



Risk, uncertainty, and option exercise

Jianjun Miao^{a,b,1}, Neng Wang^{c,d,*}

^a Department of Economics, Boston University, CEMA, Central University of Finance and Economics, 270 Bay State Road, Boston, MA 02215, United States

^b Xinhua School of Finance and Insurance, Zhongnan University of Economics and Law, China

^c Columbia Business School, NBER, 3022 Broadway, Uris Hall 812, New York, NY 10027, United States

^d Shanghai University of Finance and Economics, China

ARTICLE INFO

Article history:

Received 19 October 2009

Accepted 8 September 2010

Available online 19 November 2010

JEL classification:

D81

G31

Keywords:

Ambiguity

Multiple-priors utility

Real options

Optimal stopping problem

ABSTRACT

Many economic decisions can be described as an option exercise or optimal stopping problem under uncertainty. Motivated by experimental evidence such as the Ellsberg Paradox, we follow Knight (1921) and distinguish risk from uncertainty. To capture this distinction, we adopt the multiple-priors utility model. We show that the impact of ambiguity on the option exercise decision depends on the relative degrees of ambiguity about continuation payoffs and termination payoffs. Consequently, ambiguity may accelerate or delay option exercise. We apply our results to investment and exit problems, and show that the myopic NPV rule can be optimal for an agent having an extremely high degree of ambiguity aversion.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Many economic decisions can be described as dynamic discrete choices under uncertainty. Examples are abundant. An investor may decide whether and when to invest in a project. A firm may decide whether and when to enter or exit an industry. It may also decide whether and when to default on debt. A worker may decide whether and when to accept a job offer or quit his job. All these decisions share three characteristics. First, the decision is irreversible to some extent. Second, there is uncertainty about future rewards. Third, agents have some flexibility in choosing the timing of the decision. These three characteristics imply that waiting has positive value. Importantly, all the preceding problems can be viewed as a problem in which an agent decides when to exercise an “option” analogous to an American call option—it has the right but not the obligation to buy an asset at some future time of its choosing. This real-options approach has been widely applied in economics and finance (see Dixit and Pindyck, 1994). The aim of this paper is to analyze an option exercise problem where there is a distinction between risk and uncertainty in the sense often attributed to Knight (1921), and where agents’ attitudes toward uncertainty play a nontrivial role.

The standard real-options approach to investment under uncertainty can be summarized as “a theory of optimal inertia.” Dixit (1992) argues that “firms that refuse to invest even when the currently available rates of return are far in excess of the cost of capital may be optimally waiting to be surer that this state of affairs is not transitory. Likewise, farmers who carry large losses may be rationally keeping their operation alive on the chance that the future may be brighter.”

* Corresponding author at: Columbia Business School, NBER, 3022 Broadway, Uris Hall 812, New York, NY 10027, United States. Tel.: +1 212 854 3869. E-mail addresses: miaoj@bu.edu (J. Miao), nw2128@columbia.edu (N. Wang).

¹ Tel.: +1 617 353 6675.

However, the standard real-options approach rules out the situation where agents are unsure about the likelihoods of states of the world. It typically adopts strong assumptions about agents' beliefs. For example, according to the rational expectations hypothesis, agents know the objective probability law of the state process and their beliefs are identical to this probability law. Alternatively, according to the Bayesian approach, an agent's beliefs are represented by a subjective probability measure or Bayesian prior. There is no meaningful distinction between *risk*, where probabilities are available to guide choice, and *uncertainty*, where information is too imprecise to be summarized adequately by probabilities. By contrast, Knight (1921) emphasizes this distinction and argues that uncertainty is more common in decision-making settings. Henceforth, we refer to such uncertainty as *Knighitian uncertainty* or *ambiguity*. For experimental evidence, the Ellsberg Paradox suggests that people prefer to act on known rather than unknown or ambiguous probabilities.² Ellsberg-type behavior contradicts the Bayesian paradigm, i.e., the existence of a single probability measure underlying choices.

To incorporate Knighitian uncertainty or ambiguity, we adopt the recursive multiple-priors utility model developed by Epstein and Wang (1994). In that model, the agent's beliefs are represented by a collection of sets of one-step-ahead conditional probabilities. These sets of one-step-ahead conditionals capture both the degree of ambiguity and ambiguity aversion.³ The axiomatic foundation for the recursive multiple-priors utility model is laid out by Epstein and Schneider (2003). Their axiomatization is based on the static multiple-priors utility model proposed by Gilboa and Schmeidler (1989).

We describe an ambiguity averse agent's option exercise decision as an optimal stopping problem using the Epstein and Wang (1994) utility model. We then characterize the optimal stopping rules. The standard real-options approach emphasizes the importance of risk in determining option value and timing of option exercise. An important implication is that an increase in risk in the sense of mean preserving spread raises option value and delays option exercise. Recognizing the difference between risk and ambiguity, we conduct comparative statics analysis with respect to the set of one-step-ahead conditionals.

In our model, the agent is ambiguous about a state process which influences the continuation and termination payoffs. Importantly, we distinguish between two cases according to whether or not the agent is still ambiguous about the termination payoff after he exercises the option. This distinction is critical since it may generate opposite comparative statics results. We show that for both cases, ambiguity lowers the option value. Moreover, if there is no uncertainty after option exercise, a more ambiguity averse agent will exercise the option earlier. However, if he is also ambiguous about termination payoffs after option exercise, he may exercise the option later.⁴ This is because ambiguity lowers the termination payoff and this effect may dominate the decrease in the option value.

We provide two applications – real investment and firm exit – to illustrate our results. The real investment decision is an example where an agent decides if and when to exercise an option to pursue upside potential gains. Entry and job search are similar problems. Under a specification of the set of priors, we show explicitly that if the investment project generates a stream of future uncertain profits and if the agent is ambiguous about these profits, then a more ambiguity averse agent invests relatively later. By contrast, if the investment payoff is delivered in lump sum and uncertainty is fully resolved upon the option exercise, then a more ambiguity averse agent makes the investment sooner.

The exit problem represents an example where an agent decides if and when to exercise an option to avoid downside potential losses. Other examples include default and liquidation decisions. We show that the exit timing depends crucially on whether the owner of a firm is ambiguous about the outside value. This ambiguity may dominate the effect of ambiguity about the profit opportunities if stay in business. Consequently, an ambiguity averse owner may be hesitant to exit, even though it has lower option value. For both investment and exit problems, we solve some examples explicitly under some specification of the set of priors. We show that the myopic net present value (NPV) rule can be optimal for an agent having an extremely high degree of ambiguity aversion.

Our paper contributes to the literature on applications of decision theory to macroeconomics and finance surveyed recently by Backus et al. (2004). The idea of ambiguity aversion and the multiple-priors utility model have been applied to asset pricing and portfolio choice problems in a number of papers.⁵ A different approach based on robust control theory is proposed by Hansen and Sargent and their coauthors.⁶ They emphasize “model uncertainty” which is also motivated in part by the Ellsberg Paradox. We refer readers to Epstein and Schneider (2003) for further discussion on these two approaches.

Our paper is related to Nishimura and Ozaki (2004, 2007). Nishimura and Ozaki (2004) apply the Choquet expected utility model proposed by Schmeidler (1989) to study a job search problem.⁷ They assume that workers are ambiguous about the wage offer. They show that ambiguity reduces the reservation wage and speeds up job acceptance. Nishimura and Ozaki

² One way to describe this paradox (Ellsberg, 1961) is as follows. There are two urns. The first urn contains exactly 50 red and 50 black balls. The second urn also has 100 balls (either black or red), but the exact numbers of red and black balls are not known. Subjects are offered a bet on drawing a red ball from the two urns. A majority of subjects choose from the first urn rather than the second. The paradox surfaces after a second bet is offered – a bet on a black ball – and a majority of subjects still prefers to bet on a ball from the first urn rather than from the second.

³ For a formal definition of ambiguity aversion, see Epstein (1999), Epstein and Zhang (2001), and Ghirardato and Marinacci (2002).

⁴ Miao and Wang (2007) analyze the effects of risk and risk aversion on investment timing in a continuous time model of irreversible investment under incomplete markets. They show that whether markets are complete or not after the option is exercised matters for the effects of risk aversion on the timing of option exercising.

⁵ See Chen and Epstein (2002), Epstein and Miao (2003), Epstein and Schneider (2006, 2007), Epstein and Wang (1994, 1995), Kogan and Wang (2002), Miao (2004, 2009), and Routledge and Zin (2009).

⁶ See, for example, Anderson et al. (2003) and Hansen and Sargent (2001). See Maccheroni et al. (2006) for a related axiomatization.

⁷ The Choquet expected utility model is another well known model that also addresses ambiguity. It has been applied to study wage contracts by Mukerji and Tallon (2004). Also see the references therein for other applications.

(2007) apply the continuous-time multiple-priors utility model developed by Chen and Epstein (2002) to study an irreversible investment problem. They assume that investors are ambiguous about the investment payoff. They show that ambiguity delays investment. Our paper reconciles these conflicting results in a general unified framework of the optimal stopping problem. In particular, both the job search and irreversible investment problems belong to our general “investment” problem analyzed in Section 3.1. The generality of our framework also allows us to address a different type of option exercise problem such as exit analyzed in Section 3.2. The job search problem in Nishimura and Ozaki (2004) corresponds to our investment problem with a one-time lump-sum payoff and the irreversible investment problem in Nishimura and Ozaki (2007) corresponds to our investment problem with a flow payoff. Our Proposition 6 reconciles the conflicting results in the preceding two papers.

Our paper is also related to Riedel (2009). Riedel (2009) considers a set of time-consistent priors over the full state space as in Epstein and Schneider (2003) and then extends dynamic programming theory to multiple priors. He studies an optimal stopping problem similar to our investment problem with a one-time lump-sum payoff. Unlike his approach, we apply the recursive approach of Epstein and Wang (1994) by specifying sets of one-step-ahead probabilities in an infinite-horizon Markovian setting. In addition, our framework includes some optimal stopping problems outside the scope of Riedel (2009).

The remainder of this paper proceeds as follows. Section 2 presents the model and results. Section 3 applies the model to investment and exit problems. Section 4 uses a job matching example based on Jovanovic (1979) to discuss the role of learning under ambiguity and an alternative smooth ambiguity model of Klibanoff et al. (2005, 2009). Section 5 concludes. Proofs are relegated to the Appendix.

2. The model

In this section, we first introduce the multiple-priors utility model in Section 2.1. We then present a baseline setup of the optimal stopping problem in Section 2.2. After that, we present characterization and comparative statics results in Section 2.3. We finally consider extensions in Section 2.4.

