Risk-Sensitive Querying for Adapting Web Service Compositions

John Harney and Prashant Doshi

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

Abstract. Environments in which Web service compositions (WSCs) operate are often dynamic. To remain optimal, WSCs must adapt to the inherent changes that are present in dynamic environments. While previous WSC adaptation methods have yielded promising results, they fail to accomodate the risk preferences introduced by process designers. Risk preferences may influence which Web services are assembled and selected for execution in WSCs and thus must be considered in adaptation. Previously, the value of changed information (VOC) had been proposed to select those services for querying whose revised information is expected to bring about the most change in the composition. In this paper, we enhance VOC so that risk preferences are considered. We augment our previous model so that generalized utility functions (modeling different risk preferences) may integrated into standard value functions that determine optimal Web services to be used in the WSC. We demonstrate that our choice of which WS to query as prescribed by VOC may change if we consider risk preferences as well.

1 Introduction

Environments in which Web service compositions (WSCs) operate are often dynamic. Such environments introduce the problem of data volatility, where component Web services' (WSs') non-functional, or Quality of Service (QoS), parameters change during the execution of the WSC. Without considering data volatility, WSCs may become sup-optimal.

Recent approaches [1, 2, 3] have adopted sophisticated methods to address the data volatility problem and have shown promising results. These approaches maintain up-to-date knowledge of component WSs' parameters and hence sustain optimality in the presence of data volatility by selecting component WSs that are currently of the greatest *value* to the WSC in the changed environment. The value of a service is oftened viewed as a measure of how "good" or "desirable" that service may be to the WSC in comparison to other potential candidate services that may be selected. Often services that are expected to incur the least cost are chosen to be used in the WSC.

One particular aspect of the adaptation problem that is overlooked, however, is the lack of consideration for a process designer's *risk-preferences* and its subsequent impact on the WSC. Risk preferences can best be described as the definition of "optimality" of a process to its designer. Traditional techniques that do not explicitly consider risk preferences could be seen as generating compositions with *risk-neutral* preferences in mind. They base their decisions strictly upon maximization of value functions (or the minimization of cost functions). Not all preferences, however, share a similar view of optimality. A designer's preferences may be conservative by nature – she is willing to sacrifice some cost in exchange for more stability and reduced risk of incurring greater costs. Such preferences are said to be *risk-averse*. Other preferences may have opposite tendencies – they make decisions that could yield large possible gains, at the risk of incurring heavier losses. They are referred to as *risk-seeking* preferences.

As a concrete example, let us assume that a risk-averse manufacturer intends to use a WSC to execute its supply chain process. It must choose between supplier A (i.e. WS_A) with a cost of \$2000 and 50% availability, and supplier B (i.e. WS_B) with cost \$5000 and 100% availability. Given these parameters of suppliers A and B, a traditional, risk-neutral WSC would choose WS_A, because the expected cost of using WS_A would be less than using WS_B over the long term. Despite the lower expected cost of WS_A, risk-averse manufacturers would likely select WS_B. It would choose the more stable WS that was the "sure thing" at the expense of incurring the extra cost so that it could avoid a possible disastrous scenario of invoking WS_A several times only to have it fail repeatedly.

Changes in risk-neutral WSCs may impact the composition in different ways than changes in risk-sensitive WSCs. Using the example above, the risk-averse manufacturer will be more sensitive to a change in the availability of WS_B than if it had risk-neutral preferences. Identification of the WSs whose changes greatly affect the WSC are vital to creating a feasible adaptation scheme. Indeed, a consistent adaptation method is needed to accomodate all risk preferences.

