Risk and Time: Theory, Evidence and Economic Implications

Thomas Epper
University of St. Gallen

Introduction

Thomas Epper's research is mainly concerned with the formation, modeling and measurement of economic preferences, and, in particular, risk, ambiguity, time and social preferences. He uses decision theoretic, experimental and econometric methods to understand how economic agents make decisions and why they fail to choose optimal alternatives in some situations. In his recent empirical research, Thomas Epper combines data from large-scale experiments of a broad population with public register data on wealth, health, education and crime.

This document focusses on one particular strand of Thomas Epper's research. It discusses important interactions of risk and time preferences. These interactions have direct implications for understanding the demand for insurance, investment and savings behavior.

Most important decisions we face in our daily lives entail uncertain and delayed benefits and costs. Examples are how much of our income we spend, invest or save for the future, and whether we engage in actions that generate immediate gratification, but have potentially harmful consequences in the longer run, such as smoking, drinking, eating unhealthy meals, or polluting the environment. It is therefore no surprise that attitudes towards risk and time constitute an essential building block of economic theory. Specifically, the success of economic theory with regard to how well it predicts, explains or guides choices, critically hinges on assumptions about the structure and heterogeneity of preferences, and their interaction across choice domains.

Because uncertainty and delay often arise simultaneously, plausible theories of decision making should be capable of explaining behavior when both risk and delay are involved. The standard approach to model individual decisions treats these two domains as separable, however. Separable means that this approach suggests that there are no interactions between the two choice domains.

To evaluate situations that contain both dimensions, economists tyically combine models from both domains. The usual way of combining neglects the fact that risk and time are indeed two closely intertwined domains. While striking parallels between economic models for risk and time have been recognized (see e.g. Prelec and Loweenstein, 1991), it first required elaborate experimental research (see the literature cited below) to understand the relevance and shape of the domain interactions.

In brief, the empirical research challenges the standard modeling assumptions in both domains, as well as the assumption that the two domains are separable. It shows that risk tolerance and patience interact in important ways, and - as illustrated below - that these close interrelations also have important implications on how we make economic decisions and how we can improve them.

The idea that risk and time are two closely intertwined domains makes intuitively sense: By definition, every risky choice requires that there is at least a small delay between the point in time the alternatives are assessed and the point in time uncertainty is resolved. If, instead, these two points in time would collapse, the consequences of one's actions would be immediately known, and thus there would be no more uncertainty involved. Similarly,

every good or service that does not materialize in the immediate present, but gets promised to materialize at some point in the future, is uncertain by its very nature. In other workds, only what we hold in our hands can ultimately be certain - all promised future benefits or costs have the chance to disappear.[1]

The remainder of this document is structured as follows: The next section outlines the standard economic approach to model choice under risk and choice over time. We then briefly discuss some evidence refuting the traditional models. An illustration on how this empirical evidence can be incorporated into a more realistic economic model follows. The last section discusses some key implications of the empirical evidence and this model.

The standard approach of modeling risk and time

By and large, the literatures of mathematically modeling risk and time preferences evolved separtely from each other. The economic standard theory of choice under risk, expected utility theory (von Neumann and Morgenstern, 1944) presumes that alternatives are evaluated by the sum of the outcome utilities u weighted by the probability p they occur. In economic theory, each alternative under consideration is assigned value. It is then assumed that the agent picks the alternative with the highest value from the choice set. Formally, we can write the value of a risky alternative, L, as $V(L) = \sum_i p_i u_i$, with $\sum_i p_i = 1$ and $p_i > 0$. This theory is viewed as a normative benchmark, i.e. it says how an agent should evaluate alternatives if information about outcomes and probabilities are objectively known. The shape of the utility function u is the model's sole parameter; it governs the agent's risk tolerance. Expected utility theory has been widely used to model economic decisions, such as investment choices, and it is still the main workhorse for modeling insurance choices. The main reason for this, is the model's tractability and simplicity.[2] On the empirical side, however, it's descriptive validity has been challenged by some well-documented violations of it's key assumptions (see e.g. Allais, 1953; Kahneman and Tversky, 1979).

