future. This concept has set the foundation for various delay discounting models on intertemporal decisions, including the discounted utility model, which economists have long held to be the normative model. According to the discounted utility model, the subjective valuation of a payoff declines as a function of delay duration at a constant rate. This relationship can be captured by an exponential discount function:

$$D(t) = \exp(-kt), \tag{B1}$$

where t represents the delay duration and k is a parameter for the constant discount rate. The current subjective valuation of a payoff is then:

$$V = D(t) \times u(x), \tag{B2}$$

where $u(x)$ is the utility of the payoff.

Psychologists, however, have found that a hyperbolic discount function can provide a better description of the empirical data than the exponential

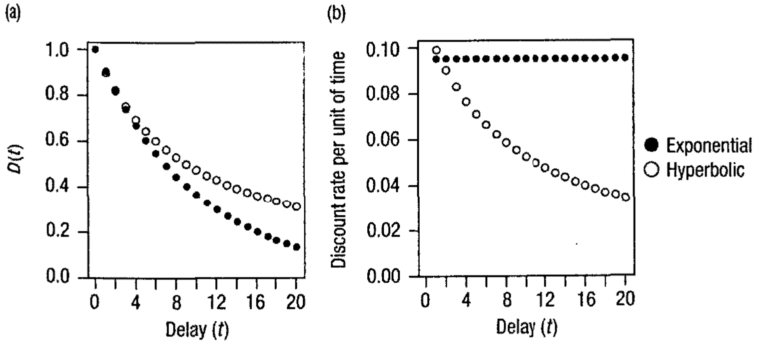

Figure 9.B1

Examples of (a) exponential and hyperbolic discount functions and (b) the corresponding discount rate per unit of time. Although the two functions shown here have similar discount rates per unit of time for shorter delays, the discount rate per unit of time decreases according to the hyperbolic discount function but stays constant according to the exponential discount function. As a result, as delay gets longer, the hyperbolic discount function entails a lower overall discount rate than the exponential discount function does.

(continued)

Box 9.1 (continued)

function (for both humans and animals). This function suggests a decreasing discount rate per unit of time as the delay gets longer. In other words, postponing an immediate payoff by a certain period of time has a larger impact on its subjective valuation than postponing it further into the future by the same period of time. A number of hyperbolic discount functions have been proposed to capture this property; the simplest and presumably most commonly used one is

$$D(t) = \frac{1}{1 + kt}, \tag{B3}$$

where *t* again represents the delay duration and *k* is a discount rate parameter (Mazur, 1987). Figure 9.B1 shows an example of both discount functions and the corresponding discount rate per unit of time as the delay gets longer.

function. This result, which holds if the actual hazard rate is uncertain, is fundamentally different from the normative solution under certainty (i.e., if the hazard rate is fixed and known), where an exponential discount function with a constant discount rate per unit of time is more appropriate. Such a function guarantees consistent preferences between options at different time points as time elapses (i.e., no preference reversals).

In sum, the behavioral regularities that are captured by a hyperbolic discount function may be due to uncertainty. If there is uncertainty about the materialization of outcomes, hyperbolic discounting is more appropriate, but under conditions of certainty, exponential discounting is more suitable.

9.2.2 How Large Is the Future Outcome?

Another source of uncertainty in intertemporal choice relates to how *much* the decision maker will benefit from a delayed option. MBA students at the start of their program will find it difficult to accurately predict the salaries they will be offered after graduating (especially if their studies span an economic crisis). The impact of such uncertainty might be manifested in one of two ways. On the one hand, if people are averse to uncertainty about the magnitude of future outcomes (as they are to ambiguity in the odds of obtaining a virtually immediate outcome; Ellsberg, 1961), that uncertainty may make a delayed option less attractive. On the other hand, delay could encourage people to take an optimistic view of uncertain magnitudes, thereby making an uncertain future outcome appear more attractive (e.g., Onay, La-Ornual, & Öncüler, 2013). There may also be individual differences in how delay impacts the level of optimism toward an uncertain outcome, with optimists expecting an uncertain future outcome to turn out well, and pessimists expecting it to turn out badly. Such individual differences could give rise to opposing preferences among delayed outcomes.

