

# Risk Sensitive Value of Changed Information for Selective Querying of Web Services

John Harney and Prashant Doshi

THINC Lab, Dept. of Computer Science,  
University of Georgia, Athens, GA 30602  
{jfh,pdoshi}@cs.uga.edu

**Abstract.** A key challenge associated with compositions is that they must often function in volatile environments, where the parameters of the component Web services may change during execution. Failure to adapt to such changes may result in sub-optimal compositions. Value of changed information (VOC) offers a principled and recognized approach for selectively querying component services for their revised information. It does so in a rational (risk neutral) way. However, risk preferences often constitute an important part of the organization's decision analysis cycle and determine its desired business goals. We show how VOC may be generalized to consider preferences such as risk seeking and aversion using a utility based approach. Importantly, considerations of risk preferences lead to different services being used in the compositions and selected for querying for revised information. This is intuitive and provides evidence toward the validity of our approach for modeling risk preferences in VOC.

## 1 Introduction

A key benefit of Web services (WS) is the potential for automatically and flexibly formulating compositions of services resulting in integrated software processes. Many of the proposed approaches for optimally composing WSs utilize fixed information about the component WSs available at design time [1, 18, 21, 24]. A notable exception is [14] where parameters of WSs are obtained again just before execution. However, the approaches predominantly result in compositions that may soon become suboptimal if the environment is volatile. For example, a product may go out of stock affecting the usefulness of the corresponding WS, the network bandwidth may fluctuate affecting the WS response time, or costs of using a service may increase. Such changes may negatively impact the performance of the composition, making it critical that the compositions identify and adapt to changes appropriately to maintain optimality.

This is not a new problem and several adaptation techniques have been proposed [4, 7, 9, 17] to improve composition performance in the presence of volatile information. They seek to maintain an updated model of the composition environment, some of them by querying for revised information [9, 17]. One such approach uses the *value of changed information (VOC)* to decide which WS to query for its revised parameters [8, 9]. VOC is a principled and recognized way for selectively querying WSs participating in a composition. Recently, He et al. [10] expanded the applicability of VOC to multiple

different workflow patterns. Experimental results reveal that VOC based querying leads to adaptive compositions that perform better than those that do not adapt, in simulations.

Approaches such as VOC aim to adapt compositions in a rational – risk neutral – manner. However, risk preferences often strongly influence a modern organization’s decision analysis cycle and play a pivotal role in determining the goals of its many different business processes. This was demonstrated in a recent survey conducted by Corner and Corner [6], which found that more than a quarter of all business processes show some sensitivity to risk. For example, an organization may be conservative by nature, willing to sacrifice some cost in exchange for more stability and reduced risk of incurring greater costs in the future. In contrast to such *risk aversion*, *risk-seeking* behavior involves making decisions that could yield large gains at the risk of incurring heavier losses.

Clearly, a comprehensive approach to composition and adaptation should allow for considerations of risk preferences. This is because distinct preferences toward risk could significantly affect which WSs are selected in the composition and how the composition is adapted. In this paper, we generalize VOC toward modeling risk preferences in deciding which WS to query for revised information. Although we focus on VOC, our general approach toward considering risk is applicable to other adaptation techniques as well. We model risk preferences using *utility functions* [22] – a well-known way of modeling different attitudes toward risk. However, its principled integration within the traditional VOC is not trivial.

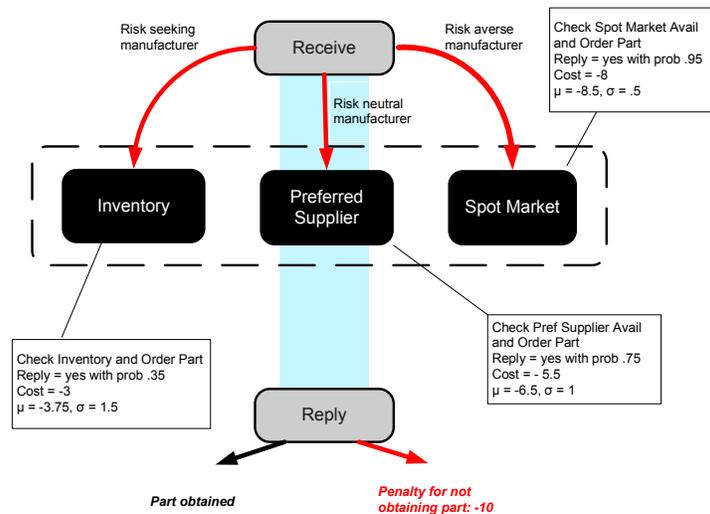
We show that considerations of risk preferences affect which WSs are included for participation in compositions, and which WSs are selected for querying of revised information. This in turn affects how compositions are adapted. We experimentally demonstrate that while the rational costs of the adapted compositions may be high, its utility to the risk-sensitive designer is optimal. This conforms to the intuition generally held about risk and its role in automation. It also provides evidence about the validity of our risk-sensitive model. The outcome is a novel approach for guiding the querying of WSs in volatile environments that models and acts on risk preferences as well.