In Harney and Doshi [4], a WSC is adapted in the presence of data volatility by selectively querying services for their revised parameters and using the new information to make better decisions. Our approach, which we called the *value of changed information (VOC)*, computes the trade-off between the cost of querying for up-to-date information and the value of expected change in the WSC that the revised information will bring. The model parameters are updated and the WSC is composed again, only if the VOC is greater than the query cost. In computing the VOC, the approach utilizes stochastic models of volatility of each of the services' parameters. The approach is *myopic* in that only one service provider is queried at a time, and the revised information for that WS is utilized which leads to the maximum VOC. The maximum VOC identifies the WS whose changed parameters will potentially have the greatest impact on the the WSC.

In this paper, we show how VOC may be made aware of different risk preferences. Extending the VOC in such a way comes with attendant challenges: (a) We must introduce new parameters that can model risk-preferences, and (b) we must transform our previous VOC model so as to include the parameters required to model risk-preferences. Given approaches that address both these challenges, we show that VOC, which identifies the WSs whose changes will impact the WSC the greatest, may be different for different risk preferences. This ensures that the WSCs adapt in accordance to the risk preferences that have been introduced to thme by their designers. We show that our approach works well for WSCs that have strictly risk-averse, risk-neutral, and risk-seeking preferences. We also include a general model that may be adapted in designing WSCs with more complex risk-preferences.

2 Scenario: Supply Chain with Risk Preferences

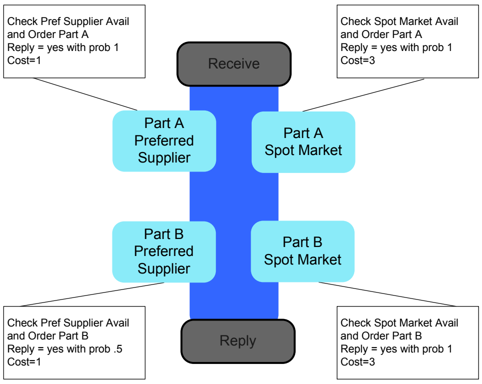

Fig. 1. A simple supply chain WSC implemented by a manufacturer. The manufacturer may choose between obtaining parts from a preferred supplier service provider or the spot market service provider. The preferred supplier can supply parts cheaper but in a less reliable manner than the spot market.

To illustrate how risk preferences influence adaptation in WSCs, we give an example scenario (shown in Fig. 1) of a simplified supply chain process of a manufacturer who wishes to implement it as a WSC. In order to assemble its product, the manufacturer must secure parts (parts A and B) from external vendors. The manufacturer must optimally decide between two different providers – their Preferred Supplier's WSs or the Spot Market WSs – from which to ascertain their parts. The Preferred Supplier can supply parts to the manufacturer cheaply, but occasionally may not be able to suitably satisfy a given order for the part. On the other hand, the required parts are consistently available from the Spot Market service, but more expensive to obtain.

A cost concious manufacturer would normally select component WSs that would minimize the costs in the long term. If she were to exhibit risk-neutral behavior, she would optimize the process based on the cost and availability metrics given in Fig. 1, and use the Preferred Supplier WS, as its *expected* cost would on average be lower than using the Spot Market WS. Because of its subpar reliability, however, the Preferred Supplier's WS may have a large variance in order satisfaction from one execution of the WSC to the next. A risk-averse manufacturer views this as undesirable, likely choosing the more stable Spot Market WS.

If the Spot Market WS's availability were to suddenly decrease, the change would impact the risk-averse manufacturer in a more profound way than if it were risk-neutral. This is due to the fact that the risk-averse manufacturer currently views the Spot Market as the optimal service to invoke. Thus, it is vital that the risk-averse manufacturer know the changes occurring in the Spot Market WS. Conversely, risk-neutral manufacturers would probably be more interested in availability fluctations in the Preferred Supplier service.

3 Background

For the purpose of illustration, we select a decision-theoretic planning technique for composing WSs [5]. We then review the formulation of the value of changed information (VOC) for participating WSs, and refer the reader to Harney and Doshi [4] for more details.

