

# Web Service Derivatives

Thomas Meinel  
 Institute of Information Systems and  
 Management (IISM)  
 Englerstr. 14  
 Karlsruhe, Germany  
 meinel@iism.uni-karlsruhe.de

Benjamin Blau  
 Institute of Information Systems and  
 Management (IISM)  
 Englerstr. 14  
 Karlsruhe, Germany  
 blau@iism.uni-karlsruhe.de

## ABSTRACT

Web service development and usage has shifted from simple information processing services to high-value business services that are crucial to productivity and success. In order to deal with an increasing risk of unavailability or failure of mission-critical Web services we argue the need for advanced reservation of services in the form of derivatives.

The contribution of this paper is twofold: First we provide an abstract model of a market design that enables the trade of derivatives for mission-critical Web services. Our model satisfies requirements that result from service characteristics such as intangibility and the impossibility to inventor services in order to meet fluctuating demand. It comprehends principles from models of incomplete markets such as the absence of a tradeable underlying and consistent arbitrage-free derivative pricing.

Furthermore we provide an architecture for a Web service market that implements our model and describes the strategy space and interaction of market participants in the trading process of service derivatives. We compare the underlying pricing processes to existing derivative models in energy exchanges, discuss eventual shortcomings, and propose Wavelets as a preprocessing tool to analyze actual data and extract long- and short-term seasonalities.

## Categories and Subject Descriptors

H.3.5 [Online Information Services]: Web-based services; K.1 [The Computer Industry]: Markets; G.1.2 [Approximation]: Wavelets and fractals

## General Terms

Economics, Management, Theory

## Keywords

Derivatives, Incomplete Markets, Services Mashups, Wavelets, Web Services

## 1. INTRODUCTION

The traditional idea of hard wired value chains is continuously giving way to highly agile service value networks that enable flexible binding of service components within business processes and scenarios. The tremendous increase in

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2009, April 20–24, 2009, Madrid, Spain.

ACM 978-1-60558-487-4/09/04.

Web service offerings and their raising relevance for business processes has pushed disciplines like risk management in the focus of attention. Gartner analysts predict that “the new focus [...] will be on high-value business applications that either have to process high data volumes [...] or that have to process complex proprietary code and business algorithms, such as loan risk assessment”.

Workflows incorporate process steps loosely-coupled provisioned by decentralized third-party providers that are not under the control of the workflow owner. In this context, service level agreements and guarantees for required quality levels are crucial for business success. Contract management demands for innovative risk management concepts that account for special characteristics of electronic services and issues that may arise in these environments.

To illustrate these critical aspects consider a manager of a company that distributes flowers over the internet. As payment processing is not a core competency of the company, the board decides on the integration of third-party services into existing business processes in order to decrease costs of operation and maintenance. The diagram in Figure 1 sketches an excerpt of the service components of an exemplary payment process. The `PaymentProcessingService` facilitates service components from *Strike Iron*<sup>1</sup>, *Duo Share*<sup>2</sup> and *CDYNE*<sup>3</sup> to verify the customer’s address and credit card information. Customer data is stored and managed using a `StorageService` and a `DataBaseService` from third-parties. Exemplary services from decentralized storage providers are *Amazon S3*<sup>4</sup>, *Digital Bucket*<sup>5</sup> and *Box.net*<sup>6</sup>. Services for organizing and managing customer data are *Amazon Simple DB*<sup>7</sup> and *Long Jump DaaS*<sup>8</sup>. However, mentioned services represent the non-crucial part of the process as there exist multiple providers that offer service substitutes which makes them *replaceable*. The actual execution of the financial transaction through the `TransactionProcessingService` in contrary is highly critical for the success of the business process (*mission-critical*). These process steps are mostly characterized by a high degree of specialization, customization and scarcity and are therefore only offered by few service providers (cp. Section 2).