9.2.3 What Will It Be Worth to Me?

Even if the magnitude of a future outcome is known, people making an intertemporal choice might still face *utility uncertainty*—that is, uncertainty about how much utility, or personal pleasure, they will derive from the outcome (Loomes, Orr, & Sudgen, 2009). For example, an immediate reward of €100 can mean an enjoyable lunch for two at that new farm-to-table restaurant or a happy afternoon with the family at a local music festival—but years of inflation could reduce the same €100 to the value of a mediocre buffet meal or a pair of movie tickets. A person's social and economic status may also change over time, turning an exciting outcome such as an annual salary of €100,000 into something less impressive several years later. On an emotional level, most people would find it painful to have to wait for a desirable future outcome, but only a rare few could anticipate precisely how painful the waiting process will actually be; people may also overestimate the emotional impact of a future outcome (error in affective forecasting; e.g., T. D. Wilson & Gilbert, 2003) or wrongly estimate its experienced utility (e.g., Kahneman & Thaler, 2006). All these factors make it difficult to evaluate a future outcome, and might thus influence an intertemporal decision. An important topic for future research will be to test whether adaptive responses to utility uncertainty are actually at the root of behaviors such as the failure to delay gratification.

9.2.4 How Long Will I Have to Wait?

All the types of uncertainty we have considered so far refer to the future outcome of an option. But the length of delay until an expected outcome is realized can also be uncertain. As a result, people do not generally know *when* future outcomes will materialize (McGuire & Kable, 2013). This type of uncertainty is usually referred to as temporal uncertainty (as opposed to outcome uncertainty). One way of dealing with temporal uncertainty is to update one's beliefs about the possible delays as time passes without