3.1 Web Service Composition Using Markov Decision Processes

Decision-theoretic planners such as Markov Decision Processes (MDPs) model the composition environment, WP, using a sextuplet:

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

where $S = \Pi_{i=1}^{n} X^i$, S is the set of all possible states factored into a set, X, of n variables, $X = \{X^1, X^2, \ldots, X^n\}$; A is the set of all possible actions; T is a transition function, $T : S \times A \rightarrow \Delta(S)$, which specifies the probability distribution over the next states given the current state and action; C is a cost function, $C : S \times A \rightarrow \mathbb{R}$, which specifies the cost of performing each action from each state; H is the period of consideration over which the plan must be optimal, also known as the horizon, $0 < H \leq \infty$; and s_0 is the starting state of the process.

As an example, let us model the problem of composing the component WSs of the supply chain process scenario in Fig. 1 as a MDP. The state of the WSC is captured by the random variables – **PreferredSupplierPartA available**, **SpotMarketPartA available**, **PreferredSupplierPartB available**, and **SpotMarketPartB available** – and is a conjunction of assignments of either *Yes*, *No*, or *Unknown* to each variable. Actions are WS invocations, A={PreferredSupplierPartA, SpotMarketPartA, PreferredSupplierPartB, and SpotMarketPartB}. The cost function, C, models the cost of invoking a service. The transition function, T, models the non-deterministic effect of each action on some random variable(s). For example, invoking the PreferredSupplierPartA WS will cause **PreferredSupplierPartA available** to be assigned *Yes* with a probability of T(**PreferredSupplierPartA available**=*Yes*|PreferredSupplierPartA, **PreferredSupplierPartA available**=*Unknown*).

Once the manufacturer has modeled its WS composition problem as a MDP, it may apply standard MDP solution techniques to arrive at an optimal process. These solution techniques revolve around the use of stochastic dynamic programming [6] for calculation of the optimal policy using *value iteration*:

$$V^n(s) = \underset{a \in A}{min} \;\; Q^n(s,a) \tag{1}$$

where:

$$Q^n(s,a) = \begin{cases} C(s,a) + \sum_{s' \in S} T(s'|a,s) V^{n-1}(s) & n > 0 \\ 0 & \text{n=0} \end{cases} \tag{2}$$

where the function, $V^n : S \rightarrow \mathbb{R}$, quantifies the minimum long-term expected cost of reaching each state with n actions remaining to be performed, and $Q^n(s,a)$ is the action-value function, which represents the long-term expected cost from s on performing action a.

Once we know the expected cost associated with each state, the optimal action for each state is the one which results in the minimum expected cost.

$$\pi^*(s) = \underset{a \in A}{argmin} \;\; Q^n(s,a) \tag{3}$$

In Eq. 3, π^* is the optimal policy which is simply a mapping from states to actions, $\pi^* : S \rightarrow A$. The WSC is composed by performing the WS invocation prescribed by the policy given the state of the process and observing the results of the actions to obtain the next state. Details of the algorithm for translating the policy to the WSC are given in [5].

3.2 Value of Changed Information

As discussed previously, the parameters of the participating services may change during the life-cycle of a WSC. For example, the cost of using the Preferred Supplier A service may increase (requiring an update of C) or the probability with which the Preferred Supplier can satisfy an order for part A may reduce (requiring an update of T). For simplicity, we focus on a change in T, though our approach is generalizable to fluctuations in other model parameters as well.

VOC [4] employs a myopic approach to information revision, in which we query a single provider at a time for new information. For example, the revised information may change the following transition probability,
$T($**PreferredSupplierPartA available** $= Yes \mid$ PreferredSupplierPartA$= Unknown)$.