The remainder of this paper is organized as follows. Section 2 presents a composition problem that is simple to understand yet demonstrative of the impact of considering risk preferences. Section 3 gives a brief overview of a specific composition technique and VOC. Section 4 shows how we may extend these models to include risk preferences. In Section 5, we present our experimental results that demonstrate the validity of our approach. Section 6 gives a concise summary of research related to this paper. We conclude with Section 7 which gives a brief summary and ideas for future work.

## 2 Scenario: Risk Sensitive Supply Chain

In order to demonstrate our approach, we present a *manufacturer’s* supply chain process designed according to the supply-chain operations reference (SCOR) [25] model. We assume that the component activities are all available as WSs. As the SCOR model is highly complex, we focus on the subtask of parts procurement and assembly, where the manufacturer must obtain a specific part to finish assembling a larger product for delivery.

We illustrate the activities involved in parts procurement in Fig. 1. The manufacturer must optimally decide between three different service providers from whom to obtain specific parts. The first option is to obtain the parts from its own inventory (shown as *Inventory* in Fig. 1), an inexpensive option that would allow the manufacturer to acquire the part quickly and cheaply. The manufacturer has limited storage available, however, making this method of obtaining the part unreliable. The manufacturer may also choose to obtain the parts directly from its preferred supplier. The part is more expensive to obtain from the supplier, but its availability is significantly better than from the inventory. Finally, the manufacturer may rely on the spot market, which almost certainly guarantees that the part will be obtained, but is more expensive than the previous two options. For the sake of simplicity, we assume that the manufacturer may invoke a single WS only. If the manufacturer is unable to complete the parts procurement task (i.e., a service is used that is not able to satisfy the order for the parts), a penalty is incurred in addition to the cost of invoking the WS. This penalty is representative of the recovery costs needed to heal the process (manufacturing halts, process rollbacks, compensations, etc.). In Fig. 1, we also indicate example QoS properties of each of the available services.



**Fig. 1.** Parts procurement in a manufacturer's supply chain. The manufacturer may choose between obtaining parts from its own inventory WS, a preferred supplier service or the spot market service. The inventory supplies parts cheaply but is less available in general, while the spot market provides parts reliably but expensively. The preferred supplier has intermediate cost and availability of parts. A penalty is incurred if the process does not complete.

A rational (risk neutral) manufacturer would optimize the composition by selecting the WS that maximizes the mathematical expectation of total cost given the probability of obtaining the parts using the WS. For the quality of service (QoS) parameters given

in Fig. 1, a rational manufacturer would elect to use the preferred supplier. However, pragmatic manufacturers are not always rational. A manufacturer that is *averse* to risk would opt for a more sure bet – a reliable service whose probability of meeting the order is thought to be high – despite the potentially higher cost. For our scenario, it seems intuitive that a risk-averse manufacturer would likely select the spot market due to almost guaranteed availability of the parts. On the other hand, a *risk seeker* would likely bet on a WS that is least expensive despite its lower reliability, such as the inventory. Subsequently, although our scenario is simple, it is appropriate for demonstrating the impact that distinct risk preferences may have.

An adaptive manufacturer would seek to modify its choice based on the updated information about the component WSs. In this regard, we indicate the mean ( $\mu$ ) and standard deviation ( $\sigma$ ) that guides the Gaussian distribution of the volatile cost of each participating service, in Fig. 1. While we restrict our analysis to Gaussian distributions, it applies to other distributions as well. Per domain knowledge, we expect inventory to be most volatile (high standard deviation) while the spot market to be least volatile.

### 3 Background: Web Service Composition and VOC

We briefly describe a framework for WS composition that provides a guarantee of expected cost based optimality. Consequently, we model both the cost and availability of the services that may participate in the composition. We then describe the *value of changed information* and how it may be used to pick a WS to query for revised information. While we show its usage in the context of the previous WS composition framework, Chafle et al. [5] demonstrate that VOC may be coupled with other WS composition frameworks as well. We refer the reader to [7] and [9] for more details about the WS composition framework and VOC, respectively.

#### 3.1 WS Composition Using Markov Decision Processes

We model the problem of WS composition as a Markov decision process (MDP) [19] using the sextuplet:

$$WP = (S, A, T, R, H, s_0)$$

where  $S = \prod_{i=1}^n X^i$ ,  $S$  is the set of all possible states of the composition factored into a set,  $X$ , of  $n$  variables,  $X = \{X^1, X^2, \dots, X^n\}$ ;  $A$  is the set of all possible actions ie. WS invocations;  $T$  is the transition function,  $T : S \times A \rightarrow \Delta(S)$ , which specifies the probability distribution over the next states of the composition given the current state and action.  $T$  models the uncertain availability of the WS itself and the service it provides;  $R$  is a reward function,  $R : S \times A \times S \rightarrow \mathbb{R}$ , which specifies the reward obtained (or cost) for performing each transition that has a non-zero probability.  $R$  models the summation of the cost of invoking a WS and the cost of consuming the service it provides, and is thus often independent of the source and resulting states. For our example, the reward depends only on the action;  $H$  is the period of consideration over which the composition must be optimal, also known as the horizon. While  $1 \leq H \leq \infty$ , for the sake of understanding we use a horizon of 1 (greedy actions) in this

paper. Our approach is applicable to longer horizons in a straightforward manner; and  $s_0$  is the starting state of the composition.