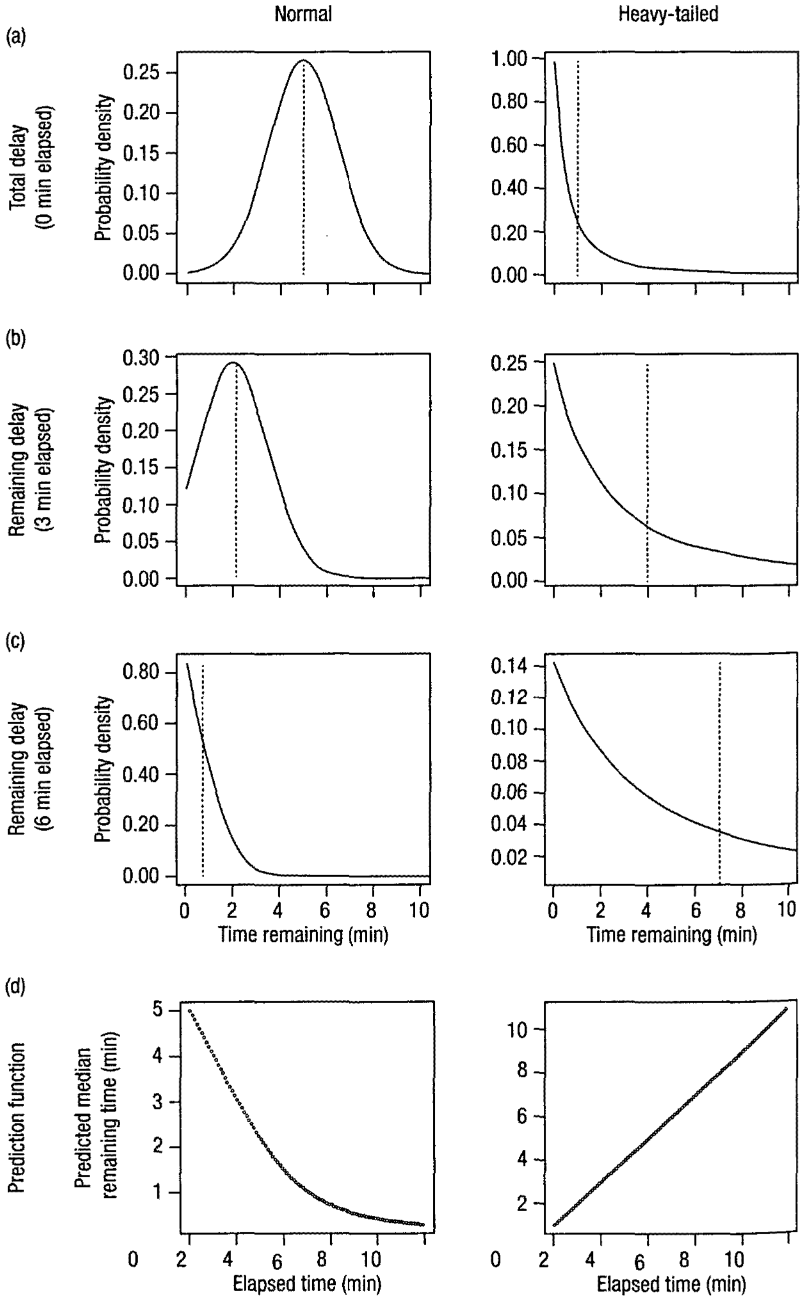
Normal
Heavy-tailed
(a)
Total delay
(0 min elapsed)
Probability density
0.25
0.20
0.15
0.10
0.05
0.00
1.00
0.80
0.60
0.40
0.20
0.00
(b)
Remaining delay
(3 min elapsed)
Probability density
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0.25
0.20
0.15
0.10
0.05
0.00
(c)
Remaining delay
(6 min elapsed)
Probability density
0.80
0.60
0.40
0.20
0.00
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0 2 4 6 8 10
Time remaining (min)
0 2 4 6 8 10
Time remaining (min)
(d)
Prediction function
Predicted median
remaining time (min)
5
4
3
2
1
10
8
6
4
2
0 2 4 6 8 10
Elapsed time (min)
0 2 4 6 8 10
Elapsed time (min)

the anticipated outcome materializing (instead of one's beliefs about the hazard rate under outcome uncertainty). This updating process might lead an individual to switch from preferring a larger payoff with an uncertain delay to preferring a smaller payoff that has been available all along (McGuire & Kable, 2012, 2013). This possibility provides a new perspective on Mischel's Marshmallow Test, where, as we highlighted above, the actual delay of the later reward is usually unknown to the decision maker. If the children decide whether to eat their treat based on their beliefs about possible delays, individual differences in how long they are willing to wait could be due to differences in these beliefs. For example, if an initial belief about possible delays follows a normal distribution, the predicted median remaining waiting time will decrease as time passes (see left column of figure 9.2). As a consequence, the larger-but-later option becomes even more attractive than the immediately available option, making it quite reasonable for the decision maker to continue waiting. In comparison, if an initial belief about possible delays follows a heavy-tailed distribution (e.g., a generalized Pareto distribution; see right column of figure 9.2), where long delays are probable, the updated median waiting time increases as time passes. In this case, the preference should shift toward the always-available option, even though this entails revoking the initial decision. To experience these different types of delays firsthand, see interactive element 9.1 (at https://taming-uncertainty.mpib-berlin.mpg.de/). The type of adjustment people actually rely on might thus depend on the features of the environment and what people know about these features.