Let $V_{\pi^*}(s|T')$ denote the expected cost of following the optimal policy, π^*, from the state s when the revised transition function, T' is used. Since T' is not known unless we query the service provider, we average over all possible values of the revised transition probability, using our belief distribution over the values. Formally,

$$EV(s) = \int_{\mathbf{p}} Pr(T'(\cdot|a,s') = \mathbf{p}) V_{\pi^*}(s|T') d\mathbf{p} \tag{4}$$

where $T'(\cdot|a, s')$ represents the distribution that may be queried and subsequently may get revised, $\mathbf{p} = \langle p_1, p_2, \ldots, p_m \rangle$ represents a possible response to the query (revised distribution), m is the number of values that the variable under question may assume, and $Pr(\cdot)$ is our *belief* over the possible distributions.

As a simple illustration, let us suppose that we intend to query the PreferredSupplierPartA WS provider for its current rate of order satisfaction. Eq. 4 becomes,
$EV(s) = \int_{\langle p_1, p_2, 1-p \rangle} Pr(T'$ (**PreferredSupplierPartA** = Yes/No/Unknown |PreferredSupplierPartA, **PreferredSupplierPartA available** = Unknown)= $\langle p_1, p_2, 1-(p_1+p_2) \rangle \rangle)$ $V_{\pi^*}(s|T') d\mathbf{p}$ given that the random variable **PreferredSupplierPartA available** assumes either *Yes*, *No*, or *Unknown* on checking the status of the Preferred Supplier.

Let $V_\pi(s|T')$ be the expected cost of following the original policy, π, from state s in the context of the revised model parameter, T'. The policy, π, is optimal in the absence of revised information. We formulate the (VOC) due to the revised transition probabilities as:

$$VOC_{T'(\cdot|a,s')}(s) = \int_{\mathbf{p}} Pr(T'(\cdot|a, s') = \mathbf{p})[V_\pi(s|T') - V_{\pi^*}(s|T')]d\mathbf{p} \qquad (5)$$

The subscript to VOC, $T'(\cdot|a, s')$, denotes the revised information inducing the change. Intuitively, Eq. 5 represents how badly, on average, the original policy, π, performs in the changed environment as formalized by the MDP model with the revised T'. We may model our beliefs over the possible parameters of the WS, $(Pr(T'(\cdot|a, s') = \mathbf{p})$, in Eq. 5) using density functions, which we may obtain from the service provider using a service-level agreement (SLA). We let the densities for the WSs take the form of *Gaussian* density functions [1].

Querying for information from service providers will often be expensive. As we employ a myopic strategy for querying, we only query the service that has the greatest VOC value, and denote that value as VOC^*. Furthermore, we query for revised information only if the VOC due to the revised information in a state is greater than the query cost. More formally, we only query if $VOC^*_{T'(\cdot|a,s')}(s) > QueryCost(T'(\cdot|a, s'))$, where $T'(\cdot|a, s')$ is the distribution we want to query.

Again, we refer the reader to our earlier work [4] for more details on WSC adaptation using the VOC-driven selective querying method.

3.3 Risk-Sensitive Decision Making

The value function described in Eq. 1 maps states of the WSC to real numbers. This number is essentially a quantification of the *utility*, or "desirability", of being in a state. Our WSC environment model, defined as an MDP, uses the

[1] We emphasize that these densities are marginalizations of the more complex ones that would account for all the factors that may influence, for example, a service's rate of order satisfaction.

maximum expected value of each state to derive its policy π^*, which ultimately determines the optimal WSs to invoke in the process. More generally, we say the MDP employs the *principle of maximum expected utility* (MEU), which states that a decision-maker is rational if and only if it chooses the action that yields the highest expected utility, averaged over all the possible outcomes of the action [7].

In general, utility can be a function of other arbitrary criteria in addition to the state of the MDP. One such criterion is money accrued, or the *wealth* of the MDP. Wealth is an important parameter to consider when judging utility and intelligent decision making. An obvious assumption about utility with respect to wealth is that it is monotonically non-decreasing – that is, the higher the level of wealth one attains, the higher the utility. Risk preferences describe how increases (or decreases) in wealth affect utility.