2.1. Multiple-priors utility

Before presenting the model, we first provide some background about multiple-priors utility. The static multiple-priors utility model of Gilboa and Schmeidler (1989) can be described informally as follows. Suppose uncertainty is represented by a measurable space (S, \mathcal{F}) . The decision-maker ranks uncertain prospects or acts, maps from S into an outcome set \mathcal{X} . Then, the multiple-priors utility $U(f)$ of any act f has the functional form:

$$U(f) = \min_{q \in \Delta} \int u(f) dq,$$

where $u : \mathcal{X} \rightarrow \mathbb{R}$ is a von Neumann–Morgenstern utility index and Δ is a subjective set of probability measures on (S, \mathcal{F}) . Intuitively, the multiplicity of priors models ambiguity about likelihoods of events and the minimum delivers aversion to such ambiguity. The standard expected utility model is obtained when the set of priors Δ is a singleton.

The Gilboa and Schmeidler model is generalized to a dynamic setting in discrete time by Epstein and Wang (1994). Their model can be described briefly as follows. The time t conditional utility from a consumption process $c = (c_t)_{t \geq 1}$ is defined by the Bellman equation

$$V_t(c) = u(c_t) + \beta \min_{q \in \mathcal{P}_t} E_t^q [V_{t+1}(c)], \quad (1)$$

where $\beta \in (0, 1)$ is the discount factor, E_t^q is the conditional expectation operator with respect to measure q , and \mathcal{P}_t is a set of one-step-ahead conditional probabilities, given information available at date t . An important feature of this utility is that it satisfies dynamic consistency because it is defined recursively in (1). Recently, Epstein and Schneider (2003) provide an axiomatic foundation for this model. They also develop a reformulation of utility closer to Gilboa and Schmeidler (1989) where there is a set of priors \mathcal{R} over the full state space implied by all histories of events,

$$V_t(c) = \min_{q \in \mathcal{R}} E_t^q \left[\sum_{s=t}^{\infty} \beta^{s-t} u(c_s) \right]. \quad (2)$$

The key to establishing this reformulation is to note that all sets of one-step-ahead conditionals, as one varies over times and histories, determines a unique set of priors \mathcal{R} over the full state space satisfying the regularity conditions defined in Epstein and Schneider (2003). Epstein and Schneider (2003) establish that if one defines multiple-priors utility according to (2), an added restriction on \mathcal{R} is needed to ensure dynamic consistency. To avoid this complication, we adopt the Epstein and Wang model in (1) and specify the sets of one-step-ahead conditionals as a primitive, instead of the set of priors over the full state space.

2.2. A baseline setup

Consider an infinite-horizon discrete-time optimal stopping problem. As explained in Dixit and Pindyck (1994), the optimal stopping problem can be applied to study an agent’s option exercise decision. The agent’s choice is binary. In each period, he decides whether to stop a process and take a termination payoff, or continue for one more period, and make the same decision in the future.

Formally, uncertainty is generated by a Markovian state process $(x_t)_{t \geq 1}$ taking values in $X = [a, b] \subset \mathbb{R}$. The probability kernel of $(x_t)_{t \geq 1}$ is given by $P : X \rightarrow \mathcal{M}(X)$, where $\mathcal{M}(X)$ is the space of probability measures on X endowed with the weak convergence topology. Continuation at date t generates a payoff $\pi(x_t)$, while stopping at date t yields a payoff $\Omega(x_t)$, where π and Ω are functions that map X into \mathbb{R} . At any time t , after observing x_t , the agent decides whether to stop or to continue. This decision is irreversible in that if the agent chooses to stop, he will not make further choices. In order to focus on the beliefs instead of tastes, we suppose that the agent is risk neutral and discounts future payoff flows according to $\beta \in (0, 1)$.

It is important to point out that, in the preceding setup, if the agent decides to stop, uncertainty is fully resolved. He faces ambiguity during periods of continuation only. In Section 2.4, we will consider a more general case where there is uncertainty about the termination payoff and the agent is ambiguous about this payoff. We will show that ambiguity may have different impact on the agent’s option exercise decision, depending on the relative degrees of ambiguity about different sources of uncertainty.

In standard models, the agent’s preferences are represented by time-additive expected utility. As in the rational expectations paradigm, P can be interpreted as the objective probability law governing the state process $(x_t)_{t \geq 1}$, and is known to the agent. The expectation in the utility function is taken with respect to this law. Alternatively, according to the Savage utility representation theorem, P is a subjective (one-step-ahead) prior and represents the agent’s beliefs. By either approach, the standard stopping problem can be described by the following Bellman equation:

$$F(x) = \max \left\{ \Omega(x), \pi(x) + \beta \int F(x')P(dx'; x) \right\}, \tag{3}$$

where the value function F can be interpreted as an option value.

To fix ideas, we make the following assumptions. These assumptions are standard in dynamic programming theory (see Stokey et al., 1989).

Assumption 1. $\pi : X \rightarrow \mathbb{R}$ is bounded, continuous, and increasing.

Assumption 2. $\Omega : X \rightarrow \mathbb{R}$ is bounded, continuous, and increasing.

Assumption 3. P is increasing and satisfies the Feller property. That is, $\int f(x')P(dx'; x)$ is increasing in x for any increasing function f and is continuous in x for any bounded and continuous function f .

The following proposition describes the solution to problem (3).

Proposition 1. *Let Assumptions 1–3 hold. Then there exists a unique bounded, continuous and increasing function F solving the dynamic programming problem (3). Moreover, if there is a unique threshold value $x^* \in X$ such that*

$$\pi(x) + \beta \int F(x')P(dx'; x) > (<) \Omega(x) \quad \text{for } x < x^* \tag{4}$$

and

$$\pi(x) + \beta \int F(x')P(dx'; x) < (>) \Omega(x) \quad \text{for } x > x^*, \tag{5}$$

then the agent continues (stops) when $x < x^*$ and stops (continues) when $x > x^*$. Finally, x^* is the solution to

$$\pi(x^*) + \beta \int F(x')P(dx'; x^*) = \Omega(x^*). \tag{6}$$

This proposition is illustrated in Fig. 1. The threshold value x^* partitions the set X into two regions—continuation and stopping regions.⁸ The top diagram of Fig. 1 illustrates the situation where

$$\pi(x) + \beta \int F(x')P(dx'; x) > \Omega(x) \quad \text{for } x < x^* \tag{7}$$

and

$$\pi(x) + \beta \int F(x')P(dx'; x) < \Omega(x) \quad \text{for } x > x^*. \tag{8}$$

In this case, we say that the continuation payoff curve crosses the termination payoff curve from above. Under this condition, the agent exercises the option when the process $(x_t)_{t \geq 1}$ first reaches the point x^* from below. The continuation region is given

⁸ For ease of presentation, we do not give primitive assumptions about the structure of these regions. See page 129 in Dixit and Pindyck (1994) for such an assumption. For the applications below, our assumptions can be easily verified.

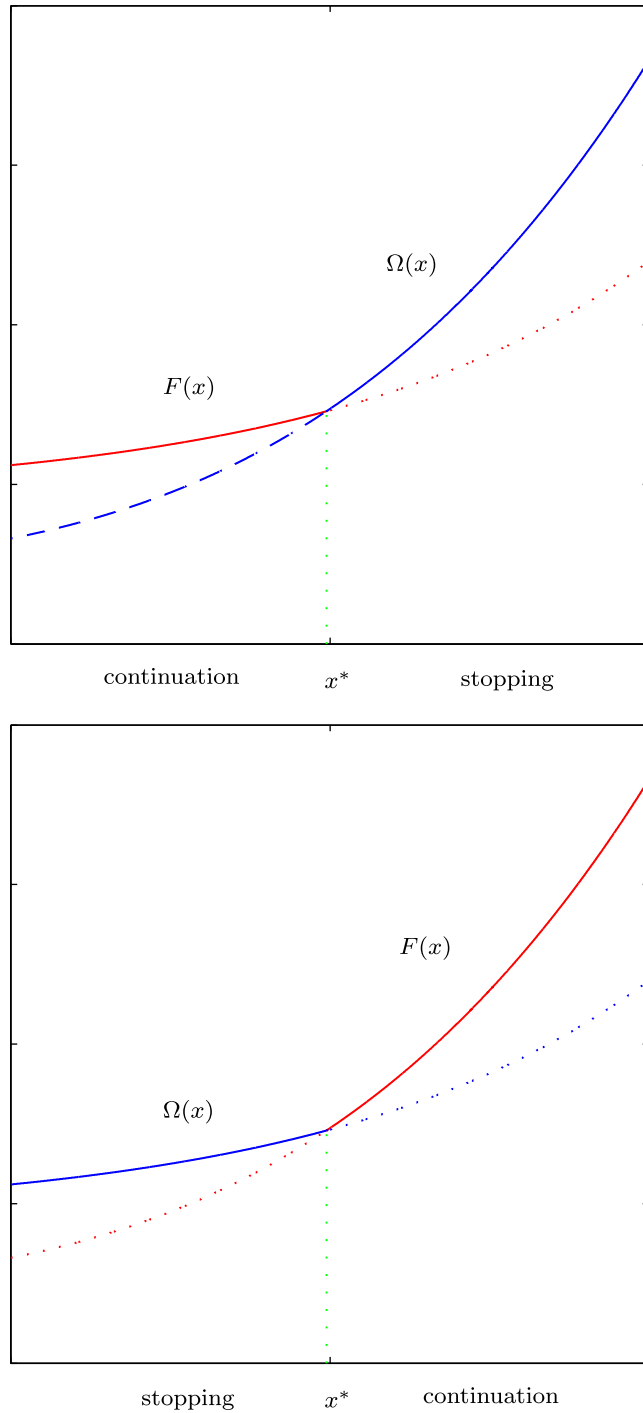


Fig. 1. Value functions and exercising thresholds in the standard model. The top diagram illustrates an option exercise problem such as investment. The bottom diagram illustrates an option exercise problem such as exit.

by $\{x \in X : x < x^*\}$ and the stopping region is given by $\{x \in X : x > x^*\}$. This case describes economic problems such as investment, where the agent’s payoff increases with the state variable. The bottom diagram of Fig. 1 depicts economic problems such as disinvestment or firm exit, where the agent’s payoff decreases in the state variable. The interpretation is similar.

