informs
http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 65, No. 4, July–August 2017, pp. 1011–1028
ISSN 0030-364X (print), ISSN 1526-5463 (online)

METHODS

Risk Aversion, Information Acquisition, and Technology Adoption

James E. Smith,[a] Canan Ulu[b]

[a] Fuqua School of Business, Duke University, Durham, North Carolina 27708; [b] McDonough School of Business, Georgetown University, Washington, DC 20057

Contact: jes9@duke.edu (JES); canan.ulu@georgetown.edu (CU)

Received: July 29, 2015
Revised: March 16, 2016; October 3, 2016
Accepted: December 17, 2016
Published Online in Articles in Advance: June 15, 2017

Subject Classifications: dynamic programming: models; decision analysis: sequential
Area of Review: Decision Analysis

https://doi.org/10.1287/opre.2017.1601

Abstract. We use a dynamic programming model to study the impacts of risk aversion on information acquisition in technology adoption decisions. In this model, the benefit of the technology is uncertain and, in each period, the decision maker (DM) may adopt the technology, reject the technology, or pay to acquire a signal about the benefit of the technology. The dynamic programming state variables are the DM's wealth and a probability distribution that describes the DM's beliefs about the benefit of the technology; these distributions are updated over time using Bayes' rule. If the signal-generating process satisfies the monotone-likelihood ratio property and the DM is risk neutral, the value functions and policies satisfy natural monotonicity properties: a likelihood-ratio improvement in the distribution on benefits leads to an increase in the value function and moves the DM away from rejection and toward adoption. With risk aversion, the value functions (now representing expected utilities) will be monotonic, but the policies need not be monotonic, even with reasonable utility functions. However, if we assume the DM exhibits decreasing absolute risk aversion and is not "too risk averse," the policies can be shown to be monotonic. Establishing these structural properties requires the use of some novel proof techniques that may prove useful in other contexts. We also study the impact of changing risk attitudes on the optimal policy.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2017.1601.

Keywords: dynamic programming • Bayesian decision analysis • sequential

1. Introduction

New technologies often promise uncertain benefits, sometimes with high costs. For example, consider the Tesla Model S luxury electric sedan. The Tesla boasts outstanding styling, performance, and efficiency (with a miles-per-gallon-equivalent of 84 miles per gallon) and earned an incredible 99 out of a possible 100 points in *Consumer Reports'* automobile ratings. However, as a new car, the Tesla's reliability is uncertain and its price tag ($70–$100,000) will give most consumers pause. Even among consumers who can afford a Tesla, many will wait to see how the car performs over the next few years to see if its benefits justify its cost. Similar dilemmas are faced by a farmer considering planting a new variety of soybeans or corn, an electric utility considering building a power plant based on a new technology, or an entrepreneur considering investing his life savings to attempt to commercialize a new idea. The stakes are high: should they act now or wait and learn more before deciding?

In this paper, we study the impact of uncertainty about the benefits of a technology on adoption and information gathering decisions, under risk aversion. Our starting point is the dynamic programming (DP) model studied in Ulu and Smith (2009, hereafter US) that builds on and generalizes the classic model of McCardle (1985). In this model, in each period, the decision maker (DM) must decide whether to adopt or reject a new technology or to wait and gather additional information about the benefits of the technology. Information gathering is costly and modeled as receiving a signal (e.g., a new review) about the benefit of the technology. After observing a signal, the DM updates her distribution on the technology's benefits using Bayes' rule. In McCardle (1985), the uncertainty in each period is described by a univariate summary statistic and consequently the DP has a univariate state variable. US (2009) considers a general model of learning where the state variable is the DM's probability distribution on the benefits of the technology; these probability distributions are ordered by likelihood-ratio (LR) dominance.

Our focus in this paper is on structural properties of the model and, in particular, the effects of risk aversion on these structural properties. Risk aversion is an important consideration in practice, particularly when the new technology (for example, the Tesla or the power plant) is expensive and when the stakes are high for the DM (as a farmer's choice of variety to plant or the entrepreneur's decision to invest his life savings in a startup). Many of the structural properties of the model generalize naturally from the risk-neutral to the risk-averse case. In particular, we can show the value functions are monotonic using essentially the same arguments as the risk-neutral case: if we assume the signal-generating process satisfies the monotone likelihood ratio (MLR) property, LR improvements in the prior lead to an increase the DM's expected utility. Convexity and some other comparative statics and convergence results also generalize in a straightforward manner to the risk-averse case.

However, the structure of the optimal policy is harder to characterize with risk aversion. If the DM is risk neutral or risk seeking, we can show that the optimal policy is monotonic in the sense that if it is optimal to adopt (or reject) with one distribution, it is optimal to adopt (reject) with all distributions that LR-dominate (or are LR-dominated by) the first distribution. This is shown using a supermodularity argument: the expected values for different actions have LR-increasing differences. With risk aversion, the corresponding utility differences need not be LR-increasing and the optimal policies need not be monotonic, even for reasonable utility functions. For example, we may have a DM for whom it is optimal to *wait* and gather more information about the technology with one prior distribution but optimal to *adopt* with a worse (i.e., a LR-dominated) prior. However, if we assume that the utility functions exhibit decreasing absolute risk aversion and are not "too risk averse" (in a sense to be made precise later), we can show that these policy differences are "sLR-increasing." This sLR-increasing property is sufficient to ensure that the utility differences are single crossing and the policies are monotonic, as in the risk-neutral case. Though the main contribution of the paper is the study of risk aversion in the technology adoption problem, this s-increasing property may prove useful when studying other DP models.

We provide a brief review of related literature in the remainder of this section. In Section 2, we describe the model and introduce a numerical example that we will use to illustrate the results of the paper. In Section 3, we briefly review some key properties of LR-dominance and LR-increasing functions. In Section 4, we discuss some results for the risk-averse model that are straightforward generalizations of the risk-neutral model. In Section 5, we study the structure of the policies in the risk-neutral and risk-averse models and show how risk aversion can lead to nonmonotonic policies. In Section 6, we develop the idea of s-increasing and sLR-increasing functions and show that s-increasing properties are preserved and propagated in recursive Bayesian models. In Section 7, we show that if the DM is not too risk averse, the policy differences will be sLR-increasing, which implies the optimal policies will be monotonic under risk aversion. In Section 8, we study the effects of changing risk attitudes on the optimal values and policies. In Section 9, we consider the model with discounting and briefly discuss the extension to include multiple information sources.

1.1. Literature Review

The literature on technology adoption is vast and spans a number of fields; Rogers (2003) provides a thorough review of the early literature. As discussed earlier, we focus on information acquisition in technology adoption decisions, following McCardle (1985) and US (2009). Jensen (1982) earlier studied a technology model where the uncertain technology value could take on just two values. There have been a number of generalizations and variations on McCardle's (1985) model. For example, Lippman and McCardle (1987) study how changes in the "informational returns to scale" in the McCardle (1985) model affect the timing of adoption decisions. Cho and McCardle (2009) study information acquisition in technology adoption decisions when there are multiple dependent technologies. The model of McCardle (1985) can also be viewed as a variation on the sequential hypothesis testing problem studied in Wald (1945) where the uncertainty is the economic benefit associated with adopting the technology rather than the truth of posited null and alternative hypotheses.

Of course, risk aversion and information gathering are central themes in decision analysis. LaValle (1968) provides an early study of the value of information with risk aversion; see also the review in Hilton (1981). The conclusions in this literature are mostly negative: e.g., there is no monotonic relationship between the degree of absolute or relative risk aversion and the value of information (Hilton 1981, Theorem 2). Others have made progress focusing on specific problems. For example, Bickel (2008) studies a single-period two-action linear-loss problem (essentially accept a risk or reject it) with exponential utility, focusing on the case with normally distributed uncertainties. Abbas et al. (2013) also study the effects of risk aversion on the value of information in a single-period problem with two actions; they find that the value of imperfect information may increase or decrease with the DM's degree of risk aversion. Here we consider a dynamic version of the two-action problem with repeated opportunities for information gathering. We consider general utilities and provide conditions that lead to some positive results about the structure of information gathering policies.

Risk aversion is also of growing interest in the broader operations research and operations management literature. For example, Eeckhoudt et al. (1995) study the impact of risk aversion in the classic newsvendor problem. Similarly, Chen et al. (2007) study the impact of risk aversion in classic dynamic inventory management models, focusing on structural properties of the models. Zhu and Kapuscinski (2015) study a multiperiod, multinational, risk-averse newsvendor facing exchange rate risks. The latter two studies use a present certainty equivalent framework (Smith 1998) that assumes an exponential utility function and leads to tractable DP models; Eeckhoudt et al. (1995) consider a more general utility analysis but in a single period setting. Our paper is analogous to these in that we study the effects of incorporating risk aversion in a classic model in the operations research literature.

Finally, this paper builds on and contributes to the literature on structural properties of DPs. Smith and McCardle (2002) synthesize some of these structural property results for DPs and relate these results to stochastic dominance. Lovejoy (1987a, 1987b) provides sufficient conditions for monotonic policies for DPs and partially observable DPs, exploiting super/submodularity results (e.g., Topkis 1978). As mentioned above, these super/submodularity arguments do not work in the technology adoption model with risk aversion. Milgrom and Shannon (1994) show that single-crossing policy differences are sufficient to obtain monotone policies, though they focus on single-period problems rather than DPs (see also Athey 2002). As we will see, the s-increasing property is convenient for use with DP models and enables the use of single-crossing arguments to establish the monotonicity of policies.

2. The Model

In this section, we begin by describing the general model. We then present a specific numerical example that we will use to illustrate our later results.

2.1. The Model

A DM is contemplating purchasing a new technology that yields an uncertain benefit denoted by $\theta \in \Theta \subseteq \mathbb{R}$; we can think of θ as representing the net present value of the stream of benefits provided by the technology. The DM's beliefs about the benefit of the technology are described by a probability distribution. For ease of notation, we will assume the DM's probability distribution is continuous and has a density π over Θ. For discrete spaces, we can interpret π as a probability mass function and consider sums instead of integrals; more general probability measures could also be considered.

Time is discrete. In each period, the DM starts with a prior distribution π and must choose whether to adopt the technology, reject it, or gather additional information:

- If she adopts the technology, she receives an uncertain benefit θ with distribution π.[1]
- If she rejects the technology, she receives nothing and stops gathering information.
- If she waits and gathers additional information, she pays c ($c \geqslant 0$) in that period and observes a signal $x \in X$, drawn with likelihood function $L(x \mid \theta)$. After observing signal x, the DM proceeds to the next period with a new distribution on benefits $\Pi(\theta; \pi, x)$ given by updating the prior π using Bayes' rule,

$$\Pi(\theta; \pi, x) = \frac{L(x \mid \theta)\pi(\theta)}{f(x; \pi)},$$

where $f(x;\pi)$ is the predictive distribution for signals x, $f(x;\pi) = \int_\theta L(x \mid \theta)\pi(\theta)\,d\theta$. We will assume $L(x \mid \theta) > 0$ for all x and θ; this ensures that the predictive distributions satisfy $f(x;\pi) > 0$ and the posterior distributions $\Pi(\theta;\pi,x)$ are well defined for all signals x.

We will write the posterior distribution as $\Pi(\pi, x)$ when we want to consider the posterior as a function of the prior π and observed signal x. Similarly, we will write the predictive distribution for signals as $f(\pi)$.

Let $u(w)$ denote the DM's utility function for wealth w. The DM's value function (or derived utility function) with k periods remaining and prior π, $U_k(w, \pi)$, can be written recursively as

$$U_0(w,\pi) = u(w),$$

$$U_k(w,\pi) = \max \begin{cases} \mathbb{E}[u(w+\tilde{\theta}) \mid \pi] & \text{(adopt)}, \\ u(w) & \text{(reject)}, \\ \mathbb{E}[U_{k-1}(w-c, \Pi(\pi,\tilde{x})) \mid f(\pi)] & \text{(wait)}. \end{cases} \quad (1)$$

Here, when waiting, the DM observes a random signal $\tilde{x}$, updates her prior π to posterior $\Pi(\pi, \tilde{x})$ after seeing the signal, and continues to the next period with her wealth reduced by the cost c of gathering information. This formulation assumes there is no discounting of costs or benefits; we will discuss an extension to a model with discounting in Section 9. The expectations in (1) can be written more explicitly as

$$\mathbb{E}[u(w+\tilde{\theta}) \mid \pi] = \int_\Theta u(w+\theta)\pi(\theta)\,d\theta \quad \text{and}$$
$$\mathbb{E}[U_{k-1}(w, \Pi(\pi,\tilde{x})) \mid f(\pi)] = \int_X U_{k-1}(w, \Pi(\pi,x)) f(x;\pi)\,dx.$$

When writing expectations, we will place tildes on the random variable involved and condition the expectation on the distribution assumed for the random variable. To ensure these expectations are well defined, we will assume that $u(w+\theta)$ is π-integrable for all w and π encountered.