In order to gain insight into the use of the above framework for WS composition, let us model the parts procurement problem of Section 2. The state of the corresponding WS composition is captured by the random variables – **Inventory available**, **Preferred-Supplier available**, **SpotMarket available** – and is a conjunction of assignments of either *Yes*, *No*, or *Unknown* to each variable. As we mentioned, actions are WS invocations:  $A = \{\text{Inventory}, \text{PreferredSupplier}, \text{SpotMarket}\}$ . The transition function,  $T$ , models the nondeterministic effect of a WS invocation. For example, invoking the PreferredSupplier WS causes **PreferredSupplier available** to be assigned *Yes* with a probability of  $T(\text{PreferredSupplier available}=\text{Yes}|\text{PreferredSupplier}, \text{Preferred-Supplier available}=\text{Unknown}) = 0.75$ . The reward function,  $R$ , models the cost of using a service. We model this cost as a combination (e.g., a sum) of the cost of invoking the WS and the cost of the parts the service offers. For example, the cost of using the inventory is  $R(\text{Inventory})$ .

Given that the manufacturer has modeled its WS composition problem in the above framework, it may apply standard MDP solution techniques to arrive at an optimal composition. Here, optimality is with respect to the expected cost summed over the horizon. These solution techniques use stochastic dynamic programming [19] for calculation of the optimal policy using *value iteration*:

$$V^n(s) = \begin{cases} 0 & s \in \mathbb{G} \\ \max_{a \in A} \sum_{s' \in S} T(s'|a, s)[R(s, a, s') + V^{n-1}(s')] & s \notin \mathbb{G} \end{cases} \quad (1)$$

where  $V^n(s)$  quantifies the maximum long-term expected reward of reaching each state  $s$  with  $n$  actions remaining to be performed, and could be written concisely as  $V^n(s) = \max_{a \in A} E_{s'}^{s, a}[r_0 + V^{n-1}(s')]$ ;  $\mathbb{G}$  represents the set of goal states, indicating that the composition has completed successfully. As we employ a one step approach,  $n = 1$  for our analysis. Here,  $V^0(s') = 0, \forall s' \in \mathbb{G}$ , and  $V^0(s') \leq 0, \forall s' \notin \mathbb{G}$ .

Because the reward is negative (cost), Eq. 1 implies that the expected value of a non-goal state is less than or equal to the expected value of a goal state. In our example,  $V^0(s') \leq 0$  for non-goal states represents the penalty of not procuring the desired parts from the selected supplier.

Once we know the expected reward (or cost) associated with each state, the optimal action for each state is the one which results in the maximum expected reward.

$$\pi^*(s) = \underset{a \in A}{\operatorname{argmax}} \sum_{s' \in S} T(s'|a, s)[R(s, a, s') + V^0(s')] \quad (2)$$

In Eq. 2,  $\pi^*$  is the optimal policy which is a mapping from states to action(s),  $\pi^* : S \rightarrow \Delta(A)$ . The WS composition is obtained by performing the WS invocation prescribed by the policy given the state of the composition and observing the result of the invocation in order to obtain the next state. Doshi et al. [7] details this procedure and provides an algorithm for translating the policy to a WS composition.

### 3.2 VOC for Selective Querying

The parameters of the participating services may change during the lifetime of a WS composition. For example, the cost of the parts from the preferred supplier service may increase (requiring an update of  $R$ ) or the probability with which the preferred supplier satisfies an order for parts may reduce (requiring an update of  $T$ ). In order to remain optimal in such volatile environments, one approach we may adopt is to query the component WSs for their revised parameters [17]. However, queries may be costly or tedious to perform. We must therefore manage them carefully.

Harney and Doshi [9] introduced the *value of changed information (VOC)* that intelligently selects which service to query for its revised information. VOC employs a myopic approach to information revision, in which we query a single provider at a time for new information. For example, the manufacturer may query for revised information about the cost of using the inventory thereby updating the reward function to  $R'(\text{Inventory})$ .

Let  $V_{\pi^*}(s|R')$  denote the expected reward of following the optimal policy,  $\pi^*$ , from the state  $s$  given updated costs,  $R'$ . Let  $V_{\pi}(s|R')$  be the expected reward of following the original policy,  $\pi$ , from state  $s$  in the context of the revised costs. The policy,  $\pi$ , is optimal in the absence of revised information. Since the updated costs are not known unless we query the service provider, we average over all possible values of the revised costs, using our subjective belief distribution over the values. We formulate the VOC pertaining to the revised costs at some state,  $s$ , as:

$$VOC_{R'}(s) = \int_r Pr(R'(s, a, s') = r)[V_{\pi^*}(s|R') - V_{\pi}(s|R')]dr \quad (3)$$

where  $R'(s, a, s')$  denotes the cost that may be queried and subsequently may get revised,  $Pr(\cdot)$  is our *belief* over the possible costs of WS  $a$ . As we mentioned previously, the cost may be independent of  $s$  and  $s'$ .