<sup>1</sup><http://strikeiron.com/>

<sup>2</sup><http://duoshare.com/>

<sup>3</sup><http://cdyne.com/>

<sup>4</sup><http://aws.amazon.com/s3/>

<sup>5</sup><http://digitalbucket.net/>

<sup>6</sup><http://box.net/>

<sup>7</sup><http://aws.amazon.com/simpledb/>

<sup>8</sup><http://longjump.com/daas/>




Figure 1: Payment Processing Service Containing Replaceable and Mission-Critical Service Components

In order to ensure the availability of these services in this oligarchy, we propose the utilization of Web service derivatives. Already established in many other business areas like stocks, commodities and energy, they have become a widely used instrument, both for risk management as well as purely investment purposes. Risk management does not only provide a hedge against resource shortage in times of (unexpected) high demand but also guarantees the delivery of the contracted services at a certain price level. The derivatives themselves are defined on the underlying dynamic price process of the respective Web service and include parameters like the date of maturity, the price function, determining the price to paid if the derivative is exercised (in case this is not mandatory), and a certain value assigned to the derivative, specifying its purchase price. It is crucial that this price is determined by following the fundamental principles of derivative pricing in order to avoid arbitrage opportunities. In our paper we analyze the first steps of how Web service derivatives need to be defined in respect to the related service characteristics and under what conditions they have to be priced.

We show certain similarities and distinctions of the derivative pricing process in stock and electricity markets and justify the need for a different approach in modeling the underlying service price process. We find that a nonparametric model fits best to the conditions given in open Web service markets and adopt the methods of Wavelet Analysis to an exemplary scenario.

## 2. REQUIREMENTS

In order analyze requirements that must be met to enable the trade of services, this section indirectly defines the service concept and differentiates it from adjacent concepts

such as goods and products through the identification of its main characteristics and their implications. In general a service is some kind of *activity* or *performance*. The result of such an activity is the *change of condition* of some person or good. This change of state is based on an agreement of the economic unit owning the good subject matter and the one providing the service [16,21].

**Definition 2.1 (Service)** *A service is an activity which an economic unit A (service provider) performs for another economic unit B (service consumer) that results in a change of state or condition of an economic unit C.*<sup>9,10</sup>

Services expose a set of unique characteristics that have strong implications from an economic perspective and allow a more or less consistent differentiation from traditional goods or products. In literature it has been argued that intangibility is the main characteristic to differentiate goods from services [33,41]. Especially in the marketing area, intangibility has been identified as the most difficult aspect of services to deal with when it comes to evaluation of service value creation as well as quality control and assurance. Focusing on economic properties and their implications for the coordination of value creation, intangibility is not the only fundamental characteristic to differentiate goods from services. The following list of service characteristics serves as a basis to derive requirements for adequate market mechanisms to coordinate value generation through services. It is neither exhaustive nor complete.

**C 2.1 (Uno-actu)** *Service production and consumption are not separable an coincide in time.*

<sup>9</sup>Economic units A, B and C are not necessarily different.

<sup>10</sup>This definition is based on [17,21]

In contrary to goods where the production, use and ownership can be separated from the economic entity itself. A service cannot be treated independently from its producer or consumer. “services involve *relationships* between producers and consumers” [22]. This implies that the process of production and consumption cannot be separated meaning that there is no producer without a consumer and the other way around (e.g. a barber can only cut hair if the customer is present at the same time which implies that there is no hair cutting activity possible without the barber or the customer being present). This principle is also called *uno-actu* which states that production coincides with consumption. This characteristic is fundamental to distinguish services from goods and it causally implicates most of the following service characteristics.

**C 2.2 (Not storable)** *Services cannot be inventoried or produced on stock.*

The main value generated by the consumption of services comes from an action or performance. Service are *ephemeral* – *transitory* and *perishable* – which implies that they cannot be stored or produced on stock. It is not possible to produce services in advance in order to meet fluctuating demand. It is of great importance to distinguish between the actual performance that leads to an immediate change in state and its effect on reality. The activity itself on the one hand cannot be produced on stock as it is intangible and perishable. The person or good that is affected by this activity on the other hand can mostly be preserved over time [17] (e.g. the actual deed of cutting hair cannot be produced on stock, whereas the change of condition – the physical cut hair – can be inventoried and exists over time).