The DM is *risk neutral* if her utility function $u(w)$ is linear, *risk averse* if $u(w)$ is concave, and *risk seeking* if $u(w)$ is convex. We let $\rho_u(w) = -u''(w)/u'(w)$ denote the *coefficient of (absolute) risk aversion* for utility function u and let $\tau_u(w) = 1/\rho_u(w) = -u'(w)/u''(w)$

be the *risk tolerance*. A utility function $u(w)$ exhibits *decreasing absolute risk aversion* (is DARA) if it is strictly increasing; concave (i.e., risk averse); twice differentiable; and $\rho_u(w)$ is decreasing or, equivalently, $\tau_u(w)$ is increasing. A utility exhibits *constant absolute risk aversion* (is CARA) if $\tau_u(w)$ is a constant; a CARA utility is either linear or exponential. A utility function u_2 is *more risk tolerant* than utility function u_1 if $\tau_{u_2}(w) \geqslant \tau_{u_1}(w)$ for all w and is *more risk averse* if the reverse inequality holds.

2.2. An Illustrative Example

We will illustrate our results using a specific numerical example. In this example, we assume that the value of the technology is $\theta = p - 0.5$ where the DM starts with a beta distribution on the parameter p; that is, $\pi(p) \propto p^{(\alpha-1)}(1-p)^{(\beta-1)}$. We assume the cost c of information is 0.01. Signals follow a Bernoulli process with positive (+) signals occurring with probability p and negative (−) signals with probability $1-p$. Applying Bayes' rule, if we start in one period with a beta prior with parameters (α, β), we enter the next period with parameters $(\alpha+1, \beta)$ if a positive signal is observed and with $(\alpha, \beta+1)$ if a negative signal is observed. Thus, the precision $(\alpha+\beta)$ increases by one in each period. The mean of the beta distribution is $\alpha/(\alpha+\beta)$ and the expected benefit of the technology is therefore $\alpha/(\alpha+\beta) - 0.5$. In this setting, rather than working with full distributions π as a state variable, we can instead work with (α, β) as the state variable or, equivalently, with the expected benefit of the technology and the precision of the distribution as state variables.

We first consider the case where the DM is risk neutral with $u(w) = w$; we will assume the initial wealth is equal to \$1.06. (This initial wealth has no impact on the optimal policy given risk neutrality.) We will consider the infinite horizon limit by taking the number of periods remaining to be very large. The optimal policy for this example is shown in Figure 1(a). Here the x-axis corresponds to the precision $(\alpha+\beta)$ and the y-axis represents the expected benefit $(\alpha/(\alpha+\beta) - 0.5)$. The DM's beliefs can be represented by a point in this figure and will move from left to right as the DM gathers information. One such path is shown as the jagged line in the figure: this DM starts with $\alpha = 2.25$ and $\beta = 1.75$ (thus precision $\alpha + \beta = 4$) and observes a signal sequence $(-,+,-,+,-,-)$ before rejecting the technology. In Figure 1(a), we see that the adoption (rejection) thresholds decrease (increase) as the precision increases, converging toward $\theta = 0$.

Figure 1(b) shows the values associated with adopting, waiting, and rejecting, as a function of the expected value of the technology given precision $\alpha+\beta = 10$; this corresponds to a vertical slice in Figure 1(a). The DM starting with precision 4 would reach this level of precision after six periods of waiting and would then have wealth 1.0. In this figure, consistent with Figure 1(a), if the expected benefit is less than approximately −0.03, it is optimal to reject the technology. If the expected benefit is more than 0.03, it is optimal to adopt the technology. Between these two levels, it is optimal to wait and gather additional information. The optimal value function (1) is the upper envelope of the three functions shown in the figure. All of these functions are increasing and convex in the expected benefit.

Figure 2 shows analogous results for the case with risk aversion. Here we assume the DM starts with

Figure 1. (Color online) Results for the Illustrative Example with Risk Neutrality

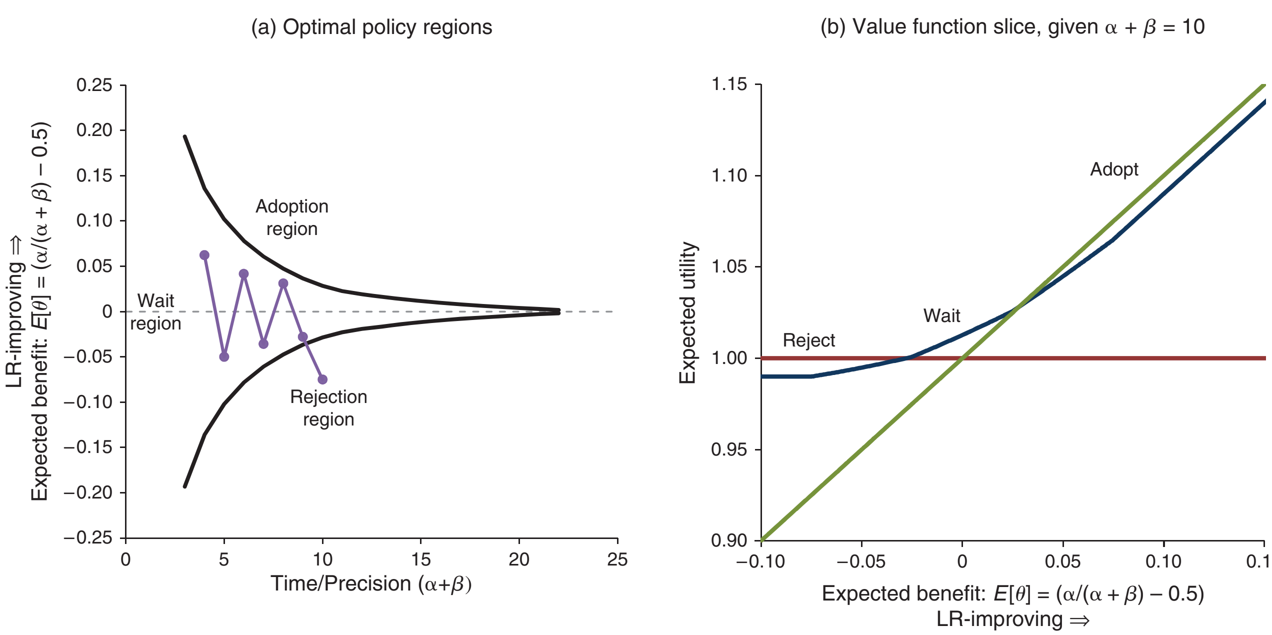

Figure 2. (Color online) Results for the Illustrative Example with Risk Aversion

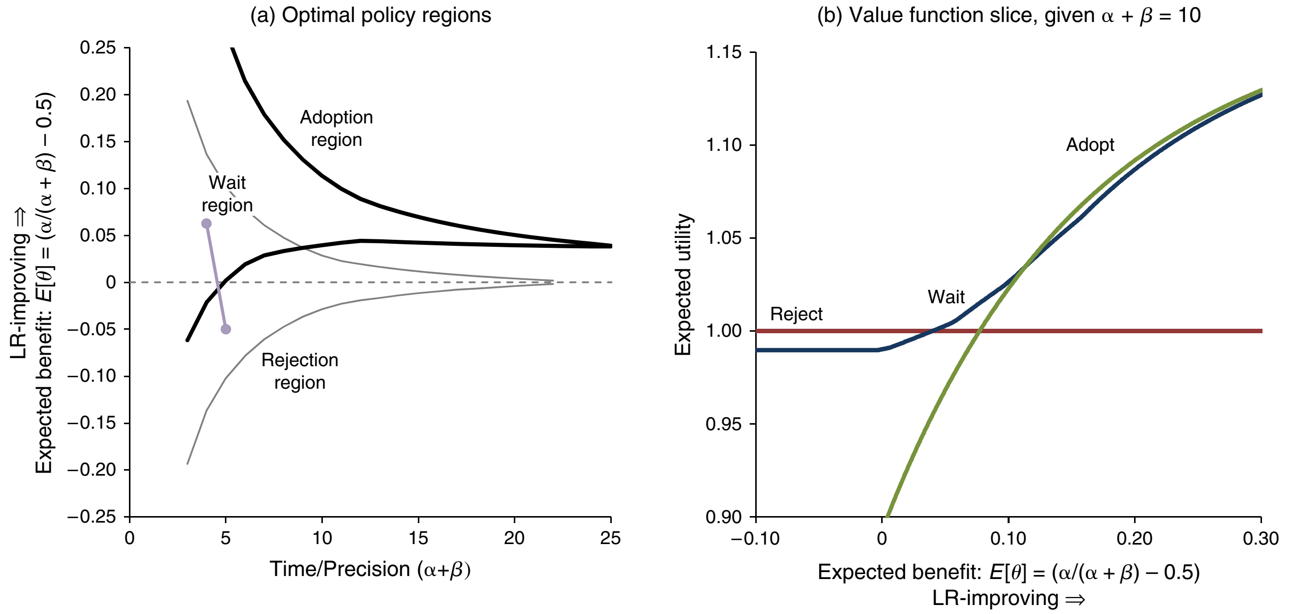

initial wealth \$1.06 (as before) and has a power utility function $u(w) = 1.2 - 0.2w^{(1-\gamma)}$ where $\gamma = 6$ is the coefficient of relative risk aversion, i.e., the DM's risk tolerance is one-sixth of her wealth, $\tau_u(w) = w/6$. (The utility function is scaled so its value and slope match that of the risk-neutral utility function $u(w) = w$ at $w = 1.0$.) Examining the optimal policy regions in Figure 2(a), we see that the risk-averse DM is more conservative in adopting (she requires a much higher expected benefit before adopting) and rejects in many scenarios where a risk-neutral DM would wait or adopt. In particular, on the example path in Figure 1(a) where the DM starts with $\alpha = 2.25$ and $\beta = 1.75$, the risk-averse DM would reject the technology after observing a single negative signal.

Figure 2(b) shows the expected utilities associated with adopting, waiting, and rejecting in the same setting as Figure 1(b). As in Figure 1(b), it is optimal to reject for low expected benefits, optimal to adopt for high expected benefits, and optimal to wait between these two levels. Again, all of these functions are increasing in the expected benefit. The value given adoption is concave in the expected benefit, reflecting the DM's risk aversion. The optimal value function (the upper envelope) and the value with waiting are convex in the lower range and concave in the higher range, where adoption is optimal or likely in future periods.

3. Preliminaries

The example of the previous section illustrates the kinds of results that one might hope to establish as general properties of the technology adoption model and as general effects of risk aversion. To study monotonicity properties of the model, we will rely heavily on the likelihood-ratio order on distributions and the monotone likelihood-ratio property for the signal's likelihood functions. In this section, we briefly review these definitions and some important related properties.

Definition 3.1 (LR-Dominance and Related Definitions). (i) π_2 likelihood-ratio (LR) dominates π_1 ($\pi_2 \succeq_{LR} \pi_1$) if $\pi_2(\theta_2)\pi_1(\theta_1) \geqslant \pi_2(\theta_1)\pi_1(\theta_2)$ for all $\theta_2 \geqslant \theta_1$.

(ii) A signal process has the monotone-likelihood-ratio (MLR) property if the signal space X is totally ordered and $L(x \mid \theta_2) \succeq_{LR} L(x \mid \theta_1)$ for all $\theta_2 \geqslant \theta_1$.

(iii) A function $V(\pi)$ defined on distributions on Θ is *LR-increasing* if $V(\pi_2) \geqslant V(\pi_1)$ whenever $\pi_2 \succeq_{LR} \pi_1$; $V(\pi)$ is *LR-decreasing* if $-V(\pi)$ is LR-increasing.

Note that if $\pi_1(\theta) > 0$ for all θ, the condition defining LR-dominance is equivalent to $\pi_2(\theta)/\pi_1(\theta)$ being increasing in θ. In the illustrative example of Section 2.2, increasing α in the beta prior while holding the precision $\alpha + \beta$ constant leads to a LR-improvement in the distribution. Thus moving upward in Figures 1(a) or 2(a) or to the right in Figure 1(b) or 2(b) corresponds to a LR-improvement in the prior. If we assume that positive signals are greater than negative signals, the Bernoulli signal process satisfies the MLR property.

The key properties of the LR-order that we will use are (i) LR-dominance implies first-order stochastic dominance; (ii) LR-dominance survives Bayesian updating; and (iii) if the signal process satisfies the MLR property, the LR-order leads to natural monotonicity properties in signals. These properties are well known.

Proposition 3.1 (Properties of the LR-Order). (i) *If* $\pi_2 \succeq_{LR} \pi_1$, *then* $\mathbb{E}[v(\tilde{\theta}) \mid \pi_2] \geqslant \mathbb{E}[v(\tilde{\theta}) \mid \pi_1]$ *for any increasing function* $v(\theta)$.

(ii) *Given any signal x, the posteriors are LR-ordered if and only if the priors are LR-ordered*: *that is*, $\pi_2 \succeq_{LR} \pi_1 \Leftrightarrow \Pi(\pi_2, x) \succeq_{LR} \Pi(\pi_1, x)$, for all x.

(iii) *If the signal process satisfies the MLR property, then*

(a) $\pi_2 \succeq_{LR} \pi_1 \Rightarrow f(\pi_2) \succeq_{LR} f(\pi_1)$

(b) *For any prior* π, $x_2 \geqslant x_1 \Leftrightarrow \Pi(\pi, x_2) \succeq_{LR} \Pi(\pi, x_1)$.