In the above model, a role for Knightian uncertainty is excluded *a priori*, either because the agent has precise information about the probability law as in the rational expectations approach or because the Savage axioms imply that the agent is indifferent to it. To incorporate Knightian uncertainty and uncertainty aversion, we follow the multiple-priors utility

approach (Gilboa and Schmeidler, 1989; Epstein and Wang, 1994) and assume that beliefs are too vague to be represented by a single probability measure and are represented instead by a set of probability measures. More formally, we model beliefs by a probability kernel correspondence $\mathcal{P} : X \rightrightarrows \mathcal{M}(X)$. Given any $x \in X$, we think of $\mathcal{P}(x)$ as the set of conditional probability measures representing beliefs about the next period's state. The multi-valued nature of \mathcal{P} reflects uncertainty aversion of preferences.

A larger set \mathcal{P} may mean the environment is more ambiguous. It may also mean the agent is more ambiguity averse (see Epstein, 1999). Thus, for the multiple-priors model, ambiguity and ambiguity aversion are confounded. In this paper, we will follow the latter interpretation when conducting comparative static analysis. We emphasize also that ambiguity and ambiguity aversion can be about any moments of the distribution including mean, variance and higher-order moments.

The stopping problem under Knightian uncertainty can be described by the following Bellman equation:

$$V(x) = \max \left\{ \Omega(x), \pi(x) + \beta \int V(x') \mathcal{P}(dx'; x) \right\}, \tag{9}$$

where we adopt the notation throughout

$$\int f(x') \mathcal{P}(dx'; x) \equiv \min_{Q(\cdot; x) \in \mathcal{P}(x)} \int f(x') Q(dx'; x), \tag{10}$$

for any Borel function $f : X \rightarrow \mathbb{R}$. Note that if $\mathcal{P} = \{P\}$, then the model reduces to the standard model (3).

To analyze problem (9), the following assumption is adopted.

Assumption 4. The probability kernel correspondence $\mathcal{P} : X \rightrightarrows \mathcal{M}(X)$ is nonempty valued, continuous, compact-valued, and convex-valued, and $P(x) \in \mathcal{P}(x)$ for any $x \in X$. Moreover, $\int f(x') \mathcal{P}(dx'; x)$ is increasing in x for any increasing function $f : X \rightarrow \mathbb{R}$.

This assumption is a generalization of Assumption 3 to correspondence. It ensures that $\int f(x') \mathcal{P}(dx'; x)$ is bounded, continuous, and increasing in x for any bounded, continuous, and increasing function $f : X \rightarrow \mathbb{R}$. A sufficient condition for the monotonicity of $\int f(x') \mathcal{P}(dx'; \cdot)$ is that given any $Q(\cdot; x) \in \mathcal{P}(x)$, $\int f(x') Q(dx'; x)$ is increasing in x for any increasing function $f : X \rightarrow \mathbb{R}$. Notice that Assumption 4 is quite general in the sense that it captures the fact that ambiguity may vary with the state. However, in some examples in Section 3, we consider an IID case in order to derive closed form solutions.

2.3. Erosion of option value

We now analyze the implications of ambiguity and ambiguity aversion on the option exercise decision for the preceding baseline model. We first characterize the solution to problem (9) in the following proposition:

Proposition 2. Let Assumptions 1–4 hold. Then there is a unique bounded, continuous, and increasing function V solving the dynamic programming problem (9). Moreover, if there exists a unique threshold value $x^{**} \in X$ such that

$$\pi(x) + \beta \int V(x') \mathcal{P}(dx'; x) > (<) \Omega(x) \quad \text{for } x > x^{**} \tag{11}$$

and

$$\pi(x) + \beta \int V(x') \mathcal{P}(dx'; x) < (>) \Omega(x) \quad \text{for } x < x^{**}, \tag{12}$$

then the agent stops (continues) when $x < x^{**}$ and continues (stops) when $x > x^{**}$. Finally, x^{**} is the solution to

$$\pi(x^{**}) + \beta \int V(x') \mathcal{P}(dx'; x^{**}) = \Omega(x^{**}). \tag{13}$$

This proposition implies that the agent's option exercise decision under Knightian uncertainty has similar features to that in the standard model described in Proposition 1. It is interesting to compare the option value and option exercise time in these two models.

Proposition 3. Let assumptions in Propositions 1 and 2 hold. Then $V \leq F$. Moreover, for both V and F , if the continuation payoff curves cross the termination payoff curves from above then $x^{**} \leq x^*$. On the other hand, if the continuation payoff curves cross the termination payoff curves from below, then $x^{**} \geq x^*$.

In the standard model, an expected utility maximizer views the world as purely risky. For the decision problems such as investment, waiting has value because the agent can avoid the downside risk, while realizing the upside potential. Similarly, for the decision problems such as exit, waiting has value because the agent hopes there is some chance that the future may be brighter. Now, if the agent has imprecise knowledge about the likelihoods of the state of the world and hence perceives the future as ambiguous, then waiting will have less value for an ambiguity averse agent because he acts on the worst scenario.

The threshold value under Knightian uncertainty can be either bigger or smaller than that in the standard model, depending on the shapes of the continuation and termination payoff curves (see Fig. 2). More specifically, if the continuation payoff curve crosses the termination payoff curve from above, then the threshold value under Knightian uncertainty is smaller than that in the standard model. The opposite conclusion can be obtained if the continuation payoff curve crosses the




Fig. 2. Comparison of the standard model and the model under Knightian uncertainty.

termination payoff curve from below. For both cases, an uncertainty averse agent exercises the option earlier than an agent with expected utility because the former has less option value.

The following proposition concerns comparative statics.

Proposition 4. *Let the assumptions in Proposition 2 hold. Consider two probability kernel correspondences \mathcal{P}_1 and \mathcal{P}_2 . Let the corresponding value functions be V^{P_1} and V^{P_2} and the corresponding threshold values be x^{P_1} and x^{P_2} . If $\mathcal{P}_1(x) \subset \mathcal{P}_2(x)$, then $V^{P_1} \geq V^{P_2}$. Moreover, if the continuation payoff curves cross the termination payoff curves from above (below), then $x^{P_1} \geq (\leq) x^{P_2}$.*

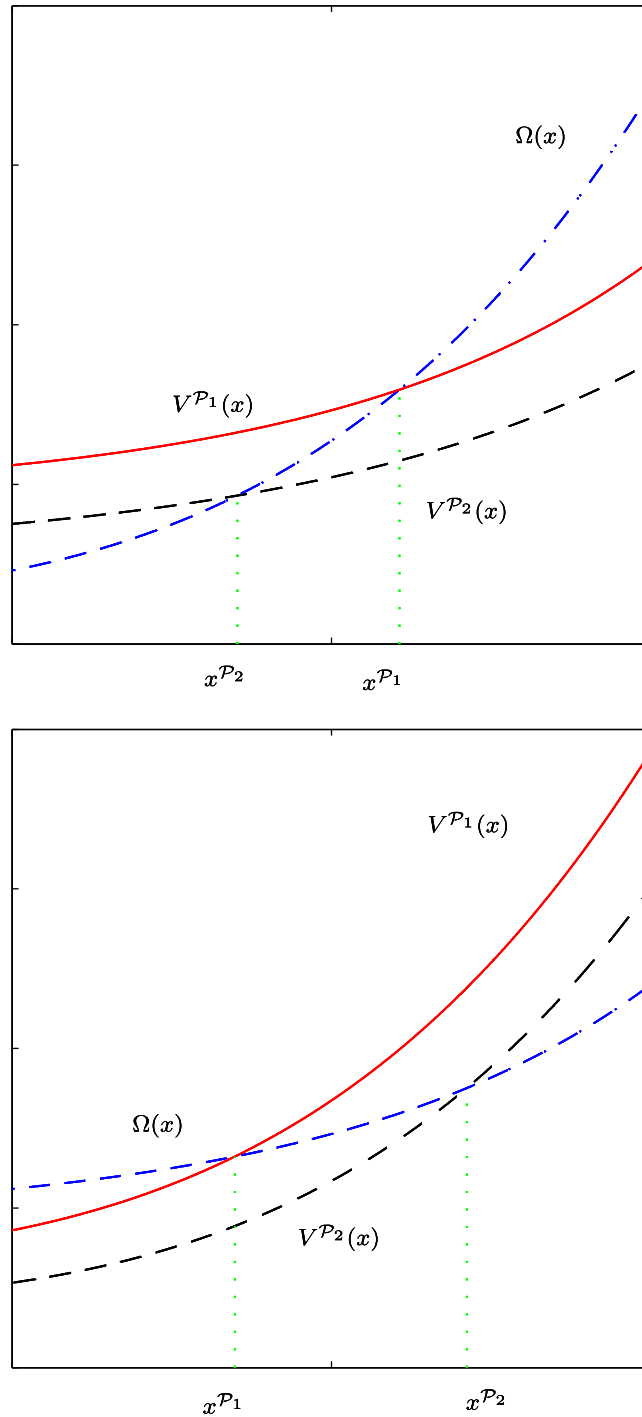


Fig. 3. Option value and exercising thresholds under Knightian uncertainty for two different sets of priors $\mathcal{P}_1 \subset \mathcal{P}_2$.

Recall that the set of priors captures both ambiguity and ambiguity aversion. Holding everything else constant, a larger set of priors delivers a lower utility to the agent because he is more ambiguity averse. Intuitively, the option value is lower if the agent is more ambiguity averse. As a result, a more ambiguity averse agent is less willing to hold the option and hence exercises the option earlier (see Fig. 3). Our interpretation of this proposition is based on the definition of absolute and comparative ambiguity aversion proposed by Ghirardato and Marinacci (2002).⁹ Their theory may also provide a behavioral

⁹ See Epstein (1999) and Epstein and Zhang (2001) for a different definition.

foundation of the interpretation that the set of priors describes the degree of ambiguity. A similar interpretation is also given in Nishimura and Ozaki (2004).

2.4. Ambiguity about termination payoffs

So far, we have assumed that once the agent exercises the option, uncertainty is fully resolved, and that the agent bears ambiguity from waiting only. As will be illustrated in the applications in Section 3, in reality there are many instances in which there is uncertainty about termination payoffs. We now show that if the agent is ambiguous about termination payoffs, ambiguity may have a different impact on the agent's option exercise decision.

To illustrate, we first consider a simple case where the termination payoff $\Omega(x_t)$ does not depend on the state x_t . Suppose it is a random variable with distribution Q . In addition to ambiguity about the state process $(x_t)_{t \geq 1}$, the agent is also ambiguous about the termination payoff Ω . He has a set of priors \mathcal{Q} over Ω . In this case, the agent is ambiguous about two different sources of uncertainty. Section 3.3 will show that an instance of this case is the exit problem. We now formally replace Assumption 2 with:

Assumption 5. Ω is a random variable with distribution Q and \mathcal{Q} is weakly compact and contains Q .