The subscript to  $VOC$ ,  $R'$ , denotes the revised information inducing the change. Intuitively, Eq. 3 represents how badly, on average, the original policy,  $\pi$ , performs in the changed environment as formalized by the MDP model with the revised  $R'$ . We may model our beliefs over the possible cost of a WS,  $Pr(R'(s, a, s') = r)$  in Eq. 3, using density functions, which could be obtained from the service provider initially in the service-level agreement or may be learnt from previous interactions. We let the densities for the WSs take the form of *Gaussian* density functions<sup>1</sup>.

Of all the participating WSs in the composition, we select the one whose possible new cost is expected to bring about the most change in the WS composition, and this change exceeds the cost of querying that provider. We first select the service provider associated with the WS for whom the VOC is maximum:

$$a = \underset{a \in A}{\operatorname{argmax}} [VOC_{R'}(s)] \quad (4)$$

Let  $VOC^*(s)$  represent the corresponding maximum VOC at state  $s$ . We query for new information only if the VOC due to the revised information in that state is greater than

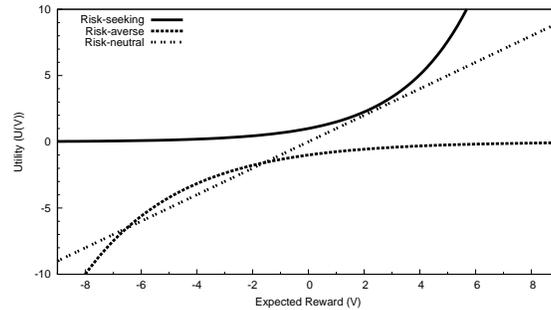
<sup>1</sup> Note that that these densities are marginalizations of the more complex ones that would account for all the factors that may influence the cost of the WS.

the query cost:  $VOC^*(s) > QueryCost(R'(a))$ . This revised parameter information is integrated into the MDP model, and the optimal policy,  $\pi^*$ , is regenerated. The WS composition resumes execution using an updated policy. Harney and Doshi [9] provide further details about the algorithm that integrates VOC within WS composition.

## 4 Risk Sensitive VOC

The traditional business analysis cycle views risk preferences as an important criteria for designing business processes [11]. Although rational process design stipulates risk indifference (often called risk neutrality), pragmatic composition design often involves either implicit or explicit considerations of risk preferences. For example, surveys [6] have found that 27% of all business processes are designed with some sensitivity to risk. Typically, risk considerations are predominant in processes involving high-stakes decisions, or those involving large sums of money or resources, in order to either avoid disastrous consequences or obtain huge financial gains [15]. A process designer may be *risk averse* – it is willing to incur some cost in exchange for more reliability and reduced risk of incurring greater costs in future. Some designers may have opposite preferences and are *risk seeking* – they make decisions that could yield large possible gains, at the risk of sustaining heavy losses.

A known way of modeling risk preferences is by adjusting the *utility function* that maps the actual expected reward to the subject’s utility [20, 22]. Risk aversion is associated with a large drop in utility for low reward (high cost) while risk-seeking behavior is thought to associate a large increase in utility for positive expected reward. On the other hand, risk neutrality involves considering the expected reward as is. We show example utility functions for the three distinct risk preferences in Fig. 2.



**Fig. 2.** Utility functions for different risk preferences. Risk aversion is modeled using a concave function while risk seeking is represented by a convex utility function. The utility function for risk neutrality is linear.

Typically, utility functions that model risk assume an exponential form. Consequently, we model risk-averse utility as,  $U^n(s) = \max_{a \in A} E_{s'}^{s,a} [-\gamma^{r_0 + V^{n-1}(s')}]$ , where  $\gamma$  is between 0 and 1 not inclusive, and  $n \geq 1$ . Parameter  $\gamma$  is often called the risk

factor and decides the degree of risk. Smaller values of  $\gamma$  signify greater aversion. Risk-seeking utility is modeled as,  $U^n(s) = \max_{a \in A} E_{s'}^{s,a} [\gamma^{r_0 + V^{n-1}(s')}]$ , where  $\gamma$  is greater than 1. Analogously, larger values of  $\gamma$  signify greater risk. If  $\gamma = 1$  we get the risk-neutral utility and  $U^n(s) = \max_{a \in A} E_{s'}^{s,a} [r_0 + V^{n-1}(s')]$ . In Fig. 2,  $\gamma = 1.5$  for the risk-seeking function while it is 0.75 for the risk-averse utility function.