These properties of the LR-order imply that expectations of increasing functions generate LR-increasing functions and the LR-increasing property is preserved by Bayesian updating. These results are critical for establishing monotonicity properties in our model and will be generalized to s-increasing functions in Section 6.

Proposition 3.2 (Properties of LR-Increasing Functions). (i) *If $v(\theta)$ is increasing in θ, then $\mathbb{E}[v(\tilde{\theta}) \mid \pi]$ is LR-increasing.*

(ii) *Suppose the signal process satisfies the MLR property and $V(\pi)$ is a LR-increasing function; then $\mathbb{E}[V(\Pi(\pi, \tilde{x})) \mid f(\pi)]$ is LR-increasing.*

Here part (i) follows from Proposition 3.1(i) and the definition of LR-increasing. Part (ii) uses each element of Proposition 3.1; see US (2009), Lemma 3.5 for a proof.

As discussed in the introduction, we will consider single-crossing properties of policy differences and define single-crossing functions as follows.

Definition 3.2 (Single-Crossing Functions). (i) $v(\theta)$ is single crossing if, for all $\theta_2 \geqslant \theta_1$, $v(\theta_1) \geqslant 0$ implies $v(\theta_2) \geqslant 0$ and $v(\theta_1) > 0$ implies $v(\theta_2) > 0$.

(ii) $V(\pi)$ is LR-single crossing if, for all $\pi_2 \succeq_{LR} \pi_1$, $V(\pi_1) \geqslant 0$ implies $V(\pi_2) \geqslant 0$ and $V(\pi_1) > 0$ implies $V(\pi_2) > 0$.

We note that single-crossing functions $v(\theta)$ generate LR-single-crossing functions $V(\pi)$ in much the same way that increasing functions generate LR-increasing functions in Proposition 3.2(i).

Proposition 3.3 (Generating Single-Crossing Functions). *If $v(\theta)$ is single crossing, then $\mathbb{E}[v(\tilde{\theta}) \mid \pi]$ is LR-single crossing.*

Proof. See, e.g., Karlin (1968, Chapter 1, Theorem 3.1, p. 21). □

4. Straightforward Properties

A number of properties of the risk-neutral model generalize in a straightforward manner to the risk-sensitive model. In this section, we focus on the monotonicity and convexity results that we will use later and briefly mention some other properties.

4.1. Monotonicity of the Value Function

We first show that if the signal process satisfies the MLR property and the DM's utility function is increasing, the optimal value function is monotonic in that LR-improvements in the prior distribution π lead to higher values. The proof is analogous to the proof for the risk-neutral case (see Proposition 3.6 in US 2009).

Proposition 4.1 (Monotonicity of the Value Function). *Suppose the DM's utility function $u(w)$ is increasing in w and the signal process satisfies the MLR property. Then, for all k and w, the value function $U_k(w, \pi)$ is LR-increasing in π.*

Proof. We show this by induction. The terminal value function, $U_0(w, \pi) = u(w)$, is independent of π and thus trivially LR-increasing for all w. The value if the DM adopts ($\mathbb{E}[u(w + \tilde{\theta}) \mid \pi]$) is LR-increasing by Proposition 3.2(i) since $u(w + \theta)$ is increasing in θ. The value if the DM rejects, $u(w)$, is independent of π and trivially LR-increasing. Now suppose $U_{k-1}(w, \pi)$ is LR-increasing for all w. By Proposition 3.2(ii), the value if the DM waits, $\mathbb{E}[U_{k-1}(w - c, \Pi(\pi, \tilde{x})) \mid f(\pi)]$, is LR-increasing. Thus, $U_k(w, \pi)$, as the maximum of three LR-increasing functions, is also LR-increasing. □

As discussed following Definition 3.1, the signal process in the illustrative example satisfies the MLR property and movements to the right in Figures 1(b) and 2(b) correspond to LR-improvements in the prior. In these figures, it is evident—both with risk neutrality and with risk aversion—that the values corresponding to adopting, rejecting, and waiting are all increasing with such improvements, so the optimal value function (the upper envelope of these functions) is also increasing in this direction.

4.2. Convexity

Convexity in the priors π follows exactly as in the case with risk neutrality. Here we do not need to place any assumptions on the utility function (e.g., the utility function need not be increasing or concave) or on the signal process (e.g., the likelihood function need not satisfy the MLR property).

Proposition 4.2 (Convexity). (i) *For all k and w, the value function $U_k(w, \pi)$ is convex in π.*

(ii) *If it is optimal to adopt (reject) with priors π_1 and π_2, then it is also optimal to adopt (reject) with prior $\alpha\pi_1 + (1 - \alpha)\pi_2$ for any α such that $0 \leqslant \alpha \leqslant 1$.*

Proof. See US (2009) Propositions 6.2 and 6.3. □

This convexity result follows because the expected utility associated with adoption is linear in the probabilities and can be interpreted as an aversion toward uncertainty about the prior π. The convex combination $\pi_\alpha = \alpha\pi_1 + (1 - \alpha)\pi_2$ can be interpreted as there being probability α of π_1 prevailing and probability $(1 - \alpha)$ of π_2 prevailing. A convex value function means that the DM would prefer to resolve this uncertainty before beginning the information-gathering process (for an expected utility of $\alpha U_k(w,\pi_1)+(1-\alpha)U_k(w,\pi_2))$ rather than begin the information-gathering process with this uncertainty unresolved (for an expected utility of $U_k(w,\pi_\alpha)$). Intuitively, if the DM knew whether π_1 or π_2 prevailed, she could make better adoption, rejection, and/or information gathering decisions.[2]

4.3. Other Straightforward Extensions

In addition to these monotonicity and convexity results, a number of other results generalize in a straightforward manner from the risk-neutral to risk-averse model. We state a few of these results informally:

- Increasing the informativeness of the signal processes (in the sense of Blackwell 1951) makes the DM better off and encourages information gathering.
- Cheaper information (reducing c) makes the DM better off, encourages information gathering, and delays adoption. If $c=0$, it is always optimal to gather additional information until period $T-1$; the DM will then decide whether to adopt or reject in period $T-1$. In this case, information is free and potentially valuable and, with no discounting, delaying adoption is not harmful.
- Increasing the number of periods remaining (k) makes the DM better off and encourages information gathering. However, there are diminishing returns to increasing the number of periods remaining:
  - —the expected utilities converge for all priors, and
  - —if the utility function is strictly increasing, the DM will almost certainly stop gathering information at some point.

Here "makes the DM better off" means the value function (weakly) increases with the given change in assumptions and "encourages information gathering" means that if it is optimal to gather more information with the initial assumption, then it remains optimal to gather information with the change. These results are discussed in US (2009) and risk aversion does not play a significant role in the proofs or interpretation.

5. Monotonicity of the Policies

We next consider how the optimal policy responds to changes in the prior on benefits. As discussed in the introduction, with risk neutrality, if the signal process satisfies the MLR property, we get monotonic policies as well as monotonic value functions: along any chain of LR-improving distributions, the optimal action moves from rejection toward adoption, perhaps passing through the information gathering region. In the illustrative example, moving up in Figure 1(a) or 2(a) corresponds to a LR-improvement and we can see that, in both the risk-neutral and risk-averse cases, the optimal policies are monotonic in this sense. The monotonicity of the rejection policy is easy to establish in both the risk-neutral and risk-averse models. The monotonicity of the adoption policy is more difficult to establish, and, as we demonstrate in Section 5.3 below, this monotonicity need not hold with risk aversion.

5.1. Rejection Policies

Monotonicity of rejection policy follows under the same conditions required to ensure that value functions are LR-increasing.

Proposition 5.1 (Monotonicity of Rejection Policies). *Suppose the DM's utility function is increasing and the signal process satisfies the MLR property. If it is optimal to reject with prior π_2, it is also optimal to reject with any prior π_1 such that $\pi_2 \succeq_{LR} \pi_1$.*

Proof. Consider the difference between the optimal value and the value given by rejecting, $F_k(w,\pi)=U_k(w,\pi)-u(w)$. If it is optimal to reject given prior π_2, we have $U_k(w,\pi_2)=u(w)$ and thus $F_k(w,\pi_2)=0$. If we assume that the DM's utility function is increasing and the signal process satisfies the MLR property, $U_k(w,\pi)$ is LR-increasing (by Proposition 4.1) and thus $U_k(w,\pi_1)\leqslant U_k(w,\pi_2)=u(w)$. However, since $U_k(w,\pi_1)$ corresponds to the optimal action and rejecting is a possible action, we know $U_k(w,\pi_1)\geqslant u(w)$. Therefore, $U_k(w,\pi_1)=u(w)$ and it is optimal to reject given π_1 as well. Thus the monotonicity of rejection policy follows under the same conditions required to ensure that value functions are LR-increasing. □

5.2. Adoption Policies

The monotonicity of the adoption policy is more difficult to establish because the optimal value and the value from adopting both potentially change with π. Let $G_k(w,\pi)=\mathbb{E}[u(w+\tilde{\theta})\mid\pi]-U_k(w,\pi)$ be the difference between the value associated with immediate adoption and the optimal value function; note that $G_k(w,\pi)\leqslant 0$. We can write $G_k(w,\pi)$ recursively as

$$G_0(w,\pi)=\mathbb{E}[u(w+\tilde{\theta})-u(w)\mid\pi],$$

$$G_k(w,\pi)=\min\begin{cases}0 \quad \text{(adopt)},\\ \mathbb{E}[u(w+\tilde{\theta})-u(w)\mid\pi] \quad \text{(reject)},\\ \mathbb{E}[u(w+\tilde{\theta})-u(w+\tilde{\theta}-c)\mid\pi]\\ \quad +\mathbb{E}[G_{k-1}(w-c,\Pi(\pi,\tilde{x}))\mid f(\pi)] \quad \text{(wait)}.\end{cases} \tag{2}$$

If we assume that information is free or that the DM is risk neutral or risk seeking, we can show the utility difference $G_k(w,\pi)$ is LR-increasing using a straightforward induction argument. The terminal utility difference, $G_0(w,\pi)$, is LR-increasing if the utility function is increasing (this follows from Proposition 3.2(i)). The terms associated with adopting and rejecting in (2) are also LR-increasing. For the induction hypothesis, suppose $G_{k-1}(w,\pi)$ is LR-increasing. Proposition 3.2(ii) implies that the expected continuation value, $\mathbb{E}[G_{k-1}(w,\Pi(\pi,\tilde{x}))\mid f(\pi)]$, is LR-increasing.

Now consider the "reward" associated with waiting in (2),

$$\mathbb{E}[u(w+\tilde{\theta}) - u(w+\tilde{\theta}-c) \mid \pi]. \tag{3}$$

If information is free (i.e., $c = 0$), this reward term reduces to 0 and is trivially LR-increasing. In the risk-neutral case with $u(w) = w$, the reward term (3) reduces to a constant c and is also trivially LR-increasing. If the DM is risk seeking (i.e., the utility function is convex), the utility difference

$$u(w+\theta) - u(w+\theta-c), \tag{4}$$

is increasing in θ and the reward term (3) is LR-increasing (again, by Proposition 3.2(i)). The sum of two LR-increasing functions is LR-increasing, so adding the reward term (3) and the continuation value, we know that the term associated with waiting in (2) is LR-increasing. Thus, the functions associated with each choice in (2) are all LR-increasing, and then $G_k(\pi)$, as the minimum of three LR-increasing functions, is LR-increasing. This leads to the following proposition.

Proposition 5.2 (Monotonicity of Adoption Policies in Special Cases). *Suppose that the signal process satisfies the MLR property, the utility function is increasing and one of the following holds*:

(i) *information is free* ($c = 0$),
(ii) *the DM is risk neutral, or*
(iii) *the DM is risk seeking.*

If it is optimal to adopt with prior π_1, *it is also optimal to adopt with any prior* π_2 *such that* $\pi_2 \succeq_{LR} \pi_1$.

Proof. Because $G_k(\pi)$ is LR-increasing (as argued before the proposition) and $G_k(\pi) \leqslant 0$, if it is optimal to adopt for π_1 (i.e., $G_k(\pi_1) = 0$), then it is also optimal to adopt for π_2 (i.e., $G_k(\pi_2) = 0$) if $\pi_2 \succeq_{LR} \pi_1$. □

If the DM is risk averse (i.e., $u(w)$ is concave) and $c > 0$, the argument underlying the result of Proposition 5.2 breaks down as the utility difference (4) is now decreasing rather than increasing and consequently the reward term (3) is LR-decreasing rather than LR-increasing. We can see the difficulty in the risk-averse case by comparing Figures 1(b) and 2(b). In the risk-neutral case, the difference between adoption and waiting is increasing when moving from left to right in Figure 1(b). In the risk-averse case in Figure 2(b), this difference is initially increasing but then decreases.

In the next subsection, we provide an example that demonstrates that the adoption policy need not be monotonic, even with a reasonable utility function. However, we will show in Section 7 that we will have monotonic policies if we assume the utility function is not "too risk averse." We first note that in the special case where there are only two possible technology values, policies will be monotonic without any special utility assumptions (beyond assuming that $u(w)$ is increasing) or assuming the signal process satisfies the MLR property. In this case, we appeal to the convexity result of Proposition 4.2 rather than increasing difference arguments to establish the monotonicity of the optimal policies.