These insights about ways to deal with temporal uncertainty suggest yet another interpretation of behavior in the Marshmallow Test. Whereas the finding that children often prefer the single, immediate marshmallow is

Figure 9.2

Updating the predicted remaining delay assuming an initial belief with normally distributed total delays (left column) and a heavy-tailed distribution of total delays (right column). (a) The initial belief distribution (normal distribution with a mean of 5 minutes and a standard deviation of 1.5 minutes). (b, c) The distributions of remaining time after 3 and 6 minutes have passed. The dashed lines in (a), (b), and (c) show the median remaining times when 0, 3, and 6 minutes, respectively, have elapsed. (d) The prediction function, expressing the relationship between elapsed time and median remaining time. As time passes, the median predicted remaining time decreases when assuming an initial belief with normally distributed total delays, and increases when assuming a heavy-tailed distribution of total delays.

often attributed to their limited self-control, it may in fact reflect an adaptive response to temporal uncertainty (McGuire & Kable, 2012, 2013). It is possible that children who have waited for a while in vain for the experimenter to come back make the reasonable inference that it will be a very long time until they get the larger reward. By extension, rather than reflecting a developing capacity of self-control, children's increased willingness to wait may be due to them accumulating experience and developing increasingly certain beliefs about possible delay lengths in the world (see also chapter 16).

9.3 Intertemporal Decisions under Temporal Uncertainty: A Description–Experience Gap?

Even in situations where the delay of an intertemporal option cannot be predicted with certainty, it is rarely the case that nothing at all is known about the delay. For instance, the probabilities of different possible, mutually exclusive delays may be—at least approximately—known. This situation is known as *timing risk*. In a study on intertemporal choice under timing risk, Onay and Öncüler (2007) asked participants to choose between a *timing lottery* with two possible delays whose probabilities were known (e.g., receiving 160 Turkish lira either in one month with a probability of .2, or in 11 months with a probability of .8) and a *sure-timing* option, where the timescale was set (e.g., receiving 160 Turkish lira in nine months). Both options offered the same hypothetical payoff, and the expected delay of the timing lottery (i.e., the average delay, with each delay being weighted by its probability) was equal to the delay of the sure-timing option. Most delay discounting models, including the discounted utility model and the hyperbolic discounting model, assume a convex discount function, leading them to predict that the discounted expected utility of the timing lottery is generally higher than that of the sure-timing option and that the timing lottery should therefore be chosen.

Onay and Öncüler (2007) found that actual choices deviated systematically from this prediction. People tended to choose the timing lottery when the shorter delay in the lottery option was less probable than the longer delay (thus suggesting risk seeking). When the shorter delay was more probable than the longer delay, however, people tended to choose the sure-timing option (thus suggesting risk aversion). One way to account for this choice pattern is to assume that the probabilities of the delays are not treated at face value but instead impact the evaluation of an option in a nonlinear fashion—specifically, consistent with an inverse S-shaped weighting function (which is typically also found in risky choice with stated probabilities; e.g., Tversky & Kahneman, 1992). This function implies that rare delays are overweighted—they have more psychological impact than they deserve given their objective probabilities (see chapter 8).

In the Onay and Öncüler (2007) study and other studies on timing risk, the possible delays and the probabilities of those delays occurring were clearly described. Outside the lab, however, such convenient descriptions are rare. As we have mentioned in the previous section, one possible recourse is for people to draw on their own beliefs about the possible delays. Another possibility is to draw on past experience with the options. This raises the question of whether a description–experience gap like that observed in risky choice also arises in intertemporal choice (see chapters 7 and 8). If so, do factors similar to those that play a role in risky choice, such as sampling error, also operate in this context? To investigate this question, we compared intertemporal choices with timing risk, in which each possible delay and its probability of occurrence is provided as a description, with intertemporal choices with *timing uncertainty*, in which possible delays and/or their probabilities of occurrence must be learned from experience and are thus at best vaguely known (Dai, Pachur, Pleskac, & Hertwig, 2018). The design of our study was similar to that in Onay and Öncüler, but with an added experience-based condition. In this condition, people learned about the possible delays of the outcomes as well as their probabilities from sequential experiential sampling—as in research on risky choice using the sampling paradigm (e.g., Hertwig & Erev, 2009; see chapter 7). The sure-timing option always entailed the same delay (e.g., 6 months), whereas the delays in the timing lottery could vary from sample to sample (e.g., 1 vs. 11 months), depending on their probability.