The way economists traditionally modeled choices over time shares similiarities to how they model choices under risk (Lowenstein and Prelec, 1992). The classic theory, discounted utility theory (Samuelson, 1952) for modeling intertemporal choices also weights outcome utility by some factor and then sums over the stream of utilities materializing at different points in time. Formally, an payment stream T is evaluated as $V(T) = \sum_t \delta^t u_t$, with $t \geq 0$ and $1 \geq \delta \geq 0$. Intuitively, this model first converts each future outcome utility into the present, and then sums over these discounted utilities. It hs two parameters: The utility function u and the discount factor δ. The assumptions put into this model appear to be stronger than those underlying expected utility theory. While most people would possibly agree that it makes sense to weight outcome utilities by their objective probabilities, it seems to be less clear why one should take the exponential function δ^t to discount the future. However, as Strotz (1955) has shown, an exponential discount factor is the only discount factor that ensures that future

---

[1]In fact, the early research on choice over time (see e.g. the contributions by Eugen von Böhm-Bawerk or Irving Fisher), already considered risk as an important determinant of discounting behavior.

[2]Some of this tractability, however, is retained for some generalizations of this model (see Machina, 1982).

actions are in line with present plans. Put differently, this particular form ensures that the agent is dynamically consistent, i.e. - assuming no other things change - that she or he would always stick with her/his plans, no matter when she/he would be reevaluates them.

This poses the question about how risk and time should be integrated into a unifying model. Specifically, how would one evaluate a risky alternative with outcomes materializing at future points in time? The literature suggests to combine the two domain models in a separable way. A risky and delayed alternative P would formally be evaluated by $V(P) = \sum_t \delta^t \sum_i p_i u_{it}$. Note that, due to the fact that we integrated both parts in a separable way, it does not matter whether we first evaluate time and then risk, or whether we first evaluate risk and then time.

Empirical evidence against the standard model

There are two streams of research refuting the standard model: A first stream makes use of experimental methods to test economic theory. We describe some key results reported in this literature below. The second stream documents real-world phenomena that are incompatible with theoretical predictions. We illustrate some specific examples in the last part of this document.

Experiments are particularly well suited for testing economic theory. In such experiments, participants face a series of decisions among sets of alternatives. The alternatives are predetermined such that the participant's choice reveals whether she or he obeys a particular theoretical assumption. To give a concrete example: Maurice Allais (1953) proposed a number of examples that permit a direct test of the key separability condition entailed in expected utility theory. It is this separability condition which ensures that utilities are weighted by the respective objective probabilities - or, put differently, that the valuation function is linear in probabilities. One of these examples confronts the subject with the following two binary choices: First, the subject is asked to choose between USD 3,000 for sure or USD 4,000 in 80 of 100 cases (and nothing otherwise). Second, the subject is asked to choose between USD 3,000 in 25 of 100 cases (and nothing otherwise) or USD 4,000 in 20 of 100 cases (and nothing otherwise). The typical response pattern documented in such pairs of choices (see e.g. Kahneman and Tversky, 1979) is as follows: Participants tend to pick the certain option in the first pair of alternatives, but they are willing to give up the 5 percentage point higher chance of winning in the second pair and opt for the alternative which potentially yields the USD 4,000. The literature refers to this finding as "certainty effect". It says that people are generally attracted by certainty; but once it is removed, they put more emphasis on the money they can gain. It can be shown (see Fehr-Duda and Epper, 2012, among others) that this choice violates expected utility theory. Specifically, this violation gives rise to nonlinear probability weighting. That is, people tend to not weigh outcome utilities by their objective probabilities, even when these probabilities are known to them. Bruhin, Fehr-Duda and Epper (2010) document that about 80% of subjects distort probabilities, and the remaining 20% do not. Fehr-Duda and Epper (2012) provide a review of the recent literature, and illustrate how the typical response pattern can be accommodated by a slight variation of the standard model of choice under risk. In particular, one can replace the objective probabilities

by decision weights. These decision weights are obtained by transforming the cumulative distribution (instead of the probability distribution). The transformation has to satisfy specific properties to map Maurice Allais' paradoxical choices. For instance, to accommodate the typical response pattern in the above pair of binary choices, it is required that the function transforming p, $w(p)$, is increasingly elastic (see e.g. Segal, 1987).[3]