Proposition 5.3 (Monotonicity of Policies with Two Outcomes). *Suppose the DM's utility function is increasing and there are two possible technology values. If it is optimal to adopt with prior* π_1, *then it is also optimal to adopt with any prior* π_2 *such that* $\pi_2 \succeq_{LR} \pi_1$. *Similarly, if it is optimal to reject with prior* π_2, *then it is also optimal to reject with any prior* π_1 *such that* $\pi_2 \succeq_{LR} \pi_1$.

Proof. Let θ_l and θ_h denote the two possible values of θ and p be the probability associated with θ_h and $(1-p)$ with θ_l. If $0 \leqslant \theta_l, \theta_h$ (or $0 \geqslant \theta_l, \theta_h$), given that the utility function is increasing, it would clearly be optimal to adopt (or reject) for any p. Now assume $\theta_l < 0 < \theta_h$. With two outcomes, increasing p is an LR-improvement and all possible priors are convex combinations of the priors with $p = 0$ and $p = 1$. If it is optimal to adopt with probability p_1, because it is optimal to adopt with probability $p = 1$, convexity of the adoption region (Proposition 4.2) implies that it is optimal to adopt with any probability p_2, such that $p_1 \leqslant p_2 \leqslant 1$. A similar argument holds for the rejection result. □

One implication of this result is that any example demonstrating the nonmonotonicity of the adoption policy must involve at least three different technology values.

5.3. Nonmonotonic Adoption Policies

We show that adoption policies need not be monotonic by considering the two-period example illustrated in the decision trees of Figure 3. Here the DM is assumed to have a logarithmic utility function with initial wealth \$23.002. There are three possible technology values (θ): low (−\$20.00), medium (−\$10.55), or high (+\$25.00). The DM can adopt or reject immediately or wait and pay \$3 to observe a signal that may be positive or negative. We consider two different priors π_1 and π_2 as shown in Table 1, corresponding to Figures 3(a) and 3(b), respectively. The likelihoods for the signals $L(x \mid \theta)$ are the same in both cases and are shown in Table 1. Here π_2 LR-dominates π_1 and the signal process satisfies the MLR property. The corresponding predictive distributions for the signals and posteriors are shown in Figures 3(a) and (b).

The expected utilities are shown below and left of the decision and chance nodes in the decision trees of Figure 3; the optimal choices are indicated with bold lines. In Figure 3(a) with prior π_1, it is optimal for the DM to adopt immediately. In Figure 3(b) with prior π_2, it is optimal for the DM to wait and decide after observing a signal, adopting if the signal is positive and rejecting

Figure 3. An Example with a Nonmonotonic Optimal Policy

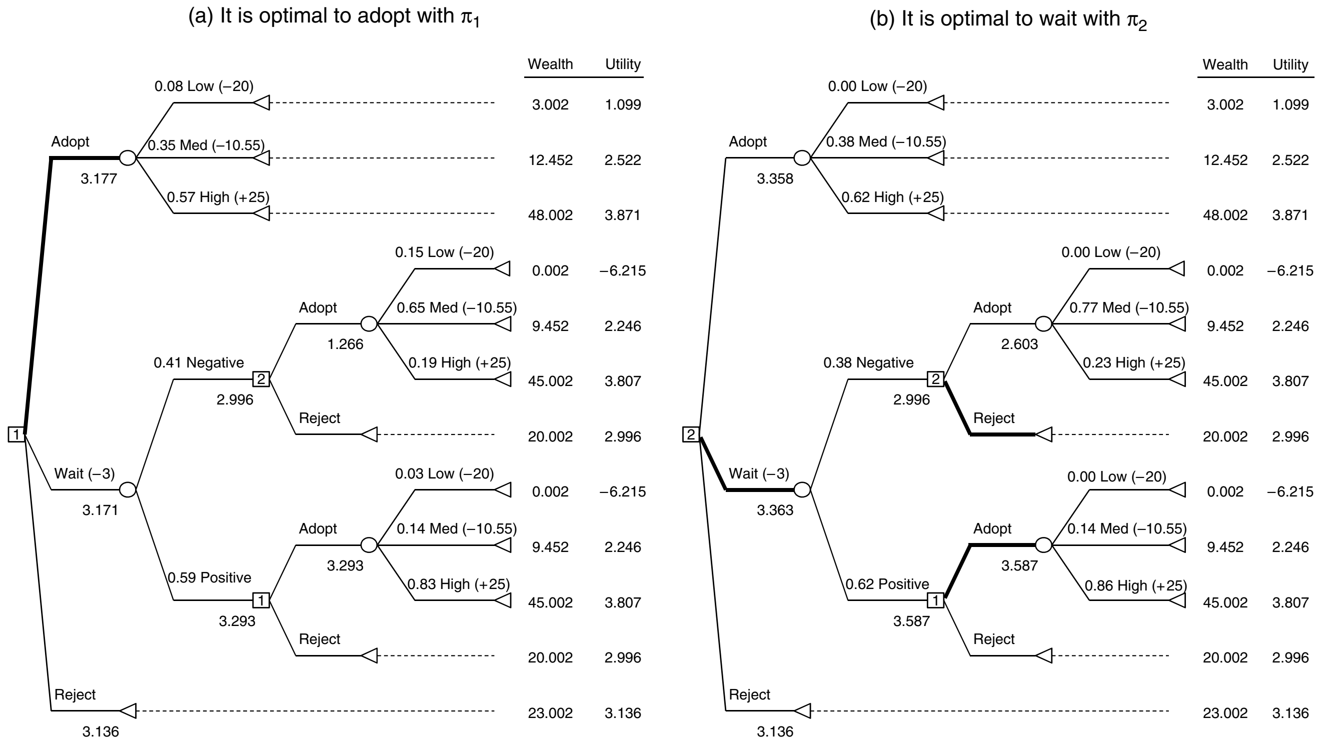

otherwise. Thus, the optimal policy for this example is not monotonic in the prior: a LR-improvement in the prior—changing from π_1 to π_2—leads the DM to switch from adopting to waiting.

It is not difficult to understand what is happening in this example. With prior π_1, there is a 0.08 probability of having the low outcome with a technology value of −\$20. With an initial wealth of \$23.002, this outcome would leave the DM with a wealth of \$3.002. If the DM waits and adopts after a positive signal, there is a reduced chance (≈ 0.03) of this low outcome. However, taking into account the costs of waiting (\$3), in this case the DM is left with a wealth of only \$0.002, which is a near ruinous outcome with a logarithmic utility function, with a large negative utility. However, with prior π_2, there is no chance of having the low outcome and thus there is no chance of having the near ruinous outcome if the DM waits; in this case, it is optimal to wait. Here, as we reduce (or eliminate) the probability of the disastrous outcome, we need the intermediate, but still negative, third outcome to make waiting attractive.

Table 1. Data for the Example of Figure 1

		Priors		Likelihood $L(x \mid \theta)$	
Scenario	Benefit θ	π_1	π_2	Negative	Positive
Low	−20.00	0.08	0.00	0.78	0.22
Medium	−10.55	0.35	0.38	0.77	0.23
High	+25.00	0.57	0.62	0.14	0.86

The essential features of this example are positive information gathering costs c and a utility function that approaches $-\infty$ as wealth approaches zero. Given any specific nonmonotonicity—that is, priors π_1 and π_2 such that $\pi_2 \succeq_{LR} \pi_1$ and it is optimal to adopt with π_1 and optimal to gather information with π_2—if we reduce the cost c (holding all else constant), for sufficiently small positive c, the nonmonotonicity will disappear, as it will become optimal to gather information with π_1 as well as π_2. As noted in Proposition 5.2, adoption policies are monotonic if $c = 0$; thus $c > 0$ is necessary to have a nonmonotonicity. Given any $c > 0$, if we have a utility function that approaches $-\infty$ as wealth approaches zero—such as the log utility in the example above or the power utility—we can construct a nonmonotonicity. Specifically, given a π_2 (and likelihood function) such that waiting is strictly preferred to adopting and adopting is strictly preferred to quitting, we can define a prior π_1 such that $\pi_2 \succeq_{LR} \pi_1$ by adding mass at a $\underline{\theta}$ such that $w + \underline{\theta} - c = 0$, taking the mass to be small enough so adopting is still preferred to quitting. The expected utility associated with adopting after waiting given $\underline{\theta}$ is $u(w + \underline{\theta} - c) = -\infty$ and waiting is unattractive; adopting is then optimal with π_1. (We formalize this construction in Online Appendix B.1.) Thus, given $c > 0$, a utility function that approaches $-\infty$ as wealth approaches zero is sufficient for creating examples with nonmonotonic policies.

6. s-Increasing and sLR-Increasing Functions

As discussed in Section 5 the policy differences $G_k(w, \pi)$ in Equation (2) are not necessarily LR-increasing with risk aversion and this can lead to nonmonotonic policies. To establish monotonicity of the policies, we will provide conditions that ensure that these differences will cross at most once as we LR-improve the prior distribution. Unfortunately, single crossing is a difficult property to work with in DP models. Whereas sums and weighted sums of increasing functions will be increasing, sums and weighted sums of single crossing functions need not be single crossing. Thus, in a DP, even if the reward function and continuation value are both single crossing, their sum (or expectation) need not be. Instead, we define new properties that are weaker than increasing and LR-increasing but still imply the desired single-crossing properties.

Definition 6.1 (s-Increasing Functions). Let $s(\theta)$ be a *scaling function* such that $s(\theta) \geqslant 0$ for all θ and $s(\theta) > 0$ for some θ and let $S(\pi) = \mathbb{E}[s(\tilde{\theta}) \mid \pi]$.

(i) *$v(\theta)$ is s-increasing if $v(\theta_2)s(\theta_1) \geqslant v(\theta_1)s(\theta_2)$ for all $\theta_2 \geqslant \theta_1$.*

(ii) *$V(\pi)$ is sLR-increasing if $V(\pi_2)S(\pi_1) \geqslant V(\pi_1)S(\pi_2)$ for all $\pi_2 \succeq_{LR} \pi_1$.*

Note that if $s(\theta) > 0$ for all θ, s-increasing is equivalent to $v(\theta)/s(\theta)$ being increasing in θ and sLR-increasing is equivalent to $V(\pi)/S(\pi)$ being LR-increasing. If the scaling function $s(\theta)$ is a positive constant, s-increasing is equivalent to the ordinary sense of increasing. However, with nonconstant scaling functions, s-increasing functions need not be increasing. Similarly, constant functions need not be s-increasing.

We can think of s-increasing as being "sort of" increasing, as the properties of increasing functions that we used to show the policy differences were increasing hold for s-increasing functions as well.

Proposition 6.1 (Properties of s-Increasing Functions). *Given a scaling function $s(\theta)$ satisfying the conditions of Definition 6.1,*

(i) *For any scalar α, $\alpha s(\theta)$ and $-\alpha s(\theta)$ are both s-increasing.*

(ii) *If $v_1(\theta)$ and $v_2(\theta)$ are s-increasing, then for any $\alpha, \beta \geqslant 0$, $\alpha v_1(\theta) + \beta v_2(\theta)$ is also s-increasing.*

(iii) *The maximum or minimum of two or more s-increasing functions is also s-increasing.*

(iv) *Pointwise limits of s-increasing functions are also s-increasing.*

(v) *Suppose $s(\theta)$ is single crossing and $v(\theta)$ is s-increasing.*

(a) *If $s(\theta) = 0$, then $v(\theta) \leqslant 0$.*

(b) *If $s(\theta_1) > 0$ and $\theta_2 \geqslant \theta_1$, then $v(\theta_1) \geqslant 0$ implies $v(\theta_2) \geqslant 0$ and $v(\theta_1) > 0$ implies $v(\theta_2) > 0$.*

These same properties hold for sLR-increasing functions with $V(\pi)$ and $S(\pi)$ replacing $v(\theta)$ and $s(\theta)$ and $\pi_2 \succeq_{LR} \pi_1$ replacing $\theta_2 \geqslant \theta_1$.

These results are straightforward to prove. Part (i) of the proposition says that functions $\alpha s(\theta)$ can be interpreted as "s-constants." Recall from Section 5.2 that the utility differences (4) are constant in the risk-neutral case but decreasing in the risk-averse case, when we wanted it to be constant or increasing. We will choose a scaling function $s(\theta)$ so that these utility differences are s-constant or s-increasing. We then use parts (ii) and (iii) above in a recursive proof to show that the policy differences are sLR-increasing. Parts (ii) and (iv) of Proposition 6.1 imply that s-increasing is a closed convex cone (C3) property in the sense of Smith and McCardle (2002); such properties arise frequently and quite naturally in stochastic DPs.