**C 2.3 (Co-production)** *Services are often co-produced by their consumers.*

According to Definition 2.1, services are deeds or actions that change the condition of another economic unit. This economic unit – often referred to as external factor – is mostly brought in by the consumer. The consumer proactively influences the service activity and might therefore influence its result and quality. The degree of customer participation and co-production in the context of different service categories is analyzed in [3]. Depending on the type of service (*i*) customer presence might be required during service delivery, (*ii*) customer inputs might be required for the actual service creation or (*iii*) customer inputs are completely mandatory. Co-production is argued to be the main characteristic to differentiate services from goods [15]. However recent production strategies of traditional goods heavily integrate customers in the production process – often referred to as mass customization [32] – which shows that co-production does not appear to be a suitable service characteristic in order to strictly distinguish services from goods.

**C 2.4 (Intangible value creation)** *Value creation through services is dominated by intangible elements.*

Some services include physical elements in the process of value creation (i.e. spare parts during a repair process). However, the most value is created in the form of intangible, immaterial elements. The consumer of a service *experiences* the performance or activity which embodies the main portion of created value [26]. Services create value when service consumers benefit from experiencing a service without

a transfer of ownership (e.g. booking a hotel room). Due to this fact, the assessment of quality and its assurance is a critical issue in the context of services as an experience or an intangible result is hard to measure and strongly depends on the economic unit to which it is provided. A continues spectrum from tangible-dominant to intangible-dominant to differentiate between goods and services is suggested in [38].

**C 2.5 (Fuzzy inputs and outputs)** *Service inputs and outputs are fuzzy and tend to vary more widely.*

Implied by the previous characteristic, it is hardly possible to control quality aspects of a service in a way that outcomes are predictable and constant over time [18]. Services are produced and consumed coincidentally and the value that is created during this process varies widely due to the lack of control instruments and various facets of service experience. This issue is even more intensified by another phenomenon that is specific to services. The quality of a service might depend on the “quality” or effort of the service consumer (e.g. in teaching or consulting) [19]. Due to the fact that the quality or effort of a service consumer is not under the control of the provider and tends to vary from individual to individual the final outcome of a service activity is fuzzy and varies more widely.

In summary, the fact that services cannot be inventoried and that time plays an important role for the value of a service has strong implications in situations with fluctuating demand and supply. Possible ways to counteract situations when demand exceeds capacity are *advanced reservation* and *dynamic pricing strategies* [6].

The dynamic price processes we assume to be given in our work are subject to deterministic seasonal as well as stochastic influences. This process is the result of the subjacent market model which may be e.g. an (double-) auction or other heuristic or optimized allocation mechanisms. In our scenario Web service derivatives may also be defined on a static price process as advance reservation makes also sense in hedging unpredictable resource shortfalls.

Despite of general service characteristics we also focus on important differentiations from a requester’s perspective. According to [30] we distinguish services based on the context of usage into *replaceable* and *mission-critical* services. A replaceable service is a service typically provided by multiple providers. Assuming all service expose the same interfaces these types of services can easily be replaced without jeopardizing business functionality. Replaceable service are typically used in processes or applications that do not suffer from unavailability over short periods of time (e.g. weather forecast). Mission-critical services on the other hand are more crucial to business functionality. Such services are mostly provided by few, highly specialized providers. Due to the high level of customization and specialization and the importance of their availability, they can hardly be replaced or substituted. These types of services are responsible for the flawless execution and behavior of whole applications or business processes.