Similar patterns are documented in the domain of intertemporal choice (see Frederick, Loewenstein and O'Donogue, 2002). Future utility is usually not discounted at a constant per-period rate, i.e. it is not discounted with an exponential discount function. Instead, the empirical literature suggests that people tend to discount at a hyperbolically declining rate. The literature refers to this behavior as "hyperbolic discounting". Subjects who are prone to such behavior exhibit a "present bias". That is, they are tempted by immediate gratification, and thus overconsume in the short run or engage in actions they might regret in the future.

Thus, we know that the two building blocks of the theory combining risk and time fail. This evidence does not only challenge the theories in both domains, but it also invalidates the standard approach to incorporating them into a model that can deal with delayed, uncertain alternatives. We therefore need to look for more realistic theories of risk and time - theories that can accommodate the violations of expected utility theory, the violations of discounted utility theory, plus potential interactions between the choice domains.

There are some good news for developing such a theory, however: Violations in the two choice domains seems to be closely related in that nonlinear probability weighting is associated with hyperbolic discounting. Epper, Fehr-Duda and Bruhin (2011) provide direct empirical evidence for this. This evidence restricts the set of feasible models considerably. In particular, we are looking for a model that explains violations in both choice domains through a common channel.

Still, some complex interactions seem to occur between the two domains (see Epper and Fehr-Duda (2015) for a review). Experimental research shows that people's risk tolerance increases with time delay (Shelley, 1994; Ahlbrecht and Weber, 1997; Sagristano, Trope, and Liberman, 2002; Noussair and Wu, 2006; Coble and Lusk,2010; Abdellaoui, Diecidue, and Öncüler 2011). Similarly, people typically prefer uncertainty to resolve at once instead of over the course in time (Gneezy and Potters, 1997; Thaler, Tversky, Kahneman, and Schwartz, 1997; Bellemare, Krause, Kroger, and Zhang, 2005; Gneezy, Kapteyn, and Potters, 2003; Haigh and List, 2005). Evidence also suggest that people have a preference for the timing of uncertainty resolution. Specifically, in many situations they prefer to learn about news later rather than earlier (Chew and Ho, 1994; Ahlbrecht and Weber, 1996; Arai, 1997; Lovallo and Kahneman, 2000; Eliaz and Schotter, 2007; von Gaudecker, van Soest, and Wengstroem, 2011). It also appears that the order of whether risk or time is evaluated first matters (Öncüler and Onay, 2009)

A coherent model incorporating the domains of risk and time should be able to accommodate all these pheonomena. The standard approach clearly fails; as do models that replace

---

[3]Another important finding is the evidence for reference dependence (see e.g. Kahneman and Tversky, 1979). A discussion of this issue is beyond the scope of this document. Cumulative prospect theory (Tversky and Kahneman, 1992) incorporates both findings, probability- and reference-dependence.

expected utility theory and discounted utility theory by more realistic domain models, while still retaining *separability*.

Modeling risk and time

As we argued before, outcomes that materialize in the future are uncertain by their very nature. An intuitive way to model future outcomes is therefore to view them as only materializing with some per-period probability $s < 100\%$. As an example, consider a certain payment of USD 1,000 materializing immediately. Once this payment gets promised for some future point in time, there may be a chance that "something goes wrong". In particular, something unrelated to the promised payment may go wrong before it materialize which reduces the likelihood of eventually collect the expected amount. It is straightforward to incorporate this intuition into the economic model. Specifically, we can assume that future outcomes only obtain with a probability s^t. The probability of obtaining it therefore depends on the time delay t. Once we make allegedly guaranteed outcomes risky, all the documented violations of expected utility theory will carry over to the time domain. Indeed, it can be shown that, in such a model and consistent with the empirical evidence (Epper, Fehr-Duda and Bruhin, 2011), Allais-type violations of expected utility theory indeed predict hyperbolic discounting, the most widely discussed violation of discounted expected utility.