As previously shown by Avila-Godoy [2], we may incorporate the utility functions in the solution of an MDP. Because the utility is an expectation of the exponential function of the value, we may rewrite Eq. 1 with  $n = 1$  as:

$$U^1(s) = \begin{cases} \iota & s \in \mathbb{G} \\ \max_{a \in A} [ \sum_{s' \in S} \gamma^{R(s,a,s')} T(s'|a,s) U^0(s') ] & s \notin \mathbb{G} \end{cases} \quad (5)$$

Here,  $\iota = \text{sgn}(\ln \gamma)$ ; in other words,  $\iota$  is 1 if  $\gamma > 1$  and -1 if  $0 < \gamma < 1$ . Notice that  $\iota$  is obtained by applying the appropriate utility function to the corresponding  $V^0(s)$  for goal states,  $\mathbb{G}$ , in Eq. 1. In case of risk aversion,  $U^0(s) = -\gamma^{V^0(s)} = -\gamma^0 = -1$ . For risk-seeking,  $U^0(s) = \gamma^{V^0(s)} = \gamma^0 = 1$ .  $U^0(s')$  is the appropriate utility function applied to  $V^0(s')$  for any state,  $s'$ . Thus,  $U^0(s') = \iota$  for all goal states and  $U^0(s') \leq \iota$  for all non-goal states. Note that the utility of non-goal states continues to be less than that of goal states.

We note that considerations of risk could impact the utility of the different states of the WS composition to the designer, and potentially which WS invocations are optimal at different states. This implies that the optimal policy,  $\pi^*$ , may be different as well leading to possibly distinct WS compositions for different utility functions.

Given the risk-sensitive composition using utility functions, we may generalize VOC to include considerations of risk preferences. In order to do this, Eq. 3 may be rewritten as:

$$\mathcal{VOC}_{R'}(s) = \int_r Pr(R'(s,a,s') = r) [U_{\pi^*}(s|R') - U_{\pi}(s|R')] dr \quad (6)$$

where  $\mathcal{VOC}$  denotes the generalized version of the traditional VOC,  $U_{\pi^*}(s|R')$  denotes the risk-sensitive utility of the state  $s$  given the optimal policy in the context of revised information and  $U_{\pi}(s|R')$  denotes the risk-sensitive utility of  $s$  given the original policy in the context of revised information. The utility function,  $U$ , is as defined in Eq. 5.

Subsequently, we select the WS to query which has the maximum  $\mathcal{VOC}$ , analogously to Section 3.2. However, the query is issued only if the  $\mathcal{VOC}^*$  is greater than the utility of the query cost to the designer. Formally, if  $\mathcal{VOC}^*(s) > U(\text{QueryCost})$ , then the query is issued to the provider whose service led to the maximum expected change.

Observe that considerations of risk preferences may alter which WS is selected for querying in comparison to risk neutrality. In our example, the composition of a risk-averse manufacturer may be more sensitive to changes in the parameters of the spot market. This is because the risk-averse manufacturer could be relying on the spot market to satisfy its parts order. This is in contrast to a risk-neutral manufacturer whose composition is expected to be most affected by changes in the parameters of the preferred supplier.

## 5 Experiments

We evaluate the performance of our risk-sensitive VOC and subsequent adaptation of the compositions in the context of our example problem domain of Section 2. Our service-oriented architecture involving VOC is identical to the one used previously by Harney and Doshi [9]. In particular, VOC computations are performed within internal WSDL WSs. The benefit of using VOC for selectively querying services for revised information and adapting compositions using the updated information has been demonstrated previously [8–10]. Specifically, compositions adapted using VOC lead to significantly better performance in volatile environments compared to compositions that are not adapted and to those that are adapted using adhoc techniques. Hence, we will not demonstrate the benefits of VOC here; rather we will focus on illustrating the influence of risk preferences on composition and adaptation, and thereby demonstrate the intuitive validity of our general approach.

### 5.1 Impact of Risk Preferences on Composition

Columns 1-3 of Table 1 summarize the cost and availability parameters of the participating supplier WSs that the manufacturer could use for procuring parts. Although the inventory is least costly, its likelihood of satisfying the order is the lowest. On the other hand, the spot market is most expensive but almost guaranteed to satisfy the manufacturer’s order.

WS	Reward	Availability	Volatility	
			Mean	Std. dev.
Inventory	-3	0.35	-3.75	1.5
Preferred Supplier	-5.5	0.75	-6.5	1
Spot Market	-8	0.95	-8.5	0.5

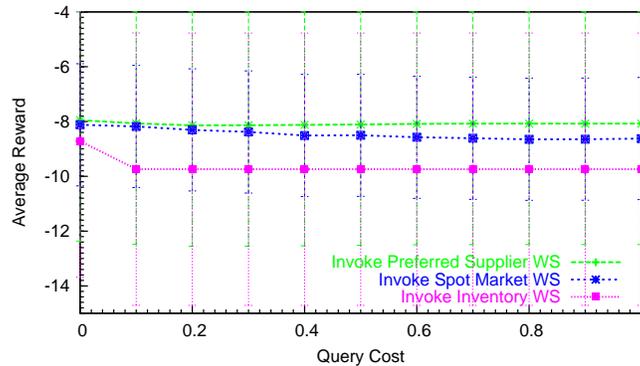
**Table 1.** Parameters of the participating WSs in our parts procurement problem domain. Existing rewards are slightly away from the means to facilitate potential querying.