This leads to the need of models for advance reservation to hedge risk of technical failures and resource shortages, based on the assumption that derivative holders are treated preferentially in respect to service requesters acquiring their needs on the spot market only. Therefore, in contrast to stock markets, derivatives might also be exercised if their strike price is above the actual spot market price, in order

to ensure fulfilment of the service request. Besides these benefits for service consumers, service providers also profit from writing (selling) derivatives on their offered services. Though they put themselves in a liable position, they are able to hedge the risks they would be exposed to by solely offering their services in a spot market. As service demand can be assumed to be unpredictable (at least to a certain degree), they face not only the risk of low prices in times of low demand, but also the opportunity costs of providing services resources eventually not being used. The sale of derivatives provide an up-front fixed income equal to the derivatives value. This compensates the provider for his liabilities and opportunity costs, thus, as the latter are known, making its risk assessable. Another benefit is the ability to improve capacity planning, as the demand for derivatives may also provide an estimator for the expected demand in the spot market.

The non-storable, non-separable nature of Web services results in an incomplete market, where the services itself can be acquired and consumed, but can not be hold for replication purposes, which is the fundamental principle of complete financial markets and arbitrage-free derivative pricing. These circumstances are similar to the ones we encounter in electricity markets [36]. Thus, the risk-neutral valuation principle of derivatives, like the celebrated *Black-Scholes pricing formula* [5] or the *Black 76* model, cf. [23], can not be applied to our scenario.

To be able to define derivatives on an underlying Web service asset offered by not just a single but rather a range of multiple providers, we require these services to be classified in a homogeneous set. Furthermore, we require the derivatives on these assets to be priced in a consistent way and independently of individual preferences, in order not to allow for arbitrage opportunities.

### 3. RELATED WORK

Today's economy is currently going through a change from a product- to a service-centric economy. This trend fosters new cooperation forms in loosely-coupled configurations of legally independent firms. Companies shift from vertical integration to horizontal specialization starting in the 1990s with outsourcing formerly internal processes becoming a more and more powerful strategic option for companies throughout all industries. Companies tend to engage in networked value creation [42] which allows participants to focus on their strengths. Partners in such ecosystems can leverage the know-how and capital assets of partners, at the same time spreading risk, sharing investment cost, and retain flexibility. By re-aggregating with partners a company can broaden its range of customer attractions. Especially in complex and highly dynamic industries, forming agile business webs, is more than an attractive strategic alternative. As [40] and [39] express it, such business webs bring together mutually networked, permanently changing, legally independent actors in customer centric, mostly heterarchical organizational forms in order to create (joint) value for customers. Specialized firms co-opetively contribute modules to an overall value proposition under the presence of network externalities.

It is natural that such decentralized scenarios require well-designed economic mechanisms to organize value creation in an efficient manner. Auction designs to allocate and price single services are presented in [2, 13] that mainly focus on

advertisement and search services. An auction for trading composite services is presented in [7]. Extending such ideas to scenarios with multiple sellers and buyers fosters the first ideas of markets for Web services have been developed in [24, 25, 28]. Nevertheless, [24, 25] assume the absence of capacity constraints and the presence of complete markets. They focus on suitable bidding languages for the exchange of information objects and the matching of demand and supply.

There can be found a lot of literature about derivatives in several markets: [23] discusses many practical aspects of financial derivatives, while [4, 12] present a more technical treatment. The latter cover most mathematical aspects used in this research, also including the incomplete market approach we adapt to our Web service scenario.

Although derivatives on electricity are frequently traded since the beginning deregulation in the 1990s, research in this area still focuses to a large extent on how to model the underlying stochastic price process, calibrate it to already established markets like the European Energy Exchange (EEX) or Nord Pool (the Nordic Power Exchange), and price the common traded electricity derivatives, like Swing Options. Incompleteness in these markets is generally represented by a set of equivalent martingale measures, which is not a singleton. [14, 35] discuss different approaches of choosing the most appropriate measure from this set.