To see how alternatives are evaluated according to this model, and how separability is dropped, consider the most simple case of a guaranteed outcome which is deferred to the future. The general case is illustrated in Epper and Fehr-Duda (2017). Assume that the utility of this outcome is $u > 0$. In case something goes wrong, the agent receives nothing instead. If the outcome obtains in the present, we have $V(P) = u$. However, if it materialize with a delay of t, we have $V(P) = w(s^t)u + (1 - w(s^t))0 = w(s^t)u$. This valuation function is easily extended to cases where there is risk in the original alternative, and where there are multiple outcomes involved (see Epper and Fehr-Duda, 2014 and 2017). Since alternatives are risky, we want to incorporate two major findings, both of which are documented in the risk literature. First, in settings where the delay is negligible, people tend to overweight the extreme tails of distributions. Second, people tend to become more risk seeking when mixing the original distributions with a worse one. Both findings refer to special cases of the examples brought up by Maurice Allais (1953). It also turns out that they can be accommodated easily within our model, by constraining the functional form of w. These constraints say how additions and scalings of the arguments of this function affect its value. The specific properties are not further elaborated in this note, but are discussed in detail in Epper and Fehr-Duda (2017).

The model featuring these properties makes predictions that fit well to reported behavior in settings where delay is negligible (see e.g. Bruhin, Fehr-Duda and Epper (2010), Epper, Fehr-Duda and Bruhin (2010), Fehr-Duda, Epper and Bruhin (2011), among others). Such a model also predicts pheonomena in the domain of choice over time, such as hyperbolic discounting (Epper and Fehr-Duda 2015). Lastly, such a model also generates the complicated interactions between the risk and the time domain. As extensively illustrated in Epper and Fehr-Duda (2012, 2015, 2017), it can generate the whole set of pheonomena outlined in the

previous section. As shown in Epper and Fehr-Duda (2017), it also produces a set of novel and untested predictions. For example, it predicts that people's preference for the timing and frequency of uncertainty resolution will be correlated, and thus is not governed by separate behavioral parameters.

Understanding real-world pheonomena

In this last section, we discuss some relevant implications of the model outlined in the previous section. Specifically, we will have a look at how these findings can help to understand disparities in the demand for insurance. The results and implications are discussed formally in Epper and Fehr-Duda (2017).

It appears to be difficult to paint a coherent picture of people's risk preferences using standard theory. In some situations their behaviors seem to be extremely risk averse while in others the opposite is the case. For example, consumers are typically willing to buy small-scale insurance at exorbitant prices (think about insurance for mobile phones or household appliance). This stays in contrast to expected utility theory, which predits that consumers should be close to risk neutral in such situations (Loomes and Segal, 1994). However, people are often reluctant to purchase adequate life insurance, and, thereby expose their loved ones to considerable poverty risk (Bernheim, Forni, Gokhale, and Kotlikoff, 2003; Cutler, Finkelstein and McGarry, 2008). To give some more examples of similar disparities: The proportion of people in the population participating on the stock market is very low in most countries (Giannetti and Koskinen, 2010). This indicates that most people are not willing to take the risks associated with such investments. Contrarily, we observe that inhabitants of disaster-prone areas are often not willing to take out insurance even if it is highly subsidized (Kunreuther, 1984; Viscusi, 2010). That is, they are willing to take substantial, low-probability risks of losing their wealth or lives.

These apparently inconsistent patterns can be understood with the model introduced in Epper and Fehr-Duda (2017). It illustrates that taking into account the timing and process of uncertainty resolution is key for understanding the disparities. Specifically, it argues that these characteristics of the alternatives under consideration determine whether unlikely events are under- or overweighted. Here we provide intuition for this result.