The agent's decision problem can be described by the following Bellman equation:

$$V(x) = \max \left\{ \min_{q \in \mathcal{Q}} \int \Omega dq, \pi(x) + \beta \int V(x') \mathcal{P}(dx'; x) \right\}. \quad (14)$$

When $\mathcal{Q} = \{Q\}$ and $\mathcal{P} = \{P\}$, the preceding problem reduces to the standard one for an expected utility maximizer.

One can prove a characterization proposition for (14) similar to Proposition 2. In particular, there is a threshold value such that the agent exercises the option the first time the process $(x_t)_{t \geq 1}$ falls below this value. However, there is no clear-cut result about comparative statics and the comparison with the standard model as in Propositions 3–4. This is because ambiguity lowers both the continuation payoff and the termination payoff. The overall impact on the option exercise decision depends on which effect dominates. If we fix ambiguity about the continuation payoff and consider the impact of ambiguity about the termination payoff only, we have the following clean comparative statics result analogous to Proposition 4.

Proposition 5. Let Assumptions 1 and 3–5 hold and fix \mathcal{P} . Consider two sets of priors \mathcal{Q}_1 and \mathcal{Q}_2 . Let the corresponding value functions be $V^{\mathcal{Q}_1}$ and $V^{\mathcal{Q}_2}$ and the corresponding threshold values be $x^{\mathcal{Q}_1}$ and $x^{\mathcal{Q}_2}$. If $\mathcal{Q}_1 \subset \mathcal{Q}_2$, then $V^{\mathcal{Q}_1} \geq V^{\mathcal{Q}_2}$ and $x^{\mathcal{Q}_1} \geq x^{\mathcal{Q}_2}$.

This proposition shows that although ambiguity about termination payoffs lowers the option value from continuation, the agent exercises the option later if he is more ambiguous about the termination payoffs. The intuition is that ambiguity lowers the termination payoff and this effect dominates.

In the previous case, there is no *future* uncertainty about termination payoffs. In reality, there are many instances in which there is ongoing uncertainty about termination payoffs after the agent exercises the option. For example, an agent decides whether and when to invest in a project which can generate a stream of future uncertain profits. Then the termination payoff $\Omega(x)$ depends on the future uncertainty about the profits generated by the project. To incorporate this case, we assume that $\Omega(x)$ satisfies the following Bellman equation:

$$\Omega(x) = \Phi(x) + \beta \int \Omega(x') \mathcal{P}(dx'; x), \quad (15)$$

where the period payoff $\Phi : X \rightarrow \mathbb{R}$ is an increasing and continuous function. Note that by the Blackwell Theorem, there is a unique bounded and continuous function Ω satisfying (15). In the standard model with expected utility, we have

$$\Omega(x) = \Phi(x) + \beta \int \Omega(x') P(dx'; x). \quad (16)$$

That is, $\Omega(x)$ is equal to the expected discounted payoffs,

$$\Omega(x) = E \left[\sum_{t=0}^{\infty} \beta^t \Phi(x_t) \mid x_0 = x \right]. \quad (17)$$

When $\Omega(x)$ is given by (15), the agent's decision problem is still described by the dynamic programming equation (4). Again, we can show that ambiguity lowers the option value $V(x)$. However, since ambiguity about the state process $(x_t)_{t \geq 0}$ lowers both the option value and the termination payoff, there is no general comparative statics result about the option exercise timing. We will illustrate this point in the next section.

3. Applications

This section applies our results to two classes of problems: real investment and firm exit. The real investment decision is an example where an agent decides whether or not to exercise an option to pursue upside potential. The exit problem represents an example where an agent decides whether or not to exercise an option to avoid downside loss.

3.1. Investment

A classic application of the option exercise problem is the irreversible investment decision.¹⁰ The standard real-options approach makes the analogy of investment to the exercising of an American call option on the underlying project.¹¹ Formally, consider an investment opportunity which generates stochastic values given by a Markov process $(x_t)_{t \geq 1}$. Investment costs $I > 0$. Then we can cast the investment problem into our framework by setting

$$\Omega(x) = x - I, \tag{18}$$

$$\pi(x) = 0. \tag{19}$$

We can also write the agent's investment decision problem as follows:

$$V(x) = \max \left\{ x - I, \beta \int V(x') \mathcal{P}(dx'; x) \right\}. \tag{20}$$

Note that, according to this setup, once the investor makes the investment, he obtains net rewards $x - I$ and uncertainty is fully resolved. In reality, it is often the case that investment rewards come from the uncertain future. For example, after the investor invests in a project or develops a new product, the project can generate a stream of future uncertain profits. The investor may be ambiguous about future profit flows. To incorporate this case, we presume that the period t profit is given by x_t . Then the discounted total project value at date t is given by

$$\Omega(x_t) = x_t + \beta \int \Omega(x') \mathcal{P}(dx'; x_t), \tag{21}$$

and the investor's decision problem is formulated as

$$V(x) = \max \left\{ \Omega(x) - I, \beta \int V(x') \mathcal{P}(dx'; x) \right\}. \tag{22}$$

The standard real-options model predicts that there is an option value of waiting, because investment is irreversible and flexibility in timing has value. Another main prediction of the standard real investment model is that an increase in risk in the sense of mean-preserving spread raises the option value and delays investment (see Dixit and Pindyck, 1994). This derives from the fundamental insight behind the option pricing theory, in that firms may capture the upside gains and minimize the downside loss by waiting for the risk of project value to be partially resolved.

While the standard real-options model predicts a monotonic relationship between investment and risk, our model makes an important distinction between risk and uncertainty. We argue that risk (which can be described by a single probability measure) and uncertainty (multiplicity of priors) have different effects on investment timing. Specifically, our model predicts that Knightian uncertainty lowers the option value of waiting (see Fig. 4). Moreover, under formulation (20), Propositions 3–4 imply that an increase in ambiguity pulls the investment trigger earlier and hence speeds up investment. These propositions also imply that the more ambiguity averse the investor is, the earlier he makes the investment.

As pointed in Section 2.4, this conclusion is not generally true if there is ongoing uncertainty about the termination payoff. For the investment problem under the formulation in (21)–(22), the investor may be ambiguous about the future profit opportunities of the investment project. He may well hesitate to invest.

We now consider a concrete parametric example to illustrate the above analysis. Recall $X = [a, b]$. Following Epstein and Wang (1994), we consider an IID ε -contamination specification of the set of one-step-ahead priors. That is, let

$$\mathcal{P}(x) = \{(1 - \varepsilon)\mu + \varepsilon m : m \in \mathcal{M}(X)\} \text{ for all } x, \tag{23}$$

where $\varepsilon \in [0, 1]$ and μ is any distribution over X . The interpretation is the following: μ represents the “true” distribution of the reward. The investor does not know this distribution precisely. With probability ε , he believes that the reward may be distributed according to some other distribution. Here ε may represent the degree of ambiguity and ambiguity aversion. This can be justified by observing that, if ε is larger, the set $\mathcal{P}(x)$ is larger in the sense of set inclusion. When $\varepsilon = 0$, $\mathcal{P}(x) = \{\mu\}$ and the model reduces to the standard one with expected utility. When $\varepsilon = 1$, the investor is completely ignorant about the “true” distribution. The following proposition characterizes the optimal investment trigger.

¹⁰ See Bernanke (1983), Brennan and Schwartz (1985), and McDonald and Siegel (1986) for important early contributions. See Dixit and Pindyck (1994) for a textbook treatment.

¹¹ We emphasize that there is an important distinction between financial options and real options. The standard method of solving financial option problems is to assume that there exist spanning assets so that the complete markets contingent claims analysis can be applied. The real-options literature sometimes adopts the same spanning assumption and then use the financial option pricing technique. However, unlike financial options, real options and the underlying investment opportunities are often not traded in the market. We thus assume there is no spanning asset and instead follow Dixit and Pindyck (1994) by solving the agent's dynamic programming problem. Unlike the complete-markets financial options analysis, in the latter real-options approach (as in Dixit and Pindyck, 1994), the mean of asset payoffs matters for the real option exercise under incomplete markets. Miao and Wang (2007) further extend the real-options analysis to explicitly allow for non-spanned risk under incomplete markets and risk-averse decision makers.




Fig. 4. Investment timing under Knightian uncertainty and in the standard model. The upper (dashed) curve corresponds to the value function $F(x)$ in the standard model. The lower (solid) curve corresponds to the value function $V(x)$ under Knightian uncertainty.

Proposition 6. Assume (23). (i) For problem (20), the investment threshold x^* satisfies the equation¹²

$$x^* - I = \frac{\beta(1-\varepsilon)}{1-\beta} \int_{x^*}^b (x-x^*) d\mu. \tag{24}$$

Moreover, x^* decreases in ε . (ii) For problem (22), the investment threshold x^* satisfies the equation

$$x^* + \frac{\beta}{1-\beta} ((1-\varepsilon)E^\mu[x] + \varepsilon a) - I = \frac{\beta(1-\varepsilon)}{1-\beta} \int_{x^*}^b (x-x^*) d\mu. \tag{25}$$

Moreover, x^* increases in ε .

The interpretation of (24) is the following. The left-hand side of (24) represents the net benefit from investment. The right-hand side represents the opportunity cost of waiting. Because waiting has positive option value, the investment threshold exceeds the investment cost I . Eq. (24) states that at the investment threshold, the investor is indifferent between investing and waiting. It is also clear from (24) that because ambiguity lowers the option value, the right-hand side of (24) is less than that in the standard model with $\varepsilon = 0$. Moreover, an increase in ε lowers the investment threshold. Thus, an increase in ambiguity speeds up investment and a more ambiguity averse investor invests relatively earlier.

The interpretation of (25) is similar. The difference is that there is future uncertainty about profit opportunities of the investment project. Under the ε -contamination specification in (23), using (21) one can verify that the value of the investment project is given by

$$\Omega(x) = x + \frac{\beta}{1-\beta} ((1-\varepsilon)E^\mu[x] + \varepsilon a). \tag{26}$$

When ε is increased, ambiguity lowers both the project value represented by the left-hand side of (25) and the option value from waiting represented by the right-hand side of (25). Proposition 6 demonstrates that the former effect dominates so that the investor delays the investment.

It is interesting to note that when ε approaches 1, the investor has no idea about the true distribution of the profit. Ambiguity erodes away completely the option value from waiting. Specifically, for problem (20) in which there is no ambiguity about termination payoff, the investment threshold becomes $x^* = I$. For problem (22) in which there is ambiguity about both termination and continuation payoffs, the investment threshold satisfies $x^* + \beta/(1-\beta)a = I$. Note that for both

¹² We have implicitly assumed that there are parameter values such that there exists an interior solution. We will not state such an assumption explicitly both in this proposition and in Proposition 7.