Proposition 6.1(v) is a single-crossing result. If $s(\theta) > 0$ for all θ (and thus $s(\theta)$ is single crossing), then this result reduces to the assertion that an s-increasing function is single crossing, in the same way that an increasing function is single crossing. Allowing $s(\theta)$ to be zero for some θ can provide some additional flexibility. In regions where $s(\theta) = 0$, an s-increasing function $v(\theta)$ must be nonpositive but no other conditions are placed on $v(\theta)$. However, if $s(\theta)$ is single crossing and $s(\theta_0) > 0$, then the single-crossing condition is "switched on" at θ_0 and $v(\theta)$ must be single crossing for $\theta > \theta_0$.

We now show that the analog of Proposition 3.2 holds for s-increasing functions; that is, expectations of s-increasing functions generate sLR-increasing functions and the sLR-increasing property is preserved by Bayesian updating. These results (and the properties of Proposition 6.1) are the key results for establishing monotonic policies in the technology adoption model with risk aversion.

Proposition 6.2 (Properties of sLR-Increasing Functions). *Let $s(\theta)$ be a scaling function as in Definition 6.1.*

(i) *If $v(\theta)$ is s-increasing, then $\mathbb{E}[v(\tilde{\theta}) \mid \pi]$ is sLR-increasing.*

(ii) *Suppose the signal process satisfies the MLR property, $V(\pi)$ is a sLR-increasing function and s is single-crossing; then $\mathbb{E}[V(\Pi(\pi, \tilde{x})) \mid f(\pi)]$ is sLR-increasing in π.*

Proof. See Appendix A.1. □

We also note that the result that Bayesian updating preserves the sLR-increasing property (part (ii) above) generalizes to the setting where θ is changing over time (as in a partially observable Markov decision process), provided these transitions satisfy the MLR property. We discuss this in more detail in Appendix B.2.

We conclude this section with a brief discussion of properties that are related to s-increasing. First, the definition of s-increasing functions is related to log-supermodularity (see, e.g., Athey 2002), which is also referred to as "totally positive of order 2" (see, e.g., Karlin 1968). A nonnegative function $f(x)$ is log-supermodular (or totally positive of order 2) if $\log f(x)$ is supermodular. In the case where $x = (i, \theta)$ with $i \in \{1, 2\}$ and $\theta_2 \geqslant \theta_1$, this is equivalent to requiring

$$f(1, \theta_2) f(2, \theta_1) \leqslant f(1, \theta_1) f(2, \theta_2).$$

If we take $f(1, \theta) = s(\theta)$ and $f(2, \theta) = v(\theta)$, this reduces to the condition defining s-increasing. However, s-increasing, unlike log-supermodularity, does not require $v(\theta)$ to be nonnegative. This is important in our application in the technology adoption model because the functions involved may be positive or negative.

A second related property is "signed-ratio monotonicity" (SRM), which was introduced in Quah and Strulovici (2012) and is used to "aggregate the single crossing property." Specifically, Quah and Strulovici (2012) show that the weighted sum of two single-crossing functions will be single crossing if and only if the two functions satisfy the SRM condition. They also provide conditions that ensure that the integral of a family of single-crossing functions will be single crossing. However, their conditions for aggregation are not satisfied in the technology adoption model and we could not apply these results in this setting.

7. Adoption Policies with Risk Aversion

We now return to the problem of showing that the adoption policies in the technology adoption model with risk aversion have a monotonic structure. Specifically, we will show that, given certain utility assumptions and a particular choice of scaling function $s(\theta)$, the difference between the value associated with immediate adoption and the optimal value function, $G_k(w, \pi)$, is sLR-increasing.

Consider the recursive form of the policy difference $G_k(w, \pi)$ given in (2). To show $G_k(w, \pi)$ is sLR-increasing, we will take the scaling function to be utility difference associated with waiting in (2):

$$s(\theta) = u(w_0 + \theta) - u(w_0 + \theta - c). \qquad (5)$$

Here $w_0 = w - kc$ is the DM's wealth if she starts with wealth w with k periods remaining and waits for k periods. If $c > 0$ and the DM's utility function $u(w)$ is strictly increasing and concave (i.e., risk averse), $s(\theta)$ is a positive and decreasing function. With this scaling function, an s-increasing function $v(\theta)$ may be decreasing when $v(\theta)$ is positive but must be increasing when $v(\theta)$ is negative.

To establish the monotonicity result for adoption policies, we will place assumptions on the DM's utility function to ensure the reward functions associated with waiting and rejecting in (2) are both sLR-increasing.

Proposition 7.1 (Utility Conditions). *Suppose $c > 0$ and the DM's utility $u(w)$ exhibits decreasing absolute risk aversion (is DARA); i.e., $\tau_u(w)$ is increasing. Define w_0 and $s(\theta)$ as in (5). For all $\Delta \geqslant 0$,*

(i) *$u(w_0 + \Delta + \theta) - u(w_0 + \Delta + \theta - c)$ is s-increasing.*

(ii) *$u(w_0 + \Delta + \theta) - u(w_0 + \Delta)$ is s-increasing if $\tau_u(w_0 + \underline{\theta} - c) \geqslant -\underline{\theta}$ where $\underline{\theta}$ is the smallest possible value of θ (or less than or equal to all possible θ) .*

If $u(w)$ is CARA, then (i) *and* (ii) *hold and no risk tolerance bound is required in* (ii).

Proof. See Appendix A.2. □

The lower bound on the risk tolerance in the assumption of part (ii)—the assumption that the DM is not "too risk averse" given the lowest possible wealth state—rules out the behavior that leads to nonmonotonic policies in the example of Section 5.3. This condition can be viewed either as limiting the worst possible outcome $\underline{\theta}$ for a given utility function or as constraining the utility function in light of the worst possible outcome.

We can now assemble the pieces and show that, given these utility assumptions, $G_k(w, \pi)$ is sLR-increasing and thus the adoption policies are monotonic. The proof proceeds as in Proposition 5.2 but with s-increasing and sLR-increasing properties replacing increasing and LR-increasing.

Proposition 7.2 (Monotonicity of Adoption Policies with Risk Aversion). *Suppose the assumptions of Proposition 7.1 are satisfied and the signal process satisfies the MLR property. If it is optimal to adopt with prior π_1, it is also optimal to adopt with any prior π_2 such that $\pi_2 \succeq_{LR} \pi_1$.*

Proof. We define the scaling function $s(\theta)$ as in Proposition 7.1 and use an induction argument to show that $G_k(w_k, \pi)$ is sLR-increasing. For the terminal case, by Proposition 7.1(ii) (with $\Delta = 0$), $u(w_0 + \theta) - u(w_0)$ is s-increasing; then by Proposition 6.2(i), $G_0(w_0, \pi)$ is sLR-increasing.

For the inductive step, let $w_k = w_0 + kc$ and assume that $G_{k-1}(w_{k-1}, \pi)$ is sLR-increasing. Let us consider the utility differences for the three different possible actions defining $G_k(w_k, \pi)$:

(i) Adopting yields 0, which is trivially sLR-increasing.

(ii) Rejecting yields $\mathbb{E}[u(w_k + \tilde{\theta}) - u(w_k) \mid \pi]$. The function $u(w_k + \theta) - u(w_k)$ is s-increasing by Proposition 7.1(ii) (with $\Delta = kc$). $\mathbb{E}[u(w_k + \tilde{\theta}) - u(w_k) \mid \pi]$ is then sLR-increasing by Proposition 6.2(i).

(iii) Waiting yields $\mathbb{E}[u(w_k + \tilde{\theta}) - u(w_k + \tilde{\theta} - c) \mid \pi]$ plus a continuation value. The function inside the expectation, $u(w_k + \theta) - u(w_k + \theta - c)$, is s-increasing by Proposition 7.1(i) (with $\Delta = kc$); its expectation $\mathbb{E}[u(w_k + \tilde{\theta}) - u(w_k + \tilde{\theta} - c) \mid \pi]$ is then sLR-increasing by Proposition 6.2(i). The continuation value $\mathbb{E}[G_{k-1}(w_{k-1}, \Pi(\pi, \tilde{x})) \mid f(\pi)]$ is sLR-increasing by the induction hypothesis and Proposition 6.2(ii). By Proposition 6.1(ii), the sum of reward and continuation value (both of which are sLR-increasing) is sLR-increasing. Thus the utility difference associated with waiting is sLR-increasing.

Because the minimum of three sLR-increasing functions is also sLR-increasing (Proposition 6.2(iii)), $G_k(w_k, \pi)$ is sLR-increasing. Note that $S(\pi) > 0$ for all π, because $s(\theta) > 0$; $S(\pi)$ is thus LR-single crossing. $G_k(w_k, \pi)$ is therefore also LR-single crossing (Proposition 6.1(v)). If it is optimal to adopt with prior π_1 (i.e., $G_k(w_k, \pi_1) = 0$), then because $G_k(w, \pi)$ is LR-single crossing, we have $G_k(w, \pi_2) \geqslant 0$. Since $G_k(w_k, \pi_2) \leqslant 0$, we must have $G_k(w_k, \pi_2) = 0$, which implies that is optimal to adopt with prior π_2. □

Thus, if the DM is not too risk averse in the states with bad outcomes, we will have policies that are monotonic in π, just as in the risk-neutral case and as one might expect to hold more broadly.

How restrictive are these utility conditions? Consider the case of a DM with a power utility $u(w) = -w^{1-\gamma}$ where $\gamma \geqslant 1$ ($\gamma = 1$ is the log utility); the risk tolerance is $\tau_u(w) = w/\gamma$. The risk tolerance bound of Proposition 7.1 can then be written as

$$\tau_u(w_0 + \underline{\theta} - c) \geqslant -\underline{\theta} \quad \Leftrightarrow \quad \frac{w_0 + \underline{\theta} - c}{\gamma} \geqslant -\underline{\theta}$$
$$\Leftrightarrow \quad w_0 - c \geqslant -(1+\gamma)\underline{\theta}. \qquad (6)$$

This condition can be interpreted as requiring the DM's wealth (net of information gathering costs) to be able to endure losses that are $(1+\gamma)$ times the worst possible technology outcome $\underline{\theta}$, without hitting the ruinous zero wealth state with $-\infty$ utility. A typical value of γ for an individual might range between 2 and 10, so the risk tolerance condition is satisfied for such an individual if the worst possible technology outcome represents less than $1/3$ to $1/11$ of the DM's wealth.

The risk tolerance bound of Proposition 7.1 can be sharpened given additional information about the utility function. In particular, the right side of the bound $(-\underline{\theta})$ is based on a first-order linear approximation (from above) where $u(w) \approx u(w + \underline{\theta}) - \underline{\theta} u'(w + \underline{\theta})$; see Equation (A.9). Since this linear approximation could be exact, we cannot improve the bound without making additional assumptions about the utility beyond it being DARA. However, if we assume a particular form for the utility function, we can improve the bound by calculating the relevant utilities exactly. For instance, if we consider a power utility with γ ranging from 2 to 10 (as above) and assume that the cost c of gathering information is less than 1% of the DM's wealth, the risk tolerance condition is satisfied if the worst possible technology outcome represents less than (approximately) 98.3% (for $\gamma = 2$) to 12.7% (for $\gamma = 10$) of the DM's wealth. Thus the risk-tolerance bound of Proposition 7.1 is conservative but can be improved given additional information about the utility function.

8. Changing Risk Attitudes

How do changes in risk tolerance affect policies? Considering the results for the illustrative example in Figure 2(a), we see that (i) if it is optimal to adopt in the risk-averse case, it is also optimal to adopt in the risk-neutral case and (ii) if it is optimal to reject in the risk-neutral case, it is optimal to reject in the risk-averse case. These results seem intuitive.

The rejection result is true in general: rejection yields a constant w and thus the certainty equivalent of rejecting is unaffected by changes in risk attitude. In contrast, adopting and waiting, as risky gambles, have larger certainty equivalents for a more risk-tolerant DM. Thus, we have the following.

Proposition 8.1 (Rejection Policies with Changing Risk Tolerance). *If it is optimal for one DM to reject given prior π, any more risk-averse DM should also reject given the same prior π.*

This implies that a risk-averse DM will reject a technology before a less risk-averse or risk-neutral DM, as was the case in the illustrative example, on the sample path shown in Figures 1(a) and 2(a).

As with changing priors, comparing the adopt and wait options with different risk tolerances is more complex as it requires comparing expected utilities or certainty equivalents for adopting and waiting, both of which will change with changes in risk tolerance. It is not too difficult to construct examples like that of Section 5.3 where increasing risk tolerance leads the DM to switch from adopting to waiting. For instance in the example of Figure 3(a), with a utility function $u(w) = \ln(w)$, as shown there, it is optimal to adopt immediately. If we instead consider the more risk tolerant utility function $u(w) = \ln(1+w)$, it is optimal to wait. Here, as discussed in Section 5.3, the cost of waiting puts the DM with a log utility in a state with near-zero wealth if the technology value turns out to be bad. However, the more risk tolerant DM with $u(w) = \ln(1 + w)$ is less sensitive to this incremental cost and finds waiting more attractive. Thus, we cannot hope for a general result that says increasing risk tolerance encourages adoption.