Our model, much like many other composition models, describes a risk-neutral preference, which utilize a linear utility function with respect to wealth. Risk-sensitive preferences, however, demonstrate non-linear utility curves with respect to wealth. The most common risk-sensitive utility functions follow exponential curves. They have the following form:

$$U_{exp}(w) = \begin{cases} \gamma^w, & \gamma > 1 \\ -\gamma^w, & 0 < \gamma < 1 \end{cases} \tag{6}$$

where γ is referred to as the risk factor. If $0 < \gamma < 1$, the utility function is concave, indicating that it is risk-averse. If $\gamma > 1$, the utility function is convex, indicating that it is risk-seeking (note that if $\gamma = 1$, it is risk-neutral). γ values that vary the furthest from 1 have higher risk-sensitivity for their respective preferences. Plot samples of exponential utility functions can be shown in Fig. 2.

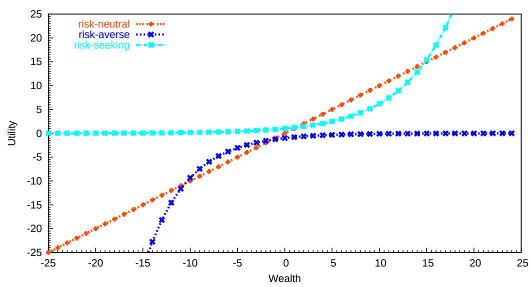

Fig. 2. Risk-seeking(light blue), risk-averse(dark blue), and risk-neutral(orange) utility functions. The risk-seeking utility function has a risk factor of 1.2 and the risk-averse utility function has a risk factor of 0.8.

The plots shown in Fig. 2 indicate two factors that are influenced by risk preferences. The first is the utility difference from one level of wealth to the next. For example, the difference in utility for the risk-neutral curve between wealth levels 0 and -5 is 5. The utility difference for the risk-averse curve between wealth levels 0 and -5 is approximately 1. The smaller difference is indicative of risk-averse attitudes – although the wealth differential is the same as risk-neutral wealth differential, the difference in utility is smaller. In other words, the gain in wealth for the risk-averse utility function has less impact on utility than the same gain in wealth for the risk-neutral. The second influential factor is that the difference in utilities across different risk preferences can be determined by the wealth level itself. Note that both the risk-averse and risk-seeking curves are steeper at differing wealth levels.

4 Selective Querying using Risk-Sensitive VOC

To incorporate a process designer's risk preferences into a composition model, it is critical that we include wealth as a component. This will allow us to use the exponential utility functions that reflect the risk preferences. As our model, and subsequently our VOC equation does not yet have a representative component for wealth, we must augment it in such a way that wealth may be included.

Augmentation of the MDP model Liu and Koenig [8, 9] have proposed a method for augmenting MDPs with wealth that use a utility function U. The model is similar to the traditional MDP in that it uses a sextuplet $\langle WP \rangle$ consisting of similar components:

$$WP = (\langle S \rangle, A, \langle T \rangle, \langle C \rangle, H, \langle s_0 \rangle)$$

$\langle\rangle$ denotes a component that has been augmented with a wealth level $w \in W$, where W is the set of all possible wealth levels that could be accumulated by the process. The augmented state, $\langle S \rangle = S \times W$, is the set of all possible states at all possible wealth levels. A is the set of possible actions (again, corresponding to WS invocations). Note that this component does not require an augmentation of wealth, as the action set remains unchanged in the transformation. $\langle T \rangle$ is the augmented transition function. It corresponds to the T function in the original model and is defined as follows:

$$\langle T \rangle(\langle s \rangle' | \langle s \rangle, a) = \{ \begin{array}{l} T(s'|s,a), \; if \quad w' = w + C(s,a,s') \\ 0, otherwise \} \end{array}$$

The augmented cost function, $\langle C \rangle$, is no longer simply the cost of invoking a service, but the *difference in utility* $(U(w') - U(w))$ between wealth level w' (reached by invoking the service) and the current wealth level w. $0 < H \leq \infty$ (i.e. the horizon) remains the period of consideration over which the plan must be optimal, and $\langle s_0 \rangle \in \langle S \rangle$ is the starting state with some initial wealth (usually assumed to be 0).