Figure 9.3 shows the proportion of choices of the sure-timing option, separated according to whether the rare delay was longer (i.e., unattractive) or shorter (i.e., attractive), and according to whether the condition was experience-based (timing uncertainty) or description-based (timing risk). As can be seen, there was indeed a gap between the timing risk and timing uncertainty conditions. When a rare delay was relatively long (i.e., 11 months with a probability of .1), people preferred the sure-timing option over the

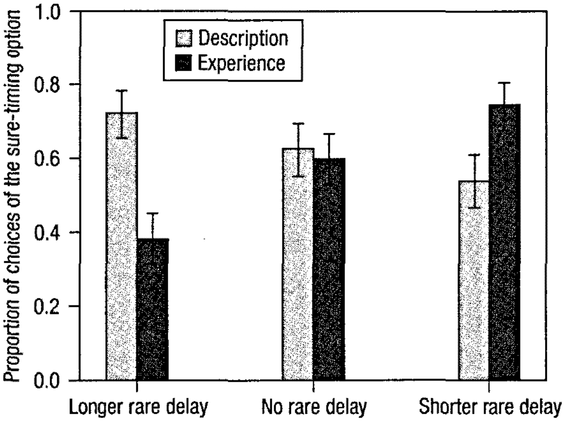

Figure 9.3
A description–experience gap in intertemporal choice. Bars show the observed choice proportions and error bars indicate 95% confidence intervals (recreated from Dai et al., 2018).

timing lottery more frequently in description than in experience. When a rare delay was relatively short (i.e., 1 month with a probability of .2), people instead preferred the sure-timing option more frequently in experience than in description. When the possible delays in the timing lottery were equally probable (and there was thus no rare event), there was no difference between the two conditions in the probability of choosing the sure-timing option. Further analyses suggested that sampling error in the experience-based condition contributed to the gap, but that difference in probability weighting (based on the experienced probabilities) also made a contribution.

9.4 Does Experiencing a Delay Always Reduce Uncertainty?

In our study (Dai et al., 2018), we implemented an experiential mode of learning to induce uncertainty about the options and to contrast this situation to a situation involving risk. Experience can, however, also decrease uncertainty. For instance, a customer deciding whether to order items using an online store's regular delivery service or to opt in for the premium, express service (at a surcharge) could order a few items with the regular service and experience how unpleasant it actually is to wait for three days (rather than just one). In this case, it is the delay length itself that is experienced. People may have a vague idea of how painful it will be to wait for a desirable outcome when the delay information is merely described, but actually experiencing the waiting period might result in a more accurate understanding of the pain involved. This in turn would render the underlying preference more consistent than in cases where the delay has not been experienced (for an initial study on how these experiences impact intertemporal choice, see Jimura, Myerson, Hilgard, Braver, & Green, 2009). The same principle applies to experiencing the actual outcome—having encountered an outcome firsthand, an individual may be able to make a more confident decision. On the other hand, experiencing a delay instead of learning it from description may also introduce perceptual uncertainty, because the objective delay length can only be estimated from experience. There are thus many ways in which experience might impact intertemporal choice, and exploring and disentangling the various influences will be an illuminating task for future research.

9.5 Conclusion

In an uncertain world, decisions about the future are inherently characterized by a lack of foreseeability—in terms of whether the expected outcome will ever actually materialize, how attractive that outcome really is, and how dreadful the waiting time will be. Nevertheless, empirical studies and theoretical work on intertemporal choice have only just begun to recognize the potentially critical role that uncertainty plays in shaping people's choices about future outcomes. Adopting such a perspective may lead to a very different interpretation of the hallmarks of intertemporal choice: what may appear irrational in a fully foreseeable and reliable world could actually represent adaptive behavior under uncertainty. This possibility calls for a shift in the science of intertemporal decision making, one that embraces the uncertainty inherent in the intertemporal choices people make. Such a paradigmatic reorientation may not yield all the answers immediately, but those answers may well be worth waiting for.