A rational, risk-neutral manufacturer should choose the preferred supplier for placing its order. This is because the expected reward of using it (see Eq. 5 with  $\gamma = 1$  which becomes identical to Eq. 1) is the largest (lowest expected cost) among all three suppliers. However, let us consider the case where the manufacturer is *risk seeking* with  $\gamma = 1.5$ . The utility of the inventory for a risk-seeking manufacturer is the highest ( $U^1(s_0) = 0.11$ ) compared to that of the inventory and the spot market. This is intuitive because the inventory is a risky bet with the potential for a large reward (low cost of using it successfully) but also low chances of winning it (low availability). In contrast, both preferred supplier and spot market represent less attractive bets (action utility of 0.08 and 0.04, respectively) for the risk seeker. On the other hand, a risk-averse manufacturer with  $\gamma = 0.75$  opts for the spot market because it represents a safe bet – its order would almost certainly be met although the cost is high ( $U^1(s_0) = -18.36$ ).

These outcomes demonstrate the impact of risk preferences while formulating compositions. Differing compositions also influence how they get adapted using VOC; we investigate this next.

## 5.2 Impact of Risk Preferences on Adaptation

As we mentioned previously, in the absence of risk preferences a rational manufacturer would choose the preferred supplier's WS to order parts. Given the volatility in the environment (see columns 4 and 5 of Table 1), let the manufacturer use a VOC-driven approach toward querying service providers for revised information. If the updated information leads to a change in the optimal policy, the composition is adapted. In this context, we show the performance of the risk-neutral manufacturer's compositions in Fig. 3. We show the average reward obtained by three distinct *adaptive* compositions that invoke the inventory, preferred supplier and spot market WSs, respectively. Here, each data point is the average of 500 executions of the composition in 1,000 simulations of our problem domain. The simulations are constructed by sampling the Gaussian distributions of parameters of the participating volatile WSs. Each of the three different compositions experienced identical simulations in order to facilitate a valid comparison. A possible VOC-driven query is issued at the starting state of the composition.

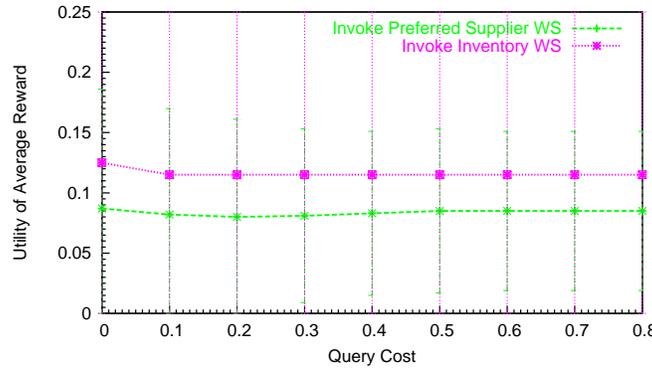


**Fig. 3.** Performance of the VOC-driven adaptive compositions for a risk-neutral manufacturer. We compare between policies prescribing the three different WSs. Higher average reward indicates better performance. The vertical bars represent the standard deviations. Large deviations are mainly due to the high penalty when the composition does not procure the parts.

We first observe that all three compositions show a drop in expected reward as the query cost increases. This is typical of VOC-driven adaptations. While adaptations are frequent when the query cost is low, they occur less as querying becomes more expensive. Notice that the composition prescribing the invocation of the preferred supplier continues to do the best. This implies that for low query costs, possible adaptations con-

tinue to outperform those in compositions that prescribed invoking other WSs. Furthermore, the risk-neutral manufacturer queried the preferred supplier the most because the associated VOC was often the largest. For large query costs, there is no adaptation and it is rational to choose the preferred supplier.

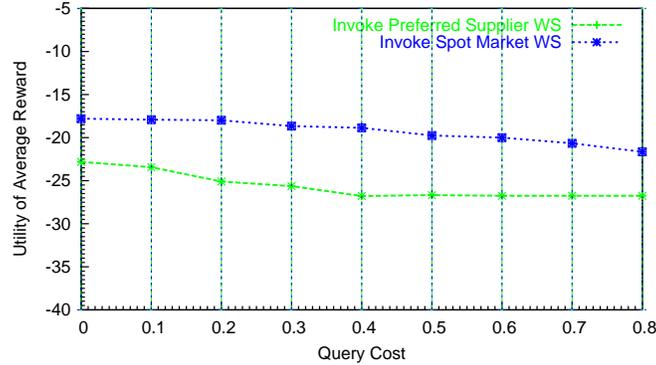
In Fig. 4, we demonstrate how the adaptive compositions of a *risk-seeking* manufacturer would perform when using VOC-driven selective querying. Our methodology for generating the data is same as before except for  $\gamma$  which is 1.5 and we show the utility of the reward obtained by the compositions. Risk-sensitive VOC was computed according to Eq. 6. As we mentioned previously, a risk-seeking manufacturer opts to invoke the inventory in comparison to a rational manufacturer who chooses the preferred supplier. Fig. 4 demonstrates that possible adaptations of this composition have a larger utility to the risk-seeking manufacturer than the adaptive preferred supplier based WS composition, which is the choice of a rational manufacturer. Thus, although the composition that prescribes the inventory performs worst in the absence of risk preferences, it becomes the optimal choice for a risk seeker.