The above situations differ in terms of *how often* information is available, and *when* uncertainty is resolved. If uncertainty resolves in multiple stages, the weights on outcome utility are compounded. The compounding of weights (all smaller than one) yields the possible outcomes to materialize with low probability. The properties of the weighting function motivated above (remember: they are just mappings of the Allais' paradoxes to the preference functional) imply that risk aversion gets amplified once such compounding takes place. As a result, the model is able to predict extreme risk aversion for situations where uncertainty is resolved gradually over time. This is exactly what happens on stock markets or when considering household appliances: The information on asset prices is readily available, for many assets even in real time. Similarly, we obtain regular feedback on the properly functioning of our household appliances. The model therefore predicts that people are reluctant to participate on the stock

market, but are willing to pay high prices for small-scale insurance. This predictions are consistent with the empirical evidence.

On the other hand, natural disasters and life insurance deal with consequences which materialize far in the future. In addition, uncertainty is typically resolved at one shot in such situations. As there is no compounding going on, but there is some uncertainty involved with the passage of time, the model predicts more risk tolerance with increasing delay. If subjects are willing to take higher risks once outcomes materialize far in the future, they will also be less willing to purchase insurance against such events. Specifically, it is predicted that they are reluctant to take actions against global warming (at least as long as they are not directly affected by its negative consequences). Once again, this model prediction is consistent with the findings documented in the literature referred above.

Besides resolving puzzles in insurance and investment choice, the model also has other relevant applications. For example, it can be used to assess policies that affect property rights or laws of contract enforcement. It might therefore help to understand heterogeneity in insurance and investment behavior across the globe.

References

Abdellaoui, M., E. Diecidue, and A. Öncüler (2011): "Risk Preferences at Different Time Periods: An Experimental Investigation," Management Science, 57(5), 975–987.

Ahlbrecht, M., and M. Weber (1997): "An Empirical Study on Intertemporal Decision Making under Risk," Management Science, 43(6), 813–826.

Allais, M. (1953): "Le Comportement de L'Homme Rationnel Devant le Risque: Critique des Postulats et Axiomes de L'Ecole Americaine," Econometrica, 21, 503—-546.

Arai, D. (1997): "Temporal Resolution of Uncertainty in Risky Choices," Acta Psychologica, 96, 15–26.

Bellemare, C., M. Krause, S. Kröger, and C. Zhang (2005): "Myopic Loss Aversion: Information Feedback vs. Investment Flexibility," Economics Letters, 87, 319–324.

Bernheim, B. D., L. Forni, J. Gokhale, and L. J. Kotlikoff (2003): "The Mismatch Between Life Insurance Holdings and Financial Vulnerabilities: Evidence from the Health and Retirement Study," American Economic Review, 93(1), 354–365.

Bruhin, A., H. Fehr-Duda, and T. Epper (2010): "Risk and Rationality: Uncovering Heterogeneity in Probability Distortion," Econometrica, 78(4), 1375–1412.

Chew, S., and J. Ho (1994): "Hope: An Empirical Study of Attitude Toward the Timing of Uncertainty Resolution," Journal of Risk and Uncertainty, 8, 267–288.

Coble, K., and J. Lusk (2010): "At the Nexus of Risk and Time Preferences: An Experimental Investigation," Journal of Risk and Uncertainty, 41(1), 67–79.

Cutler, D. M., A. Finkelstein, and K. McGarry (2008): "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," American Economic Review, 98(2), 157–162.

Eliaz, K., and A. Schotter (2007): "Experimental Testing of Intrinsic Preferences for NonInstrumental Information," American Economic Review, Papers and Proceedings, 97(2), 166–169.

Epper, T., and H. Fehr-Duda (2015): "Comment on 'Risk Preferences Are Not Time Preferences': Balancing on a Budget Line," American Economic Review, 105(7), 2261–2271.