Fig. 5. Firm exit under different degrees of Knightian uncertainty. The upper (dashed) curve corresponds to the value function $V^{\mathcal{P}_1}(x)$ and the lower (solid) curve corresponds to the value function $V^{\mathcal{P}_2}(x)$ where $\mathcal{P}_1 \subset \mathcal{P}_2$.

problems, the investor adopts the myopic NPV investment rule. Further, if the investor is also ambiguous about the future profits from the project after investment, the investor computes the NPV of the project according to the worst-case scenario, in which he believes that the cash flow in each period in the future takes the minimum value a .

3.2. Exit

Exit is an important problem in industrial organization and macroeconomics.¹³ We may describe a stylized exit model as follows. Consider a firm in an industry. The process $(x_t)_{t \geq 1}$ could be interpreted as a demand shock or a productivity shock. Staying in business at date t generates profits $\Pi(x_t)$ and incurs a fixed cost $c_f > 0$. The owner/manager may decide to exit and seek outside opportunities. Let the outside opportunity value be a constant $\gamma > 0$. Then the problem fits into our framework by setting

$$\Omega(x) = \gamma, \tag{27}$$

$$\pi(x) = \Pi(x) - c_f, \tag{28}$$

where we assume $\Pi(\cdot)$ is increasing and continuous.

According to the standard real-options approach, the exit trigger is lower than that predicted by the textbook Marshallian net present value principle. This implies that firms stay in business for a long period of time while absorbing operating losses. Only when the upside potential gain is low enough, will the firm not absorb losses and abandon operation. The standard real-options approach also predicts that an increase in risk in the sense of mean preserving spread raises the option value, and hence lowers the exit trigger. This implies that firms should stay in business longer in riskier situations, even though they suffer substantial losses. However, this prediction seems to be inconsistent with the large amount of quick exit in the IT industry in recent years.

The Knightian uncertainty theory may shed light on this issue. In recent years, due to economic recessions, firms are more ambiguous about the industry demand and their productivity. They are less sure about the likelihoods of when the economy will recover. Intuitively, the set of probability measures that firms may conceive is larger in recessions. Thus, by Proposition 4, the option value of the firm is lower and the exit trigger is higher. This induces firms to exit earlier (see Fig. 5).

The previous argument relies crucially on the fact that the outside value γ is a constant. In reality, there may be uncertainty about the outside value. For example, the outside value could represent the scrapping value of the firm and the firm is uncertain about its market value. The outside value could also represent the profit opportunity of a new business and the firm

¹³ See Hopenhayn (1992) for an industry equilibrium model of entry and exit.

is uncertain about this opportunity. Proposition 5 shows that, *ceteris paribus*, if the owner/manager is more ambiguous about the outside value, he will be more hesitant to exit. Thus, the overall effect of ambiguity on the exit timing depends on the relative degrees of ambiguity about different sources.

To illustrate, we consider a parametric example. We simply take $\Pi(x) = x \in X = [a, b]$. We still adopt the IID ε -contamination specification (23) for the process $(x_t)_{t \geq 0}$. When the outside value γ is a constant, the owner’s decision problem is described by the following Bellman equation:

$$V(x) = \max \left\{ \gamma, x - c_f + \beta \int_a^b V(x') \mathcal{P}(dx'; x) \right\}. \tag{29}$$

When the owner/manager is also ambiguous about the outside value, we adopt the following η -contamination specification for the outside value $\gamma \in [\underline{\gamma}, \bar{\gamma}]$,

$$\mathcal{Q} = \{(1 - \eta)v + \eta m : m \in \mathcal{M}([\underline{\gamma}, \bar{\gamma}])\}, \quad \eta \in [0, 1], \tag{30}$$

where v is a distribution over $[\underline{\gamma}, \bar{\gamma}]$ and may represent the “true” distribution of γ . The interpretation is that the owner/manager is not sure about the true distribution of the outside value and believes other distributions are possible with probability η . Note that η can be interpreted as a parameter measuring the degree of ambiguity and ambiguity aversion about the outside value. In this case, the owner/manager’s decision problem is described by the following Bellman equation:

$$V(x) = \max \left\{ \min_{q \in \mathcal{Q}} \int_{\underline{\gamma}}^{\bar{\gamma}} \gamma dq, x - c_f + \beta \int_a^b V(x') \mathcal{P}(dx'; x) \right\}. \tag{31}$$

The following proposition characterizes the solutions to problems (29) and (31).

Proposition 7. Assume (23) and (30). (i) For problem (29), the exit threshold x^* satisfies the equation

$$(1 - \beta)\gamma = x^* - c_f + \beta(1 - \varepsilon) \int_{x^*}^b (x - x^*) d\mu. \tag{32}$$

Moreover, x^* increases in ε . (ii) For problem (31), the exit threshold x^* satisfies the equation

$$(1 - \beta)((1 - \eta)E^v[\gamma] + \eta \underline{\gamma}) = x^* - c_f + \beta(1 - \varepsilon) \int_{x^*}^b (x - x^*) d\mu. \tag{33}$$

Moreover, x^* increases in ε and decreases in η .

The interpretation of (32) is the following. The left-hand side of (32) represents the per period outside value if the firm chooses to exit. The right-hand side represents the payoff if the firm chooses to stay. In particular, the first term represents the immediate profits and the second term represents the option value of waiting. At the exit threshold value, the owner/manager is indifferent between exit and stay. From (32), the impact of ambiguity is transparent. An increase in ε lowers the option value of stay in business by raising the exit threshold. Hence, the firm exits earlier.

The interpretation of (33) is similar. Note that, as one fixes ε and increases η , the outside value and the exit trigger are reduced. Thus, *ceteris paribus*, ambiguity about the outside value delays exit. If the owner/manager is more ambiguous about the outside value, it exits later. If one increases both ε and η , either effects may dominate and the overall effect on the exit timing depends on the relative degrees of ambiguity about these two sources of uncertainty.

Note that as in the investment problem, when ε is equal to 1, the option value of waiting to exit is equal to zero. Thus, the owner/manager just follows the simple myopic NPV rule, by using the worst-case scenario belief.

4. Discussions

Our analysis has two major limitations. First, Assumption 4 rules out the situation where ambiguity arises because the agent does not know something about the structure of payoffs. In Section 4.1, we relax this assumption. We adopt the learning model of Epstein and Schneider (2007) and derive the set of one-step-ahead conditionals \mathcal{P}_t from the agent’s learning process. We focus on the questions as to how \mathcal{P}_t evolves over time and what is the impact of ambiguity on option exercise in the long run.

The second limitation concerns the interpretation of our comparative statics results. We interpret an increase in the set of priors as either an increase in ambiguity aversion or an increase in the degree of ambiguity. This interpretation confounds ambiguity and ambiguity attitude. This limitation is inherent in the multiple-priors utility model and many other models of ambiguity such as the Choquet expected utility model of Schmeidler (1989) and the variational utility model of Maccheroni et al. (2006). In Section 4.2, we discuss a recent smooth ambiguity model of Klibanoff et al. (2005, 2009) which allows a separation between ambiguity and ambiguity attitude.¹⁴

¹⁴ The α -maxmin model of Ghirardato et al. (2004) and the neo-additive capacity model of Chateauneuf et al. (2007) also permit a limited separation.

4.1. Learning under ambiguity

Instead of considering a general setup as in Section 2, we use an example to illustrate the impact of learning on option exercise under ambiguity. The example is based on the job matching/quit model described in Section 10.10 in Stokey et al. (1989), which is a simplified discrete-time version of Jovanovic (1979). In each period, the risk-neutral worker may either continue with his current job match or quit the job and immediately obtain a constant one-time payment γ (outside option value). If he stays with his current job match, in each period, the match generates output that is either one with probability θ or zero with probability $(1-\theta)$. Let $S=\{0,1\}$ denote the state space and let state $s \in S$ denote the realized output: $s_t=1$ for one and $s_t=0$ for no output in period t . At time t , the worker's information consists of history $s^t=(s_1, s_2, \dots, s_t)$. This payoff is identically and independently distributed. The probability θ is constant but unknown. Based on past information about success, he may learn about the value of match productivity θ . The worker is infinitely lived and discounts future payoffs using the discount factor $0 < \beta < 1$.

We first consider the Bayesian learning case. Let p_t be the period- t conditional expected value of s_{t+1} given information s^t , i.e. $p_t = E_t(s_{t+1}) = E_t(\theta)$ for $t \geq 1$. Assume that the prior distribution p_0 follows a uniform distribution. Then, we have the following Bayesian belief updating rule:

$$p_{t+1} = \frac{(t+2)p_t + s_{t+1}}{t+3}, \quad t \geq 0. \tag{34}$$

Using (34) recursively, it is immediate to obtain $(t+3)p_{t+1} = 2p_0 + \sum_{j=1}^{t+1} s_j$. Let ϕ_{t+1} denote the empirical frequency: $\phi_{t+1} = \sum_{j=1}^{t+1} s_j / (t+1)$. Therefore, in the long run, the posterior probability of success is given by the empirical frequency of success, i.e. p_{t+1} converges to ϕ_∞ as $t \rightarrow \infty$ almost surely.

Next, we formulate a Bayesian worker's stationary dynamic programming problem as follows:

$$V(\phi_t) = \max\{\gamma, J(\phi_t)\}, \tag{35}$$

where

$$J(\phi_t) = p_t + \beta \left[p_t V\left(\frac{t\phi_t + 1}{t+1}\right) + (1-p_t)V\left(\frac{t\phi_t}{t+1}\right) \right]. \tag{36}$$

In the long run, the rule for job quit is given by a threshold policy with a cutoff level p^* , which solves $\gamma = p^* + \beta\gamma$. That is, $p^* = (1-\beta)\gamma$. If the empirical frequency of success ϕ is less than the cutoff value $p^* = (1-\beta)\gamma$, then the worker quits his job. Otherwise, he stays on the job.

We now turn to our model with ambiguity. Unlike the Bayesian learning model, the worker may lack confidence in his initial information about the environment. That is, he may not have a unique prior distribution over parameter values. In addition, the worker may be ambiguous about signals of payoffs given a parameter value. That is, he is unsure about how parameters are reflected in data. Thus, he considers that multiple likelihoods are possible.