**Fig. 4.** Utility of VOC-driven adaptive compositions for a risk-seeking manufacturer ( $\gamma = 1.5$ ). We compare between policies that recommend preferred supplier and inventory WSs. Higher utility indicates better performance.

We observed that for low query costs, the risk-seeking manufacturer queried the inventory the largest number of times indicating that its VOC was the largest. This is because of the high standard deviation of its cost coupled with its importance in the composition of the manufacturer. About a fraction of one-tenth of these queries led to adaptations where the preferred supplier was selected.

Finally, in Fig. 5 we demonstrate the utility to a risk-averse manufacturer of the performance of adaptive compositions. We use  $\gamma = 0.75$  (recall that risk aversion is modeled using  $\gamma < 1$ ) and each data point is generated as before. In comparison to a risk seeker, our risk-averse manufacturer continues to find the adaptive composition that recommends invoking the spot market most preferable. Its utility remains consistently high compared to that of the adaptive composition that involves the preferred supplier



**Fig. 5.** Utility of VOC-driven adaptive compositions for a risk-averse manufacturer ( $\gamma = 0.75$ ). We compare between policies that select preferred supplier and spot market WSs. Higher average utility indicates better performance.

which is the choice of a rational manufacturer. Furthermore, the risk-averse manufacturer queried the preferred supplier most number of times. This is intuitive because while the spot market is important its deviation is very low. Significant changes in preferred supplier's costs could make it the WS of choice for the risk-averse manufacturer. Despite the querying the number of times that adaptations do occur is low because of the significant value of the spot market to the risk-averse manufacturer. About one-fifth of these queries led to adaptations.

In summary, risk-sensitive VOC leads to different services being selected for querying based on the risk preferences, in comparison to queries in the absence of risk. This is in part due to the varying compositions induced by different risk preferences.

## 6 Related Work

Our paper borrows, in large part, from recent work in quantitative risk analysis. This line of research uses modern decision-theoretic planning and von-Neuman-Morgenstern utility theory concepts [22] that exist in the artificial intelligence [20] and decision analysis [12] literature. Avila-Godoy [2] derived straightforward methods to compute value functions for process designers that utilize exponential utility functions to model their risk attitudes. Liu [15] extended this line of work so that value functions may be computed using more general (i.e. non-exponential) risk-sensitive utility functions. While their focus was strictly on MDPs, these methods provide a means to include risk preferences in process creation.

Other lines of work have attempted to represent risk attitude in different ways. Kokash and D'Andrea [13] use traditional risk management strategies to derive contingency plans (such as quality of service re-negotiation or adopting other component services) when the risk of using a composition is found to be high. They operate with the notions of threats (danger sources), probabilities of threats, and their quantifiable impact on the provider of the composition (monetary losses, time losses and breach of

reputation). These threats are juxtaposed against the possible gains of the composition. Decisions are made whether to utilize a contingency plan of composition based on these comparisons. Wiesemann et al. [23] incorporate the average value at risk (AVaR) measure, widely used in economic studies, into the decision making of a WS composition. They introduce risk-preferences using the  $\beta$ -AVaR metric, which is defined as the mean value of the  $(1-\beta)$  worst losses sustained by making a particular decision. Parameter  $\beta \in [0, 1]$  represents the degree to which a decision maker considers the worst case loss of a particular decision. When  $\beta$  is 0, the decision maker is risk-neutral. As  $\beta$  increases it becomes more pessimistic, and thus, risk-averse. The  $\beta$ -AVaR metric is introduced in their value maximization equations and decisions are based on the modified equations. Although both approaches apply quantitative methods, they quantify the risks in different ways, and neither handle run-time adaptation of WS compositions in the presence of risk preferences as we do in this paper.

Finally, risk management in processes has long received attention in economic and business enterprise research communities. Traditionally, this line of research has adopted an orthogonal and external perspective to risk by seeking to identify the risks posed to business processes and outlining options that process managers may undertake to deal with the risks [16]. Unlike the formal concepts discussed in this paper, their strategies implement qualitative approaches to risk management. Human specialists oversee the processes and identify potential hazards that could upset its functionality. After the risks are identified, they are addressed in a systematic manner.

## 7 Discussion and Future Work

Real world environments in which WS compositions must function are seldom static. Parameters of WSs and the of service it provides such as the cost and availability are volatile. In this context, compositions must adapt to remain optimal. A crucial step in adaptation is finding out the updated information about the component WSs. VOC offers a principled and recognized approach for deciding which WS to query for revised information.

Although designers are often advised to be rational and objective while designing processes, risk preferences invariably play a role in the composition. Therefore, we presented an approach for considering risk preferences while composing and querying, leading to the *risk sensitive VOC*. We used utility functions to model risk preferences and showed how these may be integrated with the traditional VOC. Our experimental results on a simulated parts procurement domain confirmed our intuition about the impact of risk – its consideration leads to changes in how we compose and issue queries compared to rational behavior.

While we experimented with a simple scenario in order to promote clarity, we think that our results are indicative of more complex scenarios as well. In particular, the risk-sensitive utility as defined in Eq. 5 may be generalized straightforwardly to multiple steps and the corresponding computation of the risk-sensitive VOC continues to proceed as shown in Eq. 6.