We propose in our work a Wavelet preprocessing approach for seasonality detection to support and enhance fine granular analysis and synthesis of deterministic parts of Web service usage and price processes. A detailed discussion of the mathematical foundations is outside the scope. Wavelet Analysis can be considered being similar to Fourier Analysis. While the the Fourier expansion provides a global spectral (frequency) analysis of a given signal or time series, Wavelets are basis functions localized both in time and frequency, allowing long-term seasonality detection and the division of the original data into detail signals on different scales. These characteristics prove to be useful to enhance the accuracy of models and simulations of stochastic time series of Web service usage processes, showing repeating patterns. For an introduction and mathematical details we refer to [9–11]

### 4. WEB SERVICE DERIVATIVES MARKET

In this section we specify the abstract model for a Web service derivative market by applying the fundamentals of financial mathematics of derivative trading in incomplete markets, based on the characteristics of Web services. We identify the necessary participants and propose an architecture to realize the trade of these derivatives. Further on, we discuss the necessity of a from stock and electricity markets distinct approach. This approach utilizes Wavelets as preprocessing tools to the in these fields widely used Fourier approximation, in order to analyze and extract long- and short-term seasonalities of an observed exemplary information service usage process.

#### 4.1 Abstract Model

Derivatives in financial markets are priced according to the principles of efficiency and completeness (cf. [4]). Efficiency means that all financial instruments must be priced so as to not allow for any arbitrage opportunities, while completeness implicates that every in the market traded derivative can be replicated, i. e. there exists a self-financing trad-

ing strategy, generally by forming a portfolio of assets but not the derivative itself, replicating the behavior and outcome of the derivative at every time spot. E. g. a *European call option* is one of the most commonly traded derivatives (also: contingent claims) in financial markets, giving the owner the right, but not the obligation, to buy one asset, i. e. a certain amount of stock, at a predetermined price by a specific time, the date of maturity  $T$ . Thus, the option today is worth the by the continuously compounded risk free rate discounted difference of the yet unknown stock price at  $T$  and the option's exercise price. If this difference is not positive, the option is worth zero. This today's value is a priori derived by calculating a unique purchase price, considering the stochastic price process of the underlying asset. This is the Black-Scholes model, and it can be applied, since the primary assumption, that the market is complete, holds. Thus, in principle the call option is rendered redundant, as the exact same outcome can be reproduced by forming a *portfolio* consisting of a risk free asset (like a bank account with a certain interest rate) and the stock itself, which is then continuously adapted in time, to replicate the derivative's behavior. It is important to notice, that the price is independent of individual risk attitudes, as the stochastic process is observed under a risk neutral probability measure (also: martingale measure), which in this case is unique, thus yielding also a unique price without any arbitrage opportunities. Hence, completeness is a situational fact which is obtained from the actual market, while efficiency is the objective to be achieved in pricing the derivative, regardless of the respective underlying at hand.

As Web services, once they are acquired at a spot market, in general can neither be stored, resold, nor further actively traded, it is not possible to set up any portfolio containing the respective underlying. Thus, we deal in our scenario with an incomplete market, implying that replication of derivatives generally is not feasible. As a result, any derived prices are not unique. In order to avoid any possible arbitrage opportunities, we follow the method of pricing contingent claims in incomplete markets [4], stating, that even though prices of these claims cannot be uniquely determined, they must be at least *consistent* to each other. A detailed illustration of all mathematical aspects in this section can be found there, we omit non-essential details and function arguments.

We assume that at a specific time, mission-critical Web services are sold to all service requesters through the market mechanism for the same price, i. e. we model the price as a stochastic process  $X$ , generally given by

$$dX(t) = \mu(t, X(t)) dt + \sigma(t, X(t)) dW^P(t), \quad (1)$$

where  $\mu$  is a *local deterministic drift*,  $\sigma$  the *diffusion term* for  $W^P$ , being a scalar *Wiener process* under the objective probability measure  $P$ , and time  $t$ . Furthermore, we assume the existence of a risk free asset, like a bank account, whose dynamics are given by

$$dB(t) = rB(t) dt,$$

i. e. invested currency in this account increases accordingly to the (constant) interest rate  $r$ .

Be  $\mathcal{Y}$  any simple contingent  $T$