To capture such learning under ambiguity, we adopt the learning model of Epstein and Schneider (2007). The Epstein and Schneider (2007) model is summarized by the tuple $(\Theta, \mathcal{M}_0, \mathcal{L}, \alpha)$ where Θ is the space of parameter θ , \mathcal{M}_0 is a set of prior probability measures on Θ , \mathcal{L} is a set of likelihoods, and $\alpha \in (0, 1)$ describes the speed of learning. Define a theory as a pair (μ_0, ℓ^t) , where $\mu_0 \in \mathcal{M}_0$ and $\ell^t = (\ell_1, \dots, \ell_t) \in \mathcal{L}^t$ is a sequence of likelihoods. Let $\mu_t(\cdot; s^t, \mu_0, \ell^t)$ denote the posterior derived from the theory (μ_0, ℓ^t) by Bayes' Rule, given the history s^t . This posterior can be computed recursively as follows:

$$d\mu_t(\cdot; s^t, \mu_0, \ell^t) = \frac{\ell_t(s_t|\cdot) d\mu_{t-1}(\cdot; s^{t-1}, \mu_0, \ell^{t-1})}{\int_{\Theta} \ell_t(s_t|\theta) d\mu_{t-1}(\theta; s^{t-1}, \mu_0, \ell^{t-1})}. \tag{37}$$

The set of posteriors contains posteriors that are based on theories not rejected by a likelihood ratio test:

$$\mathcal{M}_t^\alpha(s^t) = \left\{ \mu_t(\cdot; s^t, \mu_0, \ell^t) : \mu_0 \in \mathcal{M}_0, \ell^t \in \mathcal{L}^t, \int \prod_{j=1}^t \ell(s_j|\theta) d\mu_0(\theta) \geq \alpha \max_{\tilde{\mu}_0 \in \mathcal{M}_0, \tilde{\ell}^t \in \mathcal{L}^t} \int \prod_{j=1}^t \tilde{\ell}(s_j|\theta) d\tilde{\mu}_0(\theta) \right\}. \tag{38}$$

The set of one-step-conditional beliefs is defined by

$$\mathcal{P}_t(s^t) = \left\{ p_t(\cdot) = \int \ell(\cdot|\theta) d\mu_t(\theta) : \mu_t \in \mathcal{M}_t^\alpha(s^t), \ell \in \mathcal{L} \right\}. \tag{39}$$

Our specification is as follows. The worker perceives the success probability as $(\theta + \lambda_t)$, where $\theta \in [0, 1]$ is fixed and can be learned, while λ_t is driven by many poorly understood factors and can never be learned. We assume that $\lambda_t \in [-\bar{\lambda}, \bar{\lambda}]$ where $\bar{\lambda} \in (0, 1/2)$, and $\{\lambda_t : t \geq 1\}$ is i.i.d. We set $\Theta \equiv [\bar{\lambda}, 1-\bar{\lambda}]$. The set of priors \mathcal{M}_0 on Θ consists of all Dirac measures. For simplicity, we use $\theta \in \mathcal{M}_0$ to denote a Dirac measure on $\theta \in \Theta$. The set \mathcal{L} consists of all $\ell(\cdot|\theta)$ such that

$$\ell(s = 1|\theta) = \theta + \lambda \quad \text{for } |\lambda| \leq \bar{\lambda}. \tag{40}$$

The Dirac measure specification of the set of priors implies that the posterior set $\mathcal{M}_t^\alpha(s^t)$ also consists of Dirac measures, to which we turn next.

Let λ^t denote $(\lambda_1, \dots, \lambda_t)$. The likelihood of a sample s^t under some theory, here identified with a pair (θ, λ^t) , is given by

$$L(s^t, \theta, \lambda^t) = \prod_{j=1}^t (\theta + \lambda_j)^{s_j} (1 - \theta - \lambda_j)^{1-s_j}. \tag{41}$$

For fixed θ , the theory that maximizes this likelihood is given by $\lambda_j^* = \bar{\lambda}$ if $s_j = 1$, and $\lambda_j^* = -\bar{\lambda}$ if $s_j = 0$. By (38), the posterior set can be written as

$$\mathcal{M}_t^z = \left\{ \theta : \frac{1}{t} \log L(s^t, \theta, \lambda^{*t}) \geq \max_{\tilde{\theta} \in \Theta} \frac{1}{t} \log L(s^t, \tilde{\theta}, \lambda^{*t}) + \frac{1}{t} \log \alpha \right\}. \tag{42}$$

Using the empirical frequency of success $\phi_t \equiv \sum_{i=1}^t s_i / t$, we can derive

$$\frac{1}{t} \log L(s^t, \theta, \lambda^{*t}) = f(\theta; \phi_t) \equiv \phi_t \log(\theta + \bar{\lambda}) + (1 - \phi_t) \log(1 - \theta + \bar{\lambda}). \tag{43}$$

Thus, ϕ_t is sufficient to summarize the history s^t . It follows that the likelihood ratio criterion in (42) becomes

$$f(\theta; \phi_t) \geq \max_{\tilde{\theta} \in \Theta} f(\tilde{\theta}; \phi_t) + \frac{1}{t} \log \alpha. \tag{44}$$

We can easily show that the function f is strictly concave and achieves a unique maximum at the value $\theta^*(\phi_t)$ in Θ defined by

$$\theta^*(\phi_t) = \begin{cases} \bar{\lambda} & \text{if } \phi_t < \frac{2\bar{\lambda}}{1+2\bar{\lambda}}, \\ \phi_t + \bar{\lambda}(2\phi_t - 1) & \text{if } \phi_t \in \left[\frac{2\bar{\lambda}}{1+2\bar{\lambda}}, \frac{1}{1+2\bar{\lambda}} \right], \\ 1 - \bar{\lambda} & \text{if } \phi_t > \frac{1}{1+2\bar{\lambda}}. \end{cases} \tag{45}$$

Let $\theta_2(\phi_t) > \theta_1(\phi_t)$ be the two values such that Eq. (44) holds with equality. Let $\bar{\theta}(\phi_t) > \underline{\theta}(\phi_t)$ be the two values such that $[\underline{\theta}(\phi_t), \bar{\theta}(\phi_t)] = \Theta \cap [\theta_1(\phi_t), \theta_2(\phi_t)]$. Then the set of posteriors consists of Dirac measures over values in the interval is $[\underline{\theta}(\phi_t), \bar{\theta}(\phi_t)]$. Finally, consider the set of one-step-ahead conditionals \mathcal{P}_t . By (39) and (40), we can show that this set consists of all probability measures with conditional probability of $s_{t+1} = 1$ given in the interval $[\underline{p}_t(\phi_t), \bar{p}_t(\phi_t)]$, where

$$\underline{p}_t(\phi_t) = \underline{\theta}(\phi_t) - \bar{\lambda}, \quad \bar{p}_t(\phi_t) = \bar{\theta}(\phi_t) + \bar{\lambda}. \tag{46}$$

We may simply write

$$\mathcal{P}_t(\phi_t) = [\underline{p}_t(\phi_t), \bar{p}_t(\phi_t)]. \tag{47}$$

We now formulate an ambiguity averse worker’s dynamic programming problem as follows:

$$V_t(\phi_t) = \max\{\gamma J_t(\phi_t)\}, \tag{48}$$

where

$$J_t(\phi_t) = \min_{p_t \in \mathcal{P}_t(\phi_t)} p_t + \beta \left[p_t V_{t+1} \left(\frac{t\phi_t + 1}{t+1} \right) + (1 - p_t) V_{t+1} \left(\frac{t\phi_t}{t+1} \right) \right]. \tag{49}$$

Here $V_t(\phi_t)$ is the worker’s value function. From the preceding equation, we can see that the one-step-ahead conditional probability p_t of $s_{t+1} = 1$ also represents the conditional expected return on the task at time $(t + 1)$. Thus, we may interpret the multiplicity of $\mathcal{P}_t(\phi_t)$ as ambiguity about the task return.

We are particularly interested in the long-run impact of ambiguity on option exercise. One may conjecture that ambiguity may disappear in the long run due to learning. This situation does not happen in our model. The following proposition characterizes the long-run behavior.

Proposition 8. Suppose $\lim_{t \rightarrow \infty} \phi_t = \phi$ and $0 < \gamma < 1/(1 - \beta)$.¹⁵ Then in the long run as $t \rightarrow \infty$, (i) every sequence of posteriors from the set \mathcal{M}_t^z converges to the Dirac measure on $\theta^*(\phi)$ where $\theta^*(\cdot)$ is given by (45); (ii) the set of one-step-ahead beliefs converges to $[\theta^*(\phi) - \bar{\lambda}, \theta^*(\phi) + \bar{\lambda}]$; and (iii) the value function converges to

$$V(\phi) \equiv \begin{cases} \gamma & \text{if } \theta^*(\phi) < \gamma(1 - \beta) + \bar{\lambda}, \\ \frac{\theta^*(\phi) - \bar{\lambda}}{1 - \beta} & \text{if } \theta^*(\phi) \geq \gamma(1 - \beta) + \bar{\lambda}. \end{cases} \tag{50}$$

Part (i) of Proposition 8 shows that the posterior set shrinks to a singleton $\theta^*(\phi)$ in the long run. From Eq. (45), we deduce that this singleton $\theta^*(\phi)$ is generally different from the long-run empirical frequency ϕ , in contrast to the Bayesian model. In particular, for ϕ_t in the intermediate region, i.e. $2\bar{\lambda}/(1 + 2\bar{\lambda}) \leq \phi_t \leq 1/(1 + 2\bar{\lambda})$, Eq. (45) shows that $\theta^*(\phi_t) = \phi_t + \bar{\lambda}(2\phi_t - 1)$. In the limit, $\theta^*(\phi)$ is not equal to ϕ unless $\phi = 1/2$. Only when $\phi = 1/2$, the posterior converges to the truth in that $\theta^*(\phi) = \phi$.

¹⁵ The latter assumption ensures that neither immediate quitting nor never quitting is optimal in the standard Bayesian model.

Importantly, Part (ii) of Proposition 8 shows that ambiguity persists in the long run because the set of one-step-ahead beliefs does not converge to a singleton. Instead, it converges to an interval with width $2\bar{\lambda}$. This also implies that the worker perceives that the long-run expected task return lies in an interval with width $2\bar{\lambda}$.

Finally, we explain the intuition for Part (iii) of Proposition 8 as follows. Using Eqs. (48) and (49), we deduce that the Bellman equation in the long run limit satisfies

$$V(\phi) = \max \left\{ \gamma, \min_{p \in \mathcal{P}(\phi)} p + \beta V(\phi) \right\}. \tag{51}$$

Thus, the cutoff value in the long-run limit satisfies $\gamma = \min_{p \in \mathcal{P}(\phi)} p + \beta \gamma$. Note that $\min_{p \in \mathcal{P}(\phi)} p = \theta^*(\phi) - \bar{\lambda}$ by Parts (i)–(ii) of Proposition 8. Combining these two facts, we obtain the cutoff rule expressed in $\theta^*(\phi)$ in Part (iii). The worker quits his job if and only if $\theta^*(\phi) \leq \bar{\lambda} + \gamma(1 - \beta)$. Since this cutoff value $\bar{\lambda} + \gamma(1 - \beta)$ is higher than the Bayesian cutoff value $p^* = \gamma(1 - \beta)$ by $\bar{\lambda}$, the ambiguity averse worker is more likely to quit the job than the Bayesian worker even in the long run. Even in the long run, ambiguity persists, the ambiguity averse worker undervalues the expected payoff on the job and hence prefers to quit sooner. The intuition for this result is similar to what we described earlier.