Although we used particular exponential utility functions, other forms of exponential utility functions could be used as well such as those often utilized in economics

and those which allow a switch in risk preference depending on the accumulated reward [15]. We speculate that our analysis would hold for these forms of utility functions as well. Finally, it would be interesting to demonstrate the beneficial role of risk in real-world compositions. In this regard, we are seeking a case study.

## Acknowledgment

We would like to thank Bob Bostrom, professor in the business school of the University of Georgia for useful discussions related to this work.

## References

1. V. Agarwal, G. Chafle, K. Dasgupta, N. Karnik, A. Kumar, S. Mittal, and B. Srivastava. Synth: A system for end to end composition of web services. *Journal of Web Semantics*, 3:311–339, 2005.
2. M. Avila-Godoy. Controlled Markov Chains with Exponential Risk-Sensitive Criteria: Modularity, Structured Policies and Applications. Ph.D. thesis, Department of Mathematics, University of Arizona, 1999.
3. A. Andrieux, K. Czajkowski, A. Dan, K. Keahey, H. Ludwig, T. Nakata, J. Pruyne, J. Rofrano, S. Tuecke, and M. Xu. *WS-Agreement Specification*, 2005.
4. G. Chafle, K. Dasgupta, A. Kumar, S. Mittal, and B. Srivastava. Adaptation in web service composition and execution. In *International Conference on Web Services (ICWS), Industry Track*, pages 549–557, 2006.
5. G. Chafle, P. Doshi, J. Harney, S. Mittal, and B. Srivastava. Improved Adaptation of Web Service Compositions Using Value of Changed Information. In *International Conference on Web Services (ICWS), Industry Track*, pages 784–791, 2007.
6. J. Corner and P. Corner. Characteristics of decisions in decision analysis practice. In *The Journal of Operational Research Society*, Vol. 46, pages 304–314, 2006.
7. P. Doshi, R. Goodwin, R. Akkiraju, and K. Verma. Dynamic workflow composition using markov decision processes. *Journal of Web Services Research*, 2(1):1–17, 2005.
8. J. Harney and P. Doshi. Adaptive Web Processes Using the Value of Changed Information” In *International Conference on Service-Oriented Computing (ICSOC)*, pages 179–190, 2006.
9. J. Harney and P. Doshi. Selective querying for adapting web service compositions using the value of changed information. *IEEE Transactions on Services Computing*, in press, 2008.
10. Q. He, J. Yan, H. Jin, and Y. Yang. Adaptation of Web Service Composition Based on Workflow Patterns. *International Conference on Service-Oriented Computing (ICSOC)*, pages 22–37, 2008.
11. S. Holtzman. *Intelligent Decision Systems*. Addison-Wesley, 1989.
12. C. Kirkwood. Approximating Risk Aversion in Decision Analysis Applications. *Decision Analysis*, pages 55–72, 2004.
13. N. Kokash and V. D’Andrea. Evaluating Quality of Web Services: A Risk-Driven Approach. *Business Information Systems*, 2007, pages 180–194.
14. U. Kuter, E. Sirin, D. S. Nau, B. Parsia, and J. A. Hendler. Information gathering during planning for web service composition. *Journal of Web Semantics*, 3(2-3):183–205, 2005.
15. Y. Liu. Decision-Theoretic Planning Under Risk-Sensitive Planning Objectives. Ph.D. thesis, College of Computing, Georgia Institute of Technology, 2005.
16. M. zur Muehlen and D. Ho. Risk Management in the BPM Lifecycle. *Business Process Management Workshops*, pages 454–466, 2005.

17. T.C. Au, U. Kuter, and D. Nau. Web service composition with volatile information. In *International Semantic Web Conference (ISWC)*, 2005, pages 52–66.
18. J. Pathak, S. Basu, and V. Honavar. Moscoe: A specification-driven framework for modeling web services using abstraction, composition and reformulation. In *Ph.D. Symposium, International Conference on Services-Oriented Computing (ICSOC)*, pages 1–6, 2007.
19. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley series in probability and mathematical statistics. Wiley-Interscience, 1994.
20. S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach (Second Edition)* Prentice Hall, 2003.
21. E. Sirin, B. Parsia, D. Wu, J. A. Hendler, and D. S. Nau. Htn planning for web service composition using shop2. *Journal of Web Semantics*, 1(4):377–396, 2004.
22. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
23. W. Wiesemann, R. Hochreiter, and D. Kuhn. A Stochastic Programming Approach for QoS-Aware Service Composition. In *IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008)*, pages 226–233, 2008.
24. H. Zhao and P. Doshi. Haley: An end-to-end scalable web service composition tool. In *Developer's Track, World Wide Web (WWW)*, 2008.
25. The Supply-Chain Council [http://archive.supply-chain.org/cs/root/scor\\_tools\\_resources/scor\\_model/scor\\_model](http://archive.supply-chain.org/cs/root/scor_tools_resources/scor_model/scor_model) (last accessed June 14, 2010).