Let us focus on the difference between the C component in the traditional MDP and the $\langle C \rangle$ component in the augmented MDP. We are now taking the differences in *utilities* that is a function (U) of the wealth level w. This allows preferencs with non-linear utility functions (e.g. Fig. 2) to be used.

Analogous to the traditional MDP, we would like to obtain an optimal policy $\pi^*(\langle s \rangle)$, or $\pi^*(s, w)$, for the augmented MDP. Predictably, standard value iteration techniques cannot be applied as both s *and* w must be considered. However, Liu and Koenig [8, 9] proved that a direct correspondance exists between the policies in the traditional MDP and policies in the augmented MDP. We solve the augmented MDP as if it uses a risk-neutral preference, and when translated back to a policy in the traditional MDP it will be viewed as risk-sensitive.

This policy correspondance is exploited to create an association between the values of the augmented MDP association to values of the traditional MDP. However, because of the obvious state space explosion (needing to account for all states at all possible wealth levels), solving over the augmented MDPs with general risk-sensitive utility functions may lead to intractable computations. To alleviate this, Liu and Koenig proposed a method called functional value iteration, which finds maximum *functions* over continuous values wealth. For a more detailed analysis, we refer the reader to their work [9] and leave utilizing functional value iteration for VOC as a topic for future work.

Fortunately, Avila-Godoy [10] found that the value function for exponential utility functions may be simplified to the following equation:

$$v_{exp}(s) = \max_{a \in A} \sum_{s \in S} T(s'|a, s) \gamma^{c(s,a,s')} v_{exp}(s') \quad (7)$$

where γ is the risk factor and $c(s, a, s')$ is the cost of using service a when transitioning from state s to state s'.

Risk Sensitive VOC The VOC defined in Eq. 5 exclusively models a risk-neutral preference. To utilize VOC in such a way that it may account for all risk preferences, it must be defined in terms of both state s and wealth w. Risk-sensitive VOC is defined as follows:

$$VOC_{T'(\cdot|a,s')}(s, w) = \int_{\mathbf{p}} Pr(T'(\cdot|a, s') = \mathbf{p})[V_{\pi}(s, w|T') - V_{\pi^*}(s, w|T')]d\mathbf{p} \quad (8)$$

Now, we may apply Eq. 7 to find the two value terms, $V_{\pi}(s, w|T')$ and $V_{\pi^*}(s, w|T')$, in Eq. 8 to obtain the VOC for any risk preferences that follow exponential utility functions, which, as stated previously, may model strictly risk-averse, risk-neutral, or risk-seeking risk preferences.

The introduction of w in Eq. 8 significantly impacts the resulting values that we would obtain for VOC across different risk preferences. The values, and hence the *value difference* (i.e. $V_{\pi^*}(s, w|T')$ - $V_{\pi}(s, w|T')$), are now derived from the utility function as a function of w (as stated in the $< C >$ component of the augmented MDP model). The utilities (and utility differences) are now determined with a function using gamma, rather than a simple linear relationship.

Let us illustrate the impact of integrating w into the composition using our running scenario. As stated previously, the risk-averse manufacturer follows a policy that will likely choose the Spot Market WS to secure its parts in its supply chain, while the risk-neutral manufacturer will likely choose the Preferred Supplier WS. The risk-averse manufacturer's decision is derived from its exponential utility function over w being concave (see Fig. 2 for clarity). The potential profit (i.e. wealth gained) by using the Preferred Supplier in lieu of the Spot Market is less significant, because the *utility* increase will be less significant. Thus, when changes in the rates of order satisfaction occur, the WSC will likely be more sensitive to a change in the more stable Spot Market. If the manufacturer were risk-neutral, its utility function would be linear and find that the potential wealth increase gained by invoking the Preferred Supplier is significant. Subsequently, the risk-neutral manufacturer would be more interested in changes to the rate of order satisfaction in the Preferred Supplier.