4.2. An alternative model of ambiguity

Most dynamic ambiguity models extend the multiple-priors utility model of Gilboa and Schmeidler (1989) and the Choquet expected utility model of Schmeidler (1989). Our paper follows one of such models proposed by Epstein and Wang (1994) and Epstein and Schneider (2003). All these models suffer from the limitation that ambiguity and ambiguity attitude are confounded. Recently, Klibanoff et al. (2005, 2009) propose a smooth ambiguity model that allows a separation between these two distinct concepts, which is analogous to the separation between risk and risk attitude. In this model, an agent’s preferences can be represented by the following recursive utility:

$$U_t(c; s^t) = u(c(s^t)) + \beta v^{-1} \left[\int_{\Theta} v \left(\int_S U_{t+1}(c; s^t, s_{t+1}) d\pi_{\theta}(s_{t+1}; s^t) \right) d\mu(\theta | s^t) \right]. \tag{52}$$

The interpretation is the following. The increasing function u represents risk preferences, $\pi_{\theta}(\cdot | s^t)$ represents one-step-ahead belief given the history s^t , and $\mu(\cdot | s^t)$ represents the posterior given the history s^t . Posteriors are updated using the Bayes rule. The set Θ , which describes the multiplicity of π_{θ} , captures ambiguity. Importantly, the increasing function v captures attitude towards ambiguity. Thus, ambiguity and ambiguity attitude are separated. Ju and Miao (2010) and Hayashi and Miao (2010) generalize (Klibanoff et al., 2005, 2009) model to allow for a three-way separation among intertemporal substitution, risk aversion, and ambiguity aversion.

If v is a concave function, then the agent is ambiguity averse. Comparative ambiguity aversion can be characterized by the relative concavity of v . In particular, agent A is more ambiguous averse than agent B if they share the same risk preferences u and v^A is a monotone concave transformation of v^B . The definition and characterization of the notion of “degree of ambiguity” and “comparative ambiguity” in the smooth ambiguity model have not been fully developed in the literature.¹⁶ Thus, in our application below we do not consider this issue.

We now apply the smooth ambiguity model to the job matching problem in the previous subsection. We can describe the worker’s dynamic programming problem as follows:

$$V_t(\phi_t) = \max \{ \gamma, J_t(\phi_t) \}, \tag{53}$$

where

$$J_t(\phi_t) = v^{-1} \left(\int_{\Theta} v \left(\theta + \beta \left[\theta V_{t+1} \left(\frac{t\phi_t + 1}{t+1} \right) + (1 - \theta) V_{t+1} \left(\frac{t\phi_t}{t+1} \right) \right] \right) d\mu_t(\theta) \right). \tag{54}$$

We claim that a more ambiguity averse agent quits the job sooner. This result is consistent with what we derived earlier using the multiple-priors utility model. An advantage of the smooth ambiguity model is that the comparative static result refers to comparative ambiguity attitude only, holding ambiguity fixed.

Instead of providing a complete proof of the preceding claim, which is not difficult but lengthy, we sketch the key steps of the proof here. First, by an analysis similar to that in Stokey et al. (1989, Chapter 10.10), we can show that V_t is increasing in ϕ_t . Thus, the agent’s exit decision is characterized by a trigger policy. That is, at any time t , there is a cutoff value ϕ_t^* such that if $\phi_t < \phi_t^*$ the worker quits the job. The cutoff value is determined by the equation $\gamma = J_t(\phi_t^*)$. We now conduct comparative statics analysis. Suppose agent A is more ambiguity averse than agent B in the sense that his v^A is a monotone concave transformation of v^B . We observe from Eq. (54) that J_t is the certainty equivalent of v . The standard risk analysis and the dynamic programming argument similar to the proof of Propositions 3–4 imply $J_t^A \leq J_t^B$. Therefore, we conclude that the trigger values satisfy the inequality $\phi_t^{A*} \geq \phi_t^{B*}$, implying that the more ambiguity averse agent quits job earlier.

We next turn to the long-run behavior. It is easy to show that unlike the previous model, ambiguity does not persist in the long run and $V_t(\phi_t)$ converges to a limit in the Bayesian model, if v^{-1} is Lipschitz. Klibanoff et al. (2009) establish a more

¹⁶ In a work in progress, Jewitt and Mukerji (2006) study this issue.

general result. The crucial condition is that the parameter space is bounded. They show that ambiguity may persist in the long run if the parameter space is unbounded.

5. Conclusion

Many economic decisions can be described as an option exercise problem. In this problem, uncertainty plays an important role. In standard expected utility models, there is no meaningful distinction between risk and uncertainty in the sense attributed to Knight (1921). To afford this distinction, we apply the multiple-priors utility model. We formulate the option exercise decision as a general optimal stopping problem. While the standard analysis shows that risk increases option value, we show that ambiguity lowers the option value. Moreover, the impact of ambiguity on the option exercise timing depends crucially on whether the agent has ambiguity about termination payoffs after option exercise. If uncertainty is fully resolved after option exercise, then an increase in ambiguity speeds up option exercise and a more ambiguity averse agent exercises the option earlier. However, if the agent is ambiguous about the termination payoff, then the agent may delay option exercise if this ambiguity dominates ambiguity about continuation.

We apply our general model to real investment and exit problems. For the investment problem, we show that if the project value is modeled as a lump-sum value and uncertainty over this value is fully resolved after investment, then ambiguity accelerates investment. However, if the project value is modeled as a discounted sum of future uncertain profit flows and the agent is ambiguous about these profits, then ambiguity may delay investment. For the exit problem, we presume that there are two sources of uncertainty—outside value and profit opportunities if the firm stays in business. The firm's owner/manager may be ambiguous about both sources. We show that ambiguity may delay or accelerate exit, depending on which source of ambiguity dominates. We also show that for both problems, the myopic NPV rule often recommended by the business textbooks and investment advisors may be optimal for an agent having an extremely high degree of ambiguity aversion.

Finally, using a job matching example, we analyze the impact of learning under ambiguity using the Epstein and Schneider (2007) model. We also apply the recent smooth ambiguity model of Klibanoff et al. (2005, 2009) to this example. Further analysis and applications along this line would be an interesting research topic.

Acknowledgements

We would like to thank Bart Lipman, two anonymous referees, Tim Cogley (Editor), and seminar participants at BU and Columbia for helpful comments. First Version: December 2003.

Appendix A. Proofs

Proof of Proposition 1. The proof is similar to that of Proposition 2. So we omit it. \square

Proof of Proposition 2. Let $C(X)$ denote the space of all bounded and continuous functions endowed with the sup norm. $C(X)$ is a Banach space. Define an operator T as follows:

$$Tv(x) = \max \left\{ \Omega(x), \pi(x) + \beta \int v(x') \mathcal{P}(dx'; x) \right\}, \quad v \in C(X).$$

Then it can be verified that T maps $C(X)$ into itself. Moreover, T satisfies the Blackwell sufficient condition and hence is a contraction mapping. By the Contraction Mapping Theorem, T has a unique fixed point $V \in C(X)$ which solves the problem (9) (see Stokey et al., 1989, Theorems 3.1 and 3.2).

Next, let $C'(X) \subset C(X)$ be the set of bounded continuous and increasing functions. One can show that T maps any increasing function $C'(X)$ into an increasing function in $C'(X)$. Since $C'(X)$ is a closed subset of $C(X)$, by Corollary 1 in Stokey et al. (1989, p. 52), the fixed point of T , V , is also increasing. The remaining part of the proposition is trivial and follows from similar intuition illustrated in Fig. 1. \square

Remark. The Contraction Mapping Theorem also implies that $\lim_{n \rightarrow \infty} T^n v = V$ for any function $v \in C(X)$.

Proof of Proposition 3. Let $v \in C(X)$ satisfy $v \leq F$. Since $P(x) \in \mathcal{P}(x)$,

$$\int v(x') \mathcal{P}(dx'; x) = \min_{Q(\cdot|x) \in \mathcal{P}(x)} \int v(x') Q(dx'; x) \leq \int v(x') P(dx'; x) \leq \int F(x') P(dx'; x).$$

Thus,

$$Tv(x) = \max \left\{ \Omega(x), \pi(x) + \beta \int v(x') \mathcal{P}(dx'; x) \right\} \leq \max \left\{ \Omega(x), \pi(x) + \beta \int F(x') P(dx'; x) \right\} = F(x).$$

It follows from induction that the fixed point of T , V , must also satisfy $V \leq F$. The remaining part of the proposition follows from this fact and Fig. 2. \square

Proof of Proposition 4. Define the operator $T^{P_1} : C(X) \rightarrow C(X)$ by

$$T^{P_1} v(x) = \max \left\{ \Omega(x), \pi(x) + \beta \int v(x') \mathcal{P}_1(dx'; x) \right\}, \quad v \in C(X).$$

Similarly, define an operator $T^{P_2} : C(X) \rightarrow C(X)$ corresponding to \mathcal{P}_2 . Take any functions $v_1, v_2 \in C(X)$ such that $v_1 \geq v_2$, it can be shown that $T^{P_1} v_1(x) \geq T^{P_2} v_2(x)$. By induction, the fixed points V^{P_1} and V^{P_2} must satisfy $V^{P_1} \geq V^{P_2}$. The remaining part of the proposition follows from Fig. 3. \square

Proof of Proposition 5. First, one can use the standard dynamic programming technique similar to that used in Propositions 2–4 to show that the value functions $V^{\mathcal{Q}_1}$ and $V^{\mathcal{Q}_2}$ are increasing and $V^{\mathcal{Q}_1} \geq V^{\mathcal{Q}_2}$. To show $x^{\mathcal{Q}_1} \geq x^{\mathcal{Q}_2}$, let $G^i(x) = V(x) - \min_{q \in \mathcal{Q}_i} \int \Omega dq$ for $i=1,2$. Then from (14), we can derive that

$$G^i(x) = \max \left\{ 0, \pi(x) - (1-\beta) \min_{q \in \mathcal{Q}_i} \int \Omega dq + \beta \int G^i(x') \mathcal{P}(dx'; x) \right\}.$$