Can we identify reasonable assumptions on the utility functions—for example, bounds on the risk tolerances as in Proposition 7.1—that would rule out such counterintuitive examples? The short answer is no because such examples are not limited to scenarios with extreme risk aversion. To illustrate, again consider payoffs in the simple two-period example of Figure 3(a). Here we see that gathering information reduces the probability of encountering bad outcomes, but the cost of the information makes the outcomes worse. The fact that the bad outcomes become worse can lead risk-averse DMs to choose not to pay for information in cases where more risk-tolerant DMs will gather information. For example, if we had a DM with a risk-averse utility function who is indifferent between waiting and adopting in the example of Figure 3(a) (and prefers both waiting and adopting to rejecting), we can construct a more risk-averse utility function that assigns a lower utility to the worst outcome, leaving all other utilities unchanged. The DM with this new utility function would prefer adopting to waiting, despite being more risk averse.

We can, however, provide a positive result in the case where we compare the choices of a DM with constant absolute risk aversion (a CARA DM) with those of a risk-neutral DM. The proof uses the monotonicity of the optimal policy with respect to changes in the prior.

Proposition 8.2 (Adoption Policies for CARA and Risk-Neutral DMs). *Suppose the signal process satisfies the MLR property. If it is optimal for a CARA DM to adopt given prior π, then a risk-neutral DM should also adopt given the same prior π.*

Proof. See Online Appendix B.3. □

This result does not generalize to allow comparisons between a DM with an arbitrary risk-averse utility function with a risk-neutral DM. Similarly, one might speculate that the result of Proposition 8.2 would allow comparisons between two CARA DMs. That is, one might think that if it is optimal to adopt with constant risk tolerance τ_1, then it would also be optimal to adopt with constant risk tolerance $\tau_2 \geqslant \tau_1$. However, this is not necessarily true, as we demonstrate following the proof of Proposition 8.2 in the online appendix.

9. Discounting

So far we have not considered discounting. In practice, discounting may be an important consideration: information may be inexpensive or free, but time consuming to gather. In these cases, an important "cost" of gathering information is the delay in receiving the benefit of the technology. Discounting poses some new analytic challenges, but they can be addressed using the same tools we used earlier. In this section, we consider the model with discounting and then briefly discuss an extension with multiple information sources.

9.1. Model with Discounting

We will consider the model with discounting where, as is typical in the decision analysis literature (see, e.g., Baucells and Sarin 2007), the utility function is defined on current wealth plus the net present value (NPV) of the net benefit of the technology; all costs and benefits are discounted back to the present value, which is the value with T periods to go.[3] In this formulation, note that both positive and negative benefits are discounted, so delay reduces the risks (in NPV terms) associated with adoption.

Given a per-period discount factor δ $(0 \leqslant \delta \leqslant 1)$, let $\delta_k = \delta^{T-k}$ denote the discount factor for costs incurred or benefits received with k periods to go. The value function with discounting is then

$$
\begin{aligned}
&U_0(w,\pi) = u(w),\\
&U_k(w,\pi)\\
&\quad = \max\begin{cases} \mathbb{E}[u(w+\delta_k\tilde{\theta}) \mid \pi] & \text{(adopt)}\\ u(w) & \text{(reject)}\\ \mathbb{E}[U_{k-1}(w-\delta_k c, \Pi(\pi,\tilde{x})) \mid f(\pi)] & \text{(wait).}\end{cases}
\end{aligned} \tag{7}
$$

Note that with discounting ($\delta < 1$), if information is free ($c = 0$), the reject option will be (weakly) dominated by the option to wait, but the DM still faces a tradeoff between the information provided by waiting and the "cost" associated with the delay in receiving the benefits.

In this model with discounting, the results on the monotonicity of the value function (Proposition 4.1) and convexity of the value functions and policies (Proposition 4.2) continue to hold as before. Similarly, the policy results for rejection and adoption under risk-neutrality and with two outcomes continue to hold as before (Propositions 5.1–5.3, respectively). The monotonicity of the adoption policies with risk aversion also continues to hold as before, but we need to be careful with our selection of scaling functions and in defining the utility conditions. The difference between the value associated with immediate adoption and the optimal value function, $G_k(w,\pi) = \mathbb{E}[u(w+\delta_k\tilde{\theta}) \mid \pi] - U_k(w,\pi)$, can be written recursively as

$$
\begin{aligned}
&G_0(w,\pi) = \mathbb{E}[u(w+\delta_0\tilde{\theta}) - u(w) \mid \pi],\\
&G_k(w,\pi)\\
&\quad = \min\begin{cases} 0 & \text{(adopt),}\\ \mathbb{E}[u(w+\delta_k\tilde{\theta}) - u(w) \mid \pi] & \text{(reject),}\\ \mathbb{E}[u(w+\delta_k\tilde{\theta}) - u(w+\delta_{k-1}\tilde{\theta} - \delta_k c) \mid \pi],\\ \quad +\mathbb{E}[G_{k-1}(w-\delta_k c, \Pi(\pi,\tilde{x})) \mid f(\pi)] & \text{(wait).}\end{cases}
\end{aligned} \tag{8}
$$

The utility differences associated with waiting in (8) (analogous to the utility differences (4)) are now

$$
g_k(\theta) = u(w_k + \delta_k\theta) - u(w_k + \delta_{k-1}\theta - \delta_k c), \tag{9}
$$

where $w_k = w_T - ((1-\delta_k)/(1-\delta))c$ is the NPV of the DM's wealth with k periods to go, after paying the information cost c in all previous periods. We let $\theta_0 = -c/(1-\delta)$ denote the critical value such that the utility difference $g_k(\theta)$ is positive if $\theta > \theta_0$ and negative if $\theta < \theta_0$. In the negative case, the benefit of discounting bad outcomes by one additional period exceeds the cost of gathering information for that period.

How should we choose a scaling function $s(\theta)$ in this setting? For $\theta \leqslant \theta_0$, the utility differences $g_k(\theta)$

are not positive and increasing in θ; in this region, we will take $s(\theta)$ to be a positive constant. For $\theta > \theta_0$, $g_k(\theta)$ is positive but may be increasing or decreasing. If $g_k(\theta) > 0$ and $s(\theta) > 0$ and both are differentiable, $g_k(\theta)$ will be s-increasing if and only if

$$\frac{g_k'(\theta)}{g_k(\theta)} \geqslant \frac{s'(\theta)}{s(\theta)}. \tag{10}$$

(This follows from differentiating $g_k(\theta)/s(\theta)$ and rearranging.) In Section 7, we took $s(\theta)$ to be the utility difference with zero periods remaining ($g_0(\theta)$) because without discounting and with a DARA utility function, this utility difference is decreasing most rapidly at the lowest wealth level (that is, the ratio $g_k'(\theta)/g_k(\theta)$ on the left in (10) is decreasing in k) and (10) is satisfied for all k. However, with discounting, this utility difference need not be decreasing and the most rapidly decreasing utility difference may not be the one at the lowest wealth level. We will take the scaling function to be a positive constant if the ratios $g_k'(\theta)/g_k(\theta)$ are all positive and otherwise take $s(\theta)$ to track the minimum ratio in (10). Specifically, for $\theta > \theta_0$, we define

$$\phi(\theta) = \min\left\{0, \min_k \frac{g_k'(\theta)}{g_k(\theta)}\right\}. \tag{11}$$

We then pick some immaterial constant $K > 0$ and take $s(\theta) = K$ for $\theta \leqslant \theta_0$ and, for $\theta > \theta_0$,

$$s(\theta) = K\exp\left(\int_{q=\theta_0}^{\theta} \phi(q)\, dq\right), \tag{12}$$

Note that this $s(\theta)$ is positive and (weakly) decreasing and, for $\theta > \theta_0$, $s'(\theta)/s(\theta) = \phi(\theta)$. This construction thus ensures that the utility differences $g_k(\theta)$ will satisfy (10) and hence will be s-increasing.

With this choice of scaling function, we then have the analog of Proposition 7.1.

Proposition 9.1 (Utility Conditions with Discounting). *Suppose the DM is risk averse and her utility function $u(w)$ is DARA. Define the scaling function $s(\theta)$ as in (12); then for all k,*

(i) $u(w_k + \delta_k\theta) - u(w_k + \delta_{k-1}\theta - \delta_k c)$ *is s-increasing.*

(ii) $u(w_k + \delta_k\theta) - u(w_k)$ *is s-increasing if* $\tau_u(w_k + \delta_{k-1}\theta_k^* - \delta_k c) \geqslant -\theta_k^*$ *where* $\theta_k^* = \max\{\delta_k\underline{\theta}, \theta_0\}$.

Proof. The proof uses arguments similar to those in the proof of Proposition 7.1; see Online Appendix B.4. □

The monotonicity of the optimal adoption policies follows exactly as in Proposition 7.2.

Proposition 9.2 (Monotonicity of Adoption Policies with Risk Aversion and Discounting). *Suppose the assumptions of Proposition 9.1 are satisfied and the signal process satisfies the MLR property. If it is optimal to adopt with prior π_1, it is also optimal to adopt with any prior π_2 such that $\pi_2 \succeq_{LR} \pi_1$.*

Note that with discounting the "worst possible outcome" θ_k^* that must be considered in the risk tolerance bound is the larger of the lowest possible technology outcome (in NPV terms) $\delta_k\underline{\theta}$ and $\theta_0 = -c/(1-\delta)$. If information gathering is free ($c = 0$), then the risk tolerance bound reduces to the trivial requirement $\tau_u(w_k + \delta_{k-1}\theta_k^* - \delta_k c) \geqslant 0$; i.e., any degree of risk aversion is fine and the DARA assumption is also not necessary. With $c > 0$ and $\delta = 1$, we have $\theta_0 = -\infty$ and the results of Propositions 9.1 and 9.2 reduce to the corresponding results without discounting. With a CARA DM, we did not need a risk tolerance bound to establish the s-increasing properties without discounting (see Proposition 7.1). With discounting, a risk tolerance bound is needed in the CARA case: the bound given in Proposition 9.1 is sufficient, but, as discussed at the end of Section 7, this bound can be tightened given more information about the utility function.[4]

9.2. Multiple Information Sources

There are, of course, many possible ways one could extend this model, beyond incorporating discounting. One possibility is to include multiple information gathering options, with different costs and signal processes (described by their likelihood functions). For example, one might include a costless "passive waiting" option where the DM waits without observing any useful information. Such a passive waiting option would weakly dominate quitting, as it is costless. However, without discounting, there would be no benefit to passive waiting over quitting: if the DM is passively waiting, she would never learn anything and adoption would have no chance of becoming preferred to waiting or quitting. Similarly, if the DM is risk neutral or risk seeking, there is no benefit associated with passive waiting, even with discounting. But with discounting and risk aversion, such a passive waiting option may be attractive because, as mentioned earlier, delay reduces the risks (in NPV terms) associated with adoption and information gathering. Thus it may optimal to passively wait and then gather information or adopt later, when the gamble is less risky.

With multiple information sources, the model will have the same general structure as in the single source case. The value functions will be increasing and convex (as in Propositions 4.1 and 4.2(i)) and the adoption and rejection regions will be convex (as in Proposition 4.2(ii)). Given bounds on the risk tolerances like those of Proposition 9.1, we will have monotonic rejection and adoption thresholds: along any chain of LR-improving prior distributions, we move away from the rejection region and toward the adoption region, perhaps passing through the information gathering region. This can be established by generalizing the definition of the scaling function (12) to include utility differences (9) for each information source in the minimum of (11); see the discussion in Appendix B.5 for details. However, it is difficult to say much about the choice of information sources within the information gathering region. Clearly an information source that is cheaper and more informative (in the sense of Blackwell 1951) than other sources would be preferred, but it is not clear how the DM should trade off between sources with different costs and qualities.

10. Conclusions

In this paper, we have studied the impact of risk aversion on information gathering decisions in a technology adoption model. We find that most of the structural properties from the risk-neutral case continue to hold in the risk-averse case: the value functions are increasing and convex and the optimal policies are monotonic provided the DM is not too risk averse in scenarios with bad technology outcomes. We also showed that the risk-averse DM will reject a technology before a risk-neutral DM and, if the DM has a CARA utility function, the risk-averse DM will adopt later than a risk-neutral DM.

We view these results to be mostly positive: in most practical settings the information gathering policies should behave as expected. To return to the examples mentioned in the first paragraph of the paper, in the context of a consumer contemplating purchasing a Tesla, though the stakes may be large enough to induce considerable risk aversion, information is inexpensive and such a consumer is unlikely to be ruined (or nearly so) if the car turns out to be a disappointment. In this case, positive reviews would likely encourage adoption (i.e., buying the car), as we would intuitively expect. Similarly a farmer considering planting a new variety of soybeans or corn, information gathering (e.g., planting a new variety on a test plot) may be expensive, but a poor outcome would likely not be ruinous. The same would likely be true for a large electric utility considering building a power plant based on a new technology. In these cases, policies would also likely behave as expected. However, when the downside risks are extreme and information gathering is expensive—as it may be for an entrepreneur considering investing his life savings to attempt to commercialize a new idea—the optimal policies may exhibit counterintuitive nonmonotonicities like those in the example of Section 5.3. In this case, it may be optimal for the entrepreneur to invest in information gathering with one prior (π_2 in the example) but be optimal to plunge in and adopt the new technology without gathering additional information given a worse prior (like π_1). In such a case, a failed venture may be bad, but failing after paying significant information gathering costs may be ruinous. Reducing the probability of a bad technology outcome (moving from π_1 to π_2) may then encourage the entrepreneur to gather information.