5 Experimental Results

We demonstrate empirically that WSCs with risk-neutral preferences may adapt differently than WSCs with risk-averse preferences by using the scenario described in Section 2.

We use the parameters given in the example (Preferred Supplier WS cost of 1 and 50% availability and the Spot Market WS cost of 3 and 100% availability) and solve for optimal policy π^* prior to WSC execution. We found that the risk-averse manufacturer (with risk factor 0.6) will recommend using the Spot Market for securing part A while the policy for the risk-neutral manufacturer in that same environment will recommend the preferred supplier. We varied the availabilities (in the range of 0 to 100 percent) of both the Preferred Supplier and Spot Market WSs in the initial state s_0 (i.e. before it has invoked the first service in the WSC) and plotted the landscapes of the value differences exhibited by both available WSs for each risk preference. The value difference ($V^*_\pi(s_0)$ - $V_\pi(s_0)$), as stated previously, constitutes a major component of the VOC equation.

In Fig. 3(a), we show the value differences for a manufacturer that has risk-averse preferences. Note that the area under the curve for the Spot Market WS is greater than the area for the Preferred Supplier WS. This implies that the Spot Market WS will yield higher VOC, indicating that changes in the Spot Market have a greater impact on the WSC than changes in the Preferred Supplier WS. As a result, the Spot Market will be queried for new information (as directed by the VOC). Fig. 3(b) shows that the reverse is true (albeit to a lesser degree) for a manufacturer with a risk-neutral preference. The risk-neutral manufacturer will query the Preferred Supplier WS for new information. The plots demonstrate that that impact of changes in certain component WSs' parameters is different across different risk preferences, subsequently resulting in a different WS being queried for their revised parameters.

6 Related Work

Risk management in workflows has long received attention in economic and business enterprise research communities. Traditionally, these works identify the risks that exist in business process projects and outline options that process managers may undertake to deal with the risks [11]. Unlike the concepts discussed

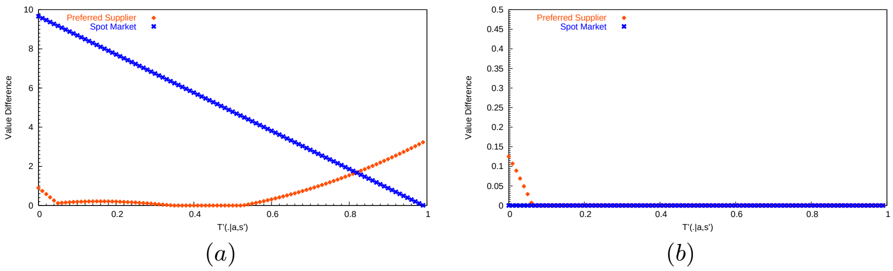

Fig. 3. Value differences obtained by a manufacturer with (a)risk-averse and (b)risk-neutral preferences for the Preferred Supplier WS (in red) and the Spot Market WS (in blue)

in this paper, their strategies implement qualitative approaches to risk management. Human specialists oversee the workflow and identify potential hazards that could upset its functionality. After the risks are identfied, they addressed in a systematic manner.