Again, by the standard dynamic programming technique, we can show that G^i is increasing and $G^2(x) \geq G^1(x)$. The threshold values $x^{\mathcal{Q}_1}$ are determined by the equation

$$0 = \pi(x) - (1-\beta) \min_{q \in \mathcal{Q}_i} \int \Omega dq + \beta \int G^i(x') \mathcal{P}(dx'; x).$$

Since $\min_{q \in \mathcal{Q}_1} \int \Omega dq \geq \min_{q \in \mathcal{Q}_2} \int \Omega dq$ and $G^2(x) \geq G^1(x)$, we have $x^{\mathcal{Q}_1} \geq x^{\mathcal{Q}_2}$. \square

Proof of Proposition 6. Because $V(x)$ is an increasing and continuous function, the optimal investment rule is described as a trigger policy whereby the investor invests the first time the process $(x_t)_{t \geq 0}$ hits a threshold value x^* . We now determine x^* and focus on problem (20). Problem (22) can be analyzed similarly. Now, for problem (20), $V(x)$ satisfies

$$V(x) = \begin{cases} x - I & \text{if } x \geq x^*, \\ \beta \int_a^b V(x') \mathcal{P}(dx'; x) & \text{if } x < x^*. \end{cases} \tag{A.1}$$

At the threshold value x^* , we have

$$x^* - I = \beta \int_a^b V(x') \mathcal{P}(dx'; x). \tag{A.2}$$

According to the IID ε -contamination specification (23), we have

$$x^* - I = \beta(1-\varepsilon) \int_a^b V(x) d\mu + \beta\varepsilon \min_{m \in \mathcal{M}([a,b])} \int_a^b V(x) dm.$$

Since the minimum of $V(x)$ is $\beta \int_a^b V(x') \mathcal{P}(dx'; x)$, which is equal to $x^* - I$ by (A.2), we can rewrite the preceding equation as

$$\begin{aligned} x^* - I &= \beta(1-\varepsilon) \int_a^b V(x) d\mu + \beta\varepsilon(x^* - I) = \beta(1-\varepsilon) \left[\int_a^{x^*} V(x) d\mu + \int_{x^*}^b V(x) d\mu \right] + \beta\varepsilon(x^* - I) \\ &= \beta(1-\varepsilon) \left[\int_a^{x^*} (x^* - I) d\mu + \int_{x^*}^b (x - I) d\mu \right] + \beta\varepsilon(x^* - I). \end{aligned}$$

Note that in the last equality, we have used (A.1) and (A.2). Rearranging yields the desired result (24). The comparative statics result follows from the implicit function theorem applied to (24). \square

Proof of Proposition 7. The proof is similar to that of Proposition 6. We consider part (i) first. The value function $V(x)$ satisfies

$$V(x) = \begin{cases} x - c_f + \beta \int_a^b V(x') \mathcal{P}(dx'; x) & \text{if } x \geq x^*, \\ \gamma & \text{if } x < x^*. \end{cases} \tag{A.3}$$

At the threshold value x^* , we have

$$\gamma = x^* - c_f + \beta \int_a^b V(x') \mathcal{P}(dx'; x). \tag{A.4}$$

Given the IID ε -contamination specification (23), we can derive

$$\gamma = x^* - c_f + \beta(1-\varepsilon) \int_a^b V(x) d\mu + \beta\varepsilon \min_{m \in \mathcal{M}([a,b])} \int_a^b V(x) dm.$$

Since the minimum of $V(x)$ is γ by (A.3), we can rewrite the preceding equation as

$$\begin{aligned} \gamma &= x^* - c_f + \beta(1-\varepsilon) \int_a^b V(x) d\mu + \beta\varepsilon\gamma = x^* - c_f + \beta(1-\varepsilon) \left[\int_a^{x^*} V(x) d\mu + \int_{x^*}^b V(x) d\mu \right] + \beta\varepsilon\gamma \\ &= x^* - c_f + \beta(1-\varepsilon) \left[\int_a^{x^*} \gamma d\mu + \int_{x^*}^b (x + \gamma - x^*) d\mu \right] + \beta\varepsilon\gamma. \end{aligned}$$

Here the last equality follows from (A.3) and (A.4). Simplifying yields (32). The comparative static result follows from simple algebra.

For part (ii), given the η -contamination specification in (30), we have

$$\min_{q \in \underline{Q}} \int_{\underline{\gamma}}^{\bar{\gamma}} \gamma dq = (1-\eta)E^V[\underline{\gamma}] + \eta\underline{\gamma}.$$

Eq. (33) follows from a similar argument for (32). The comparative statics result follows from simple algebra. \square

Proof of Proposition 8. Part (i) follows from Eq. (45). Turn to part (ii). It follows from (44) that

$$\lim_{t \rightarrow \infty} \theta(\phi_t) = \lim_{t \rightarrow \infty} \bar{\theta}(\phi_t) = \theta^*(\phi).$$

Thus, the result follows from Eqs. (46) and (47). Finally, consider part (iii). In the limit, the Bellman equations (48) and (49) become

$$V(\phi) = \max\{\gamma, J(\phi)\},$$

where

$$J(\phi) = \min_{p \in [\theta^*(\phi) - \bar{\lambda}, \theta^*(\phi) + \bar{\lambda}]} p + \beta V(\phi) = \theta^*(\phi) - \bar{\lambda} + \beta V(\phi).$$

Thus,

$$V(\phi) = \max\{\gamma, \theta^*(\phi) - \bar{\lambda} + \beta V(\phi)\}.$$

One can verify that the function given in the proposition satisfies the preceding equation. \square

References

- Anderson, E.W., Hansen, L.P., Sargent, T.J., 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1, 68–123.
- Backus, D., Routledge, B., Zin, S., 2004. Exotic preferences for macroeconomists. *NBER Macroeconomics Annual* 19, 319–390.
- Bernanke, B.S., 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* 98, 85–106.
- Brennan, M.J., Schwartz, E.S., 1985. Evaluating natural resource investments. *Journal of Business* 58, 135–157.
- Chateauneuf, A., Eichberger, J., Grant, S., 2007. Choice under uncertainty with the best and worst in mind: neo-additive capacities. *Journal of Economic Theory* 137, 538–567.
- Chen, Z., Epstein, L.G., 2002. Ambiguity, risk and asset returns in continuous time. *Econometrica* 70, 1403–1443.
- Dixit, A., 1992. Investment and hysteresis. *Journal of Economic Perspectives* 6, 107–132.
- Dixit, A., Pindyck, R., 1994. *Investment under Uncertainty*. Princeton University Press, Princeton, NJ.
- Ellsberg, D., 1961. Risk, ambiguity and the Savage axiom. *Quarterly Journal of Economics* 75, 643–669.
- Epstein, L.G., 1999. A definition of uncertainty aversion. *Review of Economic Studies* 66, 579–608.
- Epstein, L.G., Miao, J., 2003. A two-person dynamic equilibrium under ambiguity. *Journal of Economic Dynamics and Control* 27, 1253–1288.
- Epstein, L.G., Schneider, M., 2003. Recursive multiple-priors. *Journal of Economic Theory* 113, 1–31.
- Epstein, L.G., Schneider, M., 2006. Ambiguity, information quality and asset pricing. *Journal of Finance* 63, 197–228.
- Epstein, L.G., Schneider, M., 2007. Learning under ambiguity. *Review of Economic Studies* 74, 1275–1303.
- Epstein, L.G., Wang, T., 1994. Intertemporal asset pricing under Knightian uncertainty. *Econometrica* 62, 283–322.
- Epstein, L.G., Wang, T., 1995. Uncertainty, risk-neutral measures and security price booms and crashes. *Journal of Economic Theory* 67, 40–82.
- Epstein, L.G., Zhang, J., 2001. Subjective probabilities on subjectively unambiguous events. *Econometrica* 69, 265–306.
- Ghirardato, P., Maccheroni, F., Marinacci, M., 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118, 133–174.
- Ghirardato, P., Marinacci, M., 2002. Ambiguity made precise: a comparative foundation. *Journal of Economic Theory* 102, 251–289.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique priors. *Journal of Mathematical Economics* 18, 141–153.
- Hansen, L.P., Sargent, T.J., 2001. Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics* 4, 519–535.
- Hayashi, T., Miao, J., 2010. Intertemporal Substitution and Recursive Smooth Ambiguity Preferences. Working paper, Boston University.
- Hopenhayn, H., 1992. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 60, 1127–1150.
- Jewitt, I., Mukerji, S., 2006. Ordering ambiguous acts. Working paper, Oxford University.
- Jovanovic, J., 1979. Job matching and the theory of turnover. *Journal of Political Economy* 87, 972–990.
- Ju, N., Miao, J., 2010. Learning, ambiguity, and asset returns. Working paper, Boston University.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2009. Recursive smooth ambiguity preferences. *Journal of Economic Theory* 144, 930–976.
- Knight, F., 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston.
- Kogan, L., Wang, T., 2002. A simple theory of asset pricing under model uncertainty. MIT Working Paper.
- Maccheroni, F., Marinacci, M., Rustichini, A., 2006. Ambiguity aversion, robustness and the variational representation of preferences. *Econometrica* 74, 1447–1498.
- McDonald, R., Siegel, D., 1986. The value of waiting to invest. *Quarterly Journal of Economics* 101, 707–728.
- Miao, J., 2004. A note on consumption and savings under Knightian uncertainty. *Annals of Economics and Finance* 5, 299–311.
- Miao, J., 2009. Ambiguity, risk and portfolio choice under incomplete information. *Annals of Economics and Finance* 10, 257–279.

- Miao, J., Wang, N., 2007. Investment, consumption and hedging under incomplete markets. *Journal of Financial Economics* 86, 608–642.
- Mukerji, S., Tallon, J., 2004. Ambiguity aversion and the absence of wage indexation. *Journal of Monetary Economics* 51, 653–670.
- Nishimura, K.G., Ozaki, H., 2004. Search and Knightian uncertainty. *Journal of Economic Theory* 119, 299–333.
- Nishimura, K.G., Ozaki, H., 2007. Irreversible investment and Knightian uncertainty. *Journal of Economic Theory* 136, 668–694.
- Riedel, F., 2009. Optimal stopping with multiple priors. *Econometrica* 77, 857–908.
- Routledge, B., Zin, S., 2009. Model uncertainty and liquidity. *Review of Economic Dynamics* 12, 543–566.
- Schmeidler, D., 1989. Subjective probability and expected utility without additivity. *Econometrica* 57, 571–587.
- Stokey, N., Lucas, R.E., Prescott, E.C., 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.