Though the main goal of this paper was to study the technology adoption model with risk aversion, the proof techniques we used may prove useful in other contexts. The s-increasing property is a flexible generalization of the ordinary sense of increasing that is convenient for DPs and partially observable Markov DPs, as sums of s-increasing functions will be s-increasing and the sLR-increasing property survives Bayesian updating. The fact the s-increasing property implies the single-crossing condition may make it useful for studying policies in other DP models. In the technology adoption model, the desired monotonic policy results could not be established using standard increasing difference (i.e., supermodularity) arguments. However, given an appropriate scaling function, it was not difficult to show that the relevant policy differences are s-increasing; this then ensures the single-crossing properties required for the desired policy results.

Acknowledgments

The authors are grateful for the helpful feedback from the anonymous associate editor and referees.

Appendix A

A.1. Proof of Proposition 6.2

Proof. *Proposition* 6.2(i). We proceed as follows:

$$
\begin{aligned}
&S(\pi_1)V(\pi_2)-S(\pi_2)V(\pi_1)\\
&\underset{\text{(i)}}{=}\mathbb{E}[s(\tilde{\theta})\mid\pi_1]\mathbb{E}[v(\tilde{\theta})\mid\pi_2]-\mathbb{E}[s(\tilde{\theta})\mid\pi_2]\mathbb{E}[v(\tilde{\theta})\mid\pi_1]\\
&\underset{\text{(ii)}}{=}\int_{\theta_1}\int_{\theta_2} s(\theta_1)v(\theta_2)(\pi_1(\theta_1)\pi_2(\theta_2)-\pi_2(\theta_1)\pi_1(\theta_2))\,d\theta_1\,d\theta_2\\
&\underset{\text{(iii)}}{=}\iint_{\theta_2>\theta_1} s(\theta_1)v(\theta_2)(\pi_1(\theta_1)\pi_2(\theta_2)-\pi_2(\theta_1)\pi_1(\theta_2))\,d\theta_1\,d\theta_2\\
&\qquad+\iint_{\theta_2<\theta_1} s(\theta_1)v(\theta_2)(\pi_1(\theta_1)\pi_2(\theta_2)-\pi_2(\theta_1)\pi_1(\theta_2))\,d\theta_1\,d\theta_2\\
&\underset{\text{(iv)}}{=}\iint_{\theta_2>\theta_1} s(\theta_1)v(\theta_2)(\pi_1(\theta_1)\pi_2(\theta_2)-\pi_2(\theta_1)\pi_1(\theta_2))\,d\theta_1\,d\theta_2\\
&\qquad+\iint_{\theta_2>\theta_1} s(\theta_2)v(\theta_1)(\pi_1(\theta_2)\pi_2(\theta_1)-\pi_2(\theta_2)\pi_1(\theta_1))\,d\theta_1 d\theta_2\\
&\underset{\text{(v)}}{=}\iint_{\theta_2>\theta_1} \underbrace{(s(\theta_1)v(\theta_2)-s(\theta_2)v(\theta_1))}_{\text{(a)}}\\
&\qquad\cdot\underbrace{(\pi_1(\theta_1)\pi_2(\theta_2)-\pi_2(\theta_1)\pi_1(\theta_2))}_{\text{(b)}}\,d\theta_1\,d\theta_2\\
&\underset{\text{(vi)}}{\geqslant} 0
\end{aligned}
$$

Here, in the third equality, we decompose the region of integration into sets $\theta_1 < \theta_2$ and $\theta_1 > \theta_2$. (When $\theta_1 = \theta_2$, the integrand is zero.) In the fourth equality, we convert the second set into the first by changing variables $\theta_1 \to \theta_2$ and $\theta_2 \to \theta_1$. Rearranging gives the fifth equality. Here, the (a) term is nonnegative because v is s-increasing and the (b) term is nonnegative because $\pi_2 \geq_{LR} \pi_1$. The final inequality then follows and $V(\pi)$ is sLR-increasing. □

Proposition 6.2(ii). The proof of Proposition 6.2(ii) is analogous to a result of Karlin (1968, Chapter 3 Theorem 5.1). However, our assumptions are different from Karlin's and the proof has some important differences. We comment on these differences following the proof.

Proof. We first consider the case where $S(\pi_1)=0$. Since $s(\theta)$ is assumed to be single-crossing in θ, by Proposition 3.3, $S(\pi)$ is LR-single-crossing. By Proposition 6.1(v)(a), we have $V(\pi_1)\leqslant 0$. Since $S(\pi_2)\geqslant 0$, we have $V(\pi_1)S(\pi_2)\leqslant V(\pi_2)S(\pi_1)=0$, as desired. For the remainder of the proof, we assume that $S(\pi_1)>0$, which implies $S(\pi_2)>0$ since $\pi_2 \succeq_{LR} \pi_1$ and $S(\pi)$ is LR-single crossing (see Proposition 3.3).

Because $S(\pi)$ is linear in π, $\mathbb{E}[S(\Pi(\pi,\tilde{x}))\mid f(\pi)]=S(\pi)$. The desired result is then equivalent to

$$\begin{aligned}&\mathbb{E}[S(\Pi(\pi_1,\tilde{x}))\mid f(\pi_1)]\mathbb{E}[V(\Pi(\pi_2,\tilde{x}))\mid f(\pi_2)]\\&\quad-\mathbb{E}[S(\Pi(\pi_2,\tilde{x}))\mid f(\pi_2)]\mathbb{E}[V(\Pi(\pi_1,\tilde{x}))\mid f(\pi_1)]\geqslant 0 \qquad \text{(A.1)}\end{aligned}$$

To simplify, we introduce notation $V_{ij}=V(\Pi(\pi_i,x_j))$, $S_{ij}=S(\Pi(\pi_i,x_j))$, and $f_{ij}=f(x_j;\pi_i)$. Our standing assumption that $L(x\mid\theta)>0$ for all θ and x ensures that $f_{ij}>0$ and the posterior distributions $\Pi(\pi_i,x_j)$ are well defined. Our assumption that $S(\pi_1)>0$ (and hence $S(\pi_2)>0$) implies the priors have some mass in the region where $s(\theta)>0$, which implies the posteriors that $\Pi(\pi_i,x_j)$ will as well; this implies $S_{ij}>0$.

We can then rewrite (A.1) as

$$\begin{aligned}
&\mathbb{E}[S(\Pi(\pi_1,\tilde{x}))\mid f(\pi_1)]\mathbb{E}[V(\Pi(\pi_2,\tilde{x}))\mid f(\pi_2)]\\
&\quad-\mathbb{E}[S(\Pi(\pi_2,\tilde{x}))\mid f(\pi_2)]\mathbb{E}[V(\Pi(\pi_1,\tilde{x}))\mid f(\pi_1)]\\
&\underset{\text{(i)}}{=}\int_{x_1}\int_{x_2}[S_{11}f_{11}V_{22}f_{22}-S_{21}f_{21}V_{12}f_{12}]\,dx_1\,dx_2\\
&\underset{\text{(ii)}}{=}\int_{x_1}\int_{x_2}\left[\frac{V_{22}}{S_{22}}S_{11}S_{22}f_{11}f_{22}-\frac{V_{12}}{S_{12}}S_{12}S_{21}f_{12}f_{21}\right]dx_1\,dx_2\\
&\underset{\text{(iii)}}{=}\iint_{x_2>x_1}\left[\frac{V_{22}}{S_{22}}S_{11}S_{22}f_{11}f_{22}-\frac{V_{12}}{S_{12}}S_{12}S_{21}f_{12}f_{21}\right]dx_1\,dx_2\\
&\quad+\iint_{x_2<x_1}\left[\frac{V_{22}}{S_{22}}S_{11}S_{22}f_{11}f_{22}-\frac{V_{12}}{S_{12}}S_{12}S_{21}f_{12}f_{21}\right]dx_1\,dx_2\\
&\quad+\iint_{x_2=x_1}\left[\frac{V_{22}}{S_{22}}S_{11}S_{22}f_{11}f_{22}-\frac{V_{12}}{S_{12}}S_{12}S_{21}f_{21}f_{12}\right]dx_1\,dx_2\\
&\underset{\text{(iv)}}{=}\iint_{x_2>x_1}\left[\frac{V_{22}}{S_{22}}S_{11}S_{22}f_{11}f_{22}-\frac{V_{12}}{S_{12}}S_{12}S_{21}f_{12}f_{21}\right]dx_1\,dx_2\\
&\quad+\iint_{x_2>x_1}\left[\frac{V_{21}}{S_{21}}S_{12}S_{21}f_{12}f_{21}-\frac{V_{11}}{S_{11}}S_{11}S_{22}f_{11}f_{22}\right]dx_1\,dx_2\\
&\quad+\int_{x_1}\left[\frac{V_{21}}{S_{21}}S_{11}S_{21}f_{11}f_{21}-\frac{V_{11}}{S_{11}}S_{11}S_{21}f_{11}f_{21}\right]dx_1\\
&\underset{\text{(v)}}{=}\iint_{x_2>x_1}\Bigg[\underbrace{\left(\frac{V_{22}}{S_{22}}-\frac{V_{11}}{S_{11}}\right)}_{\text{(a)}}\underbrace{(S_{11}S_{22}f_{11}f_{22}-S_{12}S_{21}f_{12}f_{21})}_{\text{(b)}}\\
&\qquad+\Bigg(\underbrace{\frac{V_{22}}{S_{22}}-\frac{V_{12}}{S_{12}}}_{\text{(c)}}+\underbrace{\frac{V_{21}}{S_{21}}-\frac{V_{11}}{S_{11}}}_{\text{(d)}}\Bigg)\underbrace{S_{12}S_{21}f_{12}f_{21}}_{\text{(e)}}\Bigg]dx_1\,dx_2\\
&\quad+\int_{x_1}\underbrace{\left(\frac{V_{21}}{S_{21}}-\frac{V_{11}}{S_{11}}\right)}_{\text{(f)}}\underbrace{S_{11}S_{21}f_{11}f_{21}}_{(g)}\,dx_1\\
&\geqslant 0.
\end{aligned}$$

Here, in the third equality, we decompose the region of integration into sets $x_1<x_2$ and $x_1>x_2$ and $x_1=x_2$. (The set $x_1=x_2$ will have zero mass with continuous distributions but may have positive mass with discrete or more general distributions.) In the fourth equality, we convert the second set into the first by changing variables $x_1\to x_2$ and $x_2\to x_1$. We also rewrite the third integral taking into account the restriction to the set $x_1=x_2$ and substituting $x_2\to x_1$.

Each of the terms identified in the fifth expression above is nonnegative. Terms (e) and (g) are nonnegative because each term in the product is nonnegative. For (a), (c), (d), and (f), nonnegativity follows from the fact that V is sLR-increasing (as shown in Proposition 6.2(i)), which implies $V_{i_2,j_2}/S_{i_2,j_2}\geqslant V_{i_1,j_1}/S_{i_1,j_1}$ whenever $(i_2,j_2)\geqslant(i_1,j_1)$. Unpacking the notation, this is equivalent to

$$S(\Pi(\pi_{i_2},x_{j_2}))V(\Pi(\pi_{i_1},x_{j_1}))\leqslant S(\Pi(\pi_{i_1},x_{j_1}))V(\Pi(\pi_{i_2},x_{j_2}))$$

where $\Pi(\pi_{i_1},x_{j_1})\succeq_{LR}\Pi(\pi_{i_1},x_{j_1})$ whenever $(i_2,j_2)\geqslant(i_1,j_1)$; i.e., the posteriors are LR-dominant whenever the priors are LR-dominant and/or the signals are higher.