Recently, researchers have attempted to automatically quantify risk and incorporate them into their composition models. Kokash and D'Andrea [12] use traditional risk management strategies to derive contingency plans (such as QoS 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, breach of reputation, etc). These threats are juxtaposed against the possible gains of the composition. Decisions are made accordingly as to whether to utilize a contingency plan of composition based on these comparisons. Wiesemann et al. [13] incorporate the average Value at Risk (AVaR) measure, widely used in economic studies, into the decision making of a WSC. They introduce risk-preferences using the β-AVaR metric, which is defined as the mean value of the (1-β) worst losses sustained by making a particular decision. $\beta \in [0, 1]$ represents the degree a decision maker considers the worst case loss of a particular decision. When β is 0, the decision maker is risk-neutral. As β increases it becomes more pessimistic, and thus, risk-averse. The β-AVaR metric is introduced to their value maximization equations and decisions are based on the newly constructed equations. Although both apply quantitative approaches, both quantify the risks in different ways, and neither handle adapation of a WSC in the presence of these risk metrics as we do here.

Our paper borrows, in large part, from decision-theoretic planning and utility theory concepts that exist in the Artificial Intelligence literature [7]. The works of Liu and Koenig [9] as described above are the most current and notable in designing risk-sensitive agents. We use their work extensively in augmenting our WSC model, and subsequently the VOC equation, with a wealth parameter. Also helpful was the work of Avila-Godoy [10], who derived an easily computable value function for risk preferences that use exponential utility functions. These methods provide the background necessary to utilize a quantitative approach for adaptation in WSCs influenced by risk-sensitive preferences.

7 Conclusion

Despite the advent of successful approaches of WSCs adapting to volatile environments, the concept of adapting in the presence of risk preferences is overlooked. Risk preferences influence how compositions behave and thus must be considered in adaptation. We enhanced our previous adaptation strategy – querying for revised information using the VOC – to consider preferences in WSC adpatation and were able to show that WSCs that have different risk preferences are impacted by change in different ways.

Although our paper only outlines elementary ideas on adaptation while considering risk preferences, we believe there is great potential in understanding risk preferences and their effects on the adaptation problem in WSCs. A deeper understanding of the meaning of risk sensitivity may allow us to find correlations with volatility that exists in WSCs. Additionally, we would like to develop a more formal augmented model of the composition. It will allow us to consider WSCs with more complicated risk-sensitive attitudes.

References

[1] Harney, J., Doshi, P.: Adaptive web processes using value of changed information. In: ICSOC. (2006)

[2] Chafle, G., Dasgupta, K., Kumar, A., Mittal, S., Srivastava, B.: Adaptation in web service composition and execution. In: ICWS, Industry Track. (2006)

[3] Au, T.C., Kuter, U., Nau, D.S.: Web service composition with volatile information. In: ISWC. (2005) 52–66

[4] Harney, J., Doshi, P.: Selective querying for adapting web service compositions using the value of changed information. IEEE TSC **1**(3) (2009) 169–185

[5] Doshi, P., Goodwin, R., Akkiraju, R., Verma, K.: Dynamic workflow composition using markov decision processes. JWSR **2**(1) (2005) 1–17

[6] Puterman, M.L.: Markov Decision Processes. John Wiley & Sons, NY (1994)

[7] Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach (Second Edition). Prentice Hall (2003)

[8] Liu, Y., Koenig, S.: Functional value iteration for decision-theoretic planning with general utility functions. In: AAAI. (2006)

[9] : Risk-sensitive planning with one-switch utility functions: Value iteration. In: AAAI. (2005) 993–999

[10] Avila-Godoy, M.G.: Contolled Markov chains with Exponential Risk-Sensitive Criteria: Modularity, Structured Policies and Applications. PhD in mathematics, University of Arizona (1997)

[11] zur Muehlen, M., Ho, D.T.Y.: Risk management in the bpm lifecycle. In: Business Process Management Workshops. (2005) 454–466

[12] Kokash, N., D'Andrea, V.: Evaluating quality of web services: A risk-driven approach. In: Business Information Systems. (2007) 180–194

[13] Wiesemann, W., Hochreiter, R., Kuhn, D.: A stochastic programming approach for qos-aware service composition. In: CCGRID. (2008) 226–233