Epper, T., and H. Fehr-Duda (2015): "The Missing Link: Unifying Risk Taking and Time Discounting," Working Paper.

Epper, T., and H. Fehr-Duda (2017): "A Tale of Two Tails: Explaining the Underinsurance Puzzle," Working Paper.

Epper, T., H. Fehr-Duda, and A. Bruhin (2011): "Viewing the Future through a Warped Lens: Why Uncertainty Generates Hyperbolic Discounting," Journal of Risk and Uncertainty, 43(3), 169–203.

Fehr-Duda, H., and T. Epper (2012): "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences," Annual Review of Economics, 4, 567–593.

Frederick, S., G. F. Loewenstein, and T. O'Donoghue (2002): "Time Discounting and Time Preference: A Critical Review," Journal of Economic Literature, 40(2), 351–401.

Giannetti, M., and Y. Koskinen (2010): "Investor Protection, Equity Returns, and Financial Globalization," Journal of Financial and Quantitative Analysis, 45(1), 135–168.

Gneezy, U., A. Kapteyn, and J. Potters (2003): "Evaluation Periods and Asset Prices in a Market Experiment," Journal of Finance, LVIII(2), 821–837.

Gneezy, U., and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," Quarterly Journal of Economics, 112(2), 631–645.

Gul, F. (1991): "A Theory of Disappointment Aversion," Econometrica, 59(3), 667–686.

Haigh, M., and J. List (2005): "Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis," Journal of Finance, 60(1), 523–534.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," Econometrica, 47(2), 263–292.

Loomes, G., and U. Segal (1994): "Observing Different Orders of Risk Aversion," Journal of Risk and Uncertainty, 9, 239–256.

Lovallo, D., and D. Kahneman (2000): "Living with Uncertainty: Attractiveness and Resolution Timing," Journal of Behavioral Decision Making, 13, 179–190.

Machina, M. (1982): " 'Expected Utility' Analysis without the Independence Axiom," Econometrica, 50, 277–323.

Noussair, C., and P. Wu (2006): "Risk Tolerance in the Present and the Future: An Experimental Study," Managerial and Decision Economics, 27, 401–412.

Öncüler, A., and S. Onay (2009): "How Do We Evaluate Future Gambles? Experimental Evidence on Path Dependency in Risky Intertemporal Choice," Journal of Behavioral Decision Making, 22, 280–300.

Prelec, D., and G. F. Loewenstein (1991): "Decision Making over Time and under Uncertainty: A Common Approach," Management Science, 37(7), 770–786.

Quiggin, J. (1982): "A Theory of Anticipated Utility," Journal of Economic Behavior and Organization, 3, 323–343.

Samuelson, P. (1937): "A Note on Measurement of Utility," Review of Economic Stud- ies, 4, 155–161.

Sagristano, M., Y. Trope, and N. Liberman (2002): "Time-dependent Gambling: Odds Now, Money Later," Journal of Experimental Psychology, 131(3), 364–376.

Segal (1987): "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," International Economic Review, 28(1), 175–202.

Shelley, M. (1994): "Gain/Loss Asymmetry in Risky Intertemporal Choice," Organizational Behavior and Human Decision Processes, 59, 124–159.

Strotz, R. H. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," Review of Economic Studies, 23(3), 165-180.

Thaler, R. H., A. Tversky, D. Kahneman, and A. Schwartz (1997): "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test," Quarterly Journal of Economics, 112(2), 647–661.

Tversky, A., and D. Kahneman (1992): "Advances in Prospect Theory: Cumulative Representa- tion of Uncertainty," Journal of Risk and Uncertainty, 5, 297–323.

von Gaudecker, H.-M., A. van Soest, and E. Wengström (2011): "Heterogeneity in Risky Choice Behaviour in a Broad Population," American Economic Review, 101(2), 664–694.

von Neumann, J. and O. Morgenstern (1947): "Theory of Games and Economic Behavior," Princeton University Press, Priceton (NJ).