Nonnegativity of (b) is equivalent to showing that $S_{ij}f_{ij}$ is log-supermodular in (i,j). Unpacking the notation and using Bayes' rule, this is equivalent to the following:

$$\begin{aligned}
&S_{11}S_{22}f_{11}f_{22}-S_{12}S_{21}f_{12}f_{21}\\
&\quad=S(\Pi(\pi_1,x_1))f(x_1;\pi_1)S(\Pi(\pi_2,x_2))f(x_2;\pi_2)\\
&\qquad-S(\Pi(\pi_1,x_2))f(x_2;\pi_1)S(\Pi(\pi_2,x_1))f(x_1;\pi_2)\\
&\quad=\left(\int_\theta s(\theta)\frac{L(x_1\mid\theta)\pi_1(\theta)}{f(x_1;\pi_1)}\,d\theta\right)f(x_1;\pi_1)\\
&\qquad\cdot\left(\int_\theta s(\theta)\frac{L(x_2\mid\theta)\pi_2(\theta)}{f(x_2;\pi_2)}\,d\theta\right)f(x_2;\pi_2)\\
&\qquad-\left(\int_\theta s(\theta)\frac{L(x_1\mid\theta)\pi_2(\theta)}{f(x_1;\pi_2)}\,d\theta\right)f(x_1;\pi_2)\\
&\qquad\cdot\left(\int_\theta s(\theta)\frac{L(x_2\mid\theta)\pi_1(\theta)}{f(x_2;\pi_1)}\,d\theta\right)f(x_2;\pi_1)\\
&\quad=\left(\int_\theta s(\theta)L(x_1\mid\theta)\pi_1(\theta)\,d\theta\right)\left(\int_\theta s(\theta)L(x_2\mid\theta)\pi_2(\theta)\,d\theta\right)\\
&\qquad-\left(\int_\theta s(\theta)L(x_1\mid\theta)\pi_2(\theta)\,d\theta\right)\left(\int_\theta s(\theta)L(x_2\mid\theta)\pi_1(\theta)\,d\theta\right)\\
&\quad\geqslant 0.
\end{aligned}$$

In the third equality, we cancel the predictive distributions $f(x_i;\pi_j)$ with the denominators in the posterior distributions. The final inequality is equivalent to

$$\int_\theta L(x_i\mid\theta)\pi_j(\theta)(s(\theta)\,d\theta) \qquad \text{(A.2)}$$

being log-supermodular in (i,j), which can be established using the "basic composition formula" for log-supermodular functions (see, e.g., Karlin 1968). By assumption $L(x_i\mid\theta)$ is log-supermodular in (x_i,θ) (this is equivalent to the monotone likelihood assumption) and $\pi_j(\theta)$ is log-supermodular in (j,θ) (this is equivalent to assuming $\pi_2 \succeq_{LR} \pi_1$). Taking the measure to be $(s(\theta)\,d\theta)$, the basic composition formula then implies that (A.2) is log-supermodular or, equivalently, that the inequality above holds. This implies the (b) term above is nonnegative, thereby completing the proof. □

We now comment on the differences between this proof and that of Karlin (1968, Chapter 3 Theorem 5.1). First, Karlin requires the functions V and S to be positive. However V appears only in the numerator of fractions (in terms (a), (c), (d), and (h) above) and negative values for V do not cause difficulties. Second, Karlin assumes that S_{ij} and V_{ij} are both log-supermodular in (i,j). This is not assumed here and indeed may not be true in our setting. Specifically, Karlin uses the assumption that S_{ij} is log-supermodular in (i,j) to show term (b) above is nonnegative. We show (b) is nonnegative using the fact the likelihoods and denominators of the posterior distributions cancel; this then reduces the problem to a setting where we can use the basic composition formula to show (b) is nonnegative.

A.2. Proof of Proposition 7.1

Proof. (i) To show $u(w_0+\Delta+\theta)-u(w_0+\Delta+\theta-c)$ is s-increasing, let

$$h(\theta)=\frac{u(w_0+\Delta+\theta)-u(w_0+\Delta+\theta-c)}{u(w_0+\theta)-u(w_0+\theta-c)}. \tag{A.3}$$

We want to show $h(\theta)$ is increasing. Differentiating h and rearranging, we find that $h(\theta)$ is increasing if

$$\frac{u'(w_0+\Delta+\theta)-u'(w_0+\Delta+\theta-c)}{u(w_0+\Delta+\theta)-u(w_0+\Delta+\theta-c)} \geqslant \frac{u'(w_0+\theta)-u'(w_0+\theta-c)}{u(w_0+\theta)-u(w_0+\theta-c)}. \tag{A.4}$$

For $\Delta=0$, (A.4) holds trivially with equality.

For $\Delta>0$, (A.4) is implied by the assumption that u is DARA. To see this, define $f(u)=u'(u^{-1}(u))$ (the inverse is well defined since u is assumed to be strictly increasing) and note that $f'(u)=-\rho_u(u^{-1}(u))$, which is assumed to be increasing. Thus $f(u)$ is convex. We can rewrite (A.4) as

$$\frac{f(u_2)-f(u_1)}{u_2-u_1} \geqslant \frac{f(u_4)-f(u_3)}{u_4-u_3}. \tag{A.5}$$

where $u_2=u(w_0+\Delta+\theta)$, $u_1=u(w_0+\Delta+\theta-c)$, $u_4=u(w_0+\theta)$, and $u_3=u(w_0+\theta-c)$. Note that $u_2>u_1$, $u_4>u_3$, $u_2>u_4$, and $u_1>u_3$. We can interpret (A.5) as a comparison of slopes of two chords of a convex function, with the chord for the left term involving larger values (u_2,u_1) than the chord for the right (u_4,u_3). Since f is convex, the slopes of these chords are increasing with larger values; thus (A.5) holds.

(ii) To show $u(w_0+\Delta+\theta)-u(w_0+\Delta)$ is s-increasing, let

$$h(\theta)=\frac{u(w_0+\Delta+\theta)-u(w_0+\Delta)}{u(w_0+\theta)-u(w_0+\theta-c)}. \tag{A.6}$$

We want to show this $h(\theta)$ is increasing. Note that the denominator is positive and decreasing (since u is increasing and concave). If $\theta\geqslant 0$, the numerator in (A.6) is nonnegative and increasing in θ (since u is increasing). Thus, if $\theta\geqslant 0$, $h(\theta)$ is increasing.

Now assume that $\theta<0$. Taking the derivative of h and rearranging, we find that $h(\theta)$ is increasing if

$$\frac{u'(w_0+\theta)-u'(w_0+\theta-c)}{u(w_0+\theta)-u(w_0+\theta-c)} \geqslant \frac{u'(w_0+\Delta+\theta)}{u(w_0+\Delta+\theta)-u(w_0+\Delta)}. \tag{A.7}$$

As in the proof of part (i) of this proposition, the term on left of (A.7) can be interpreted as the slope of a chord of the convex function $f(u)=u'(u^{-1}(u))$, taken at points $u_2=u(w_0+\theta)$, $u_1=u(w_0+\theta-c)$, where $u_2>u_1$. Since f is convex, the derivative of the function f, $f'(u)=-\rho_u(u^{-1}(u))$, evaluated at the smaller value u_1 must be less than the slope of this chord. Thus we have

$$\frac{u'(w_0+\theta)-u'(w_0+\theta-c)}{u(w_0+\theta)-u(w_0+\theta-c)} \geqslant -\rho_u(w_0+\theta-c). \tag{A.8}$$

Now consider the right side of (A.7). Using a Taylor series expansion of u at $w_0+\Delta+\theta$, we can write

$$u(w_0+\Delta)=u(w_0+\Delta+\theta)-\theta u'(w_0+\Delta+\theta)+\tfrac{1}{2}\theta^2 u''(w_0^*) \tag{A.9}$$

where $w_0+\Delta+\theta \leqslant w_0^* \leqslant w_0+\Delta$ (recall that we are considering the case where $\theta<0$). We can then write the right side of (A.7) as

$$\begin{aligned}&\frac{u'(w_0+\Delta+\theta)}{u(w_0+\Delta+\theta)-u(w_0+\Delta)}\\ &\quad=\frac{u'(w_0+\Delta+\theta)}{\theta u'(w_0+\Delta+\theta)-(1/2)\theta^2 u''(w_0^*)} \leqslant \frac{1}{\theta}\end{aligned} \tag{A.10}$$

where the inequality follows because u is assumed to be concave (thus $u''(w_0^*)\leqslant 0$). Combining (A.8) and (A.10), we see that (A.7) holds if $-\rho_u(w_0+\theta-c)\geqslant 1/\theta$ or, equivalently, if $\tau_u(w_0+\theta-c)\geqslant -\theta$.

Given that the risk tolerance $\tau_u(w)$ is assumed to be increasing, we can check this risk tolerance bound by considering $\underline{\theta}$, the smallest possible value of θ, as stated in Proposition 7.1(b).

We can verify that the desired results hold with CARA utility functions u, i.e., $u(w)=-\exp(-w/\tau)$ by differentiating the functions $h(\theta)$ defined in (A.3) and (A.6). □

Endnotes

[1] In US (2009), we included a fixed cost paid at the time of adoption. This fixed cost can be incorporated into the net benefit θ, without loss of generality.

[2] Note that the value function given adoption in Figure 2(b) is concave with respect to the mean of the underlying beta distributions, holding the precision constant. In Proposition 4.2, we are considering convex combinations (mixtures) of general distributions. A convex combination of two beta distributions with the same precision would not yield a beta distribution. Thus there is no contradiction between the result of Proposition 4.2 and the concavity evident in Figure 2(b).

[3] Keeney and Raiffa (1993) describe three methods for discounting income streams in decision analysis models with risk aversion. The first is to calculate expected utilities for each period and discount these expected utilities. The second is to take certainty equivalents for each period and discount these certainty equivalents. The third approach, which we follow here, is to discount the income streams to NPVs and assess a utility function over NPVs. This third approach is typically used in decision analysis practice; see, e.g., McNamee and Celona (2005). Baucells and Sarin (2007) show that only the third approach satisfies certain desirable concavity and discounting properties; they conclude this is the "best approach" (p. 95) and discuss other issues.

[4] In the risk-seeking case, the utility differences (9) are negative for $\theta < \theta_0$ and may be increasing or decreasing in this region. However, the differences are increasing when $\theta > \theta_0$. We can establish monotonicity of adoption policies for the risk-seeking case by using a scaling function $s(\theta)$ such that $s(\theta) = 0$ for $\theta \leqslant \theta_0$ and $s(\theta)$ is a positive constant for $\theta > \theta_0$.

References

Athey S (2002) Monotone comparative statics under uncertainty. *Quart. J. Econom.* 117(1):187–223.

Abbas AE, Bakir NO, Klutke G-A, Sun Z (2013) Effects of risk aversion on the value of information in two-action decision problems. *Decision Anal.* 10(13):257–275.

Baucells M, Sarin R (2007) Evaluating time streams of income: Discounting what? *Theory Decision* 63(2):95–120.

Bickel JE (2008) The relationship between perfect and imperfect information in a two-action risk-sensitive problem. *Decision Anal.* 5(3):116–128.

Blackwell D (1951) The comparison of experiments. *Proc. Second Berkeley Sympos. Math. Statistn. Probab.* (University of California Press, Berkeley, CA), 93–102.

Chen X, Sim M, Simchi-Levi D, Sun P (2007) Risk aversion in inventory management. *Oper. Res.* 55(5):828–842.

Cho S-H, McCardle KF (2009) The adoption of multiple dependent technologies. *Oper. Res.* 57(1):157–169.

Eeckhoudt L, Gollier C, Schlesinger H (1995) The risk-averse (and prudent) newsboy. *Management Sci.* 41(5):786–794.

Hilton RW (1981) The determinants of information value: Synthesizing some general results. *Management Sci.* 27(1):57–64.

Jensen R (1982) Adoption and diffusion of an innovation of uncertain profitability. *J. Econom. Theory* 27(1):182–193.

Karlin S (1968) *Total Positivity*, Vol. 1 (Stanford University Press, Stanford, CA).

Keeney RL, Raiffa H (1993) *Decisions with Multiple Objectives: Preferences and Value Trade-offs* (Cambridge University Press, Cambridge, UK).

LaValle IH (1968) On cash equivalents and information evaluation in decisions under uncertainty: Parts I, II, and III, *J. Amer. Statist. Assoc.* 63(321):252–290.

Lippman SA, McCardle KF (1987) Does cheaper, faster, or better imply sooner in the timing of innovation decisions? *Management Sci.* 33(8):1058–1064.

Lovejoy WS (1987a) Ordered solutions for dynamic programs. *Math. Oper. Res.* 12(2):269–276.

Lovejoy WS (1987b) Some monotonicity results for partially observed Markov decision processes. *Oper. Res.* 35(5):736–743.

McCardle KF (1985) Information acquisition and the adoption of new technology. *Management Sci.* 31(11):1372–1389.

McNamee P, Celona J (2005) *Decision Analysis for the Professional* (SmartOrg, Menlo Park, CA).

Milgrom P, Shannon C (1994) Monotone comparative statics. *Econometrica* 62(1):157–180.

Quah JK-H, Strulovici B (2012) Aggregating the single crossing property. *Econometrica* 80(5):2333–2348.

Rogers EM (2003) *Diffusion of Innovations*, 5th ed. (Free Press, New York).

Smith JE (1998) Evaluating income streams: A decision analysis approach. *Management Sci.* 44(12):1690–1706.

Smith JE, McCardle KF (2002) Structural properties of stochastic dynamic programs. *Oper. Res.* 50(5):796–809.

Topkis DM (1978) Minimizing a submodular function on a lattice. *Oper. Res.* 26(2):305–321.

Ulu C, Smith JE (2009) Uncertainty, information acquisition and technology adoption. *Oper. Res.* 57(3):740–752.

Wald A (1945) Sequential tests of statistical hypotheses. *Ann. Math. Statist.* 16(2):117–186.

Zhu W, Kapuscinski R (2015) Optimal operational versus financial hedging for a risk-averse firm. Working paper, Singapore Management University.

James E. Smith is J.B. Fuqua Professor of Business Administration at the Fuqua School of Business, Duke University. His research interests are primarily in decision analysis and focus on developing methods for formulating and solving dynamic decision problems. Smith is an INFORMS Fellow.

Canan Ulu is an associate professor at Georgetown's McDonough School of Business. Her research focuses on Bayesian learning in sequential decision problems, decision making under multiple criteria and improving decision analysis methods using behavioral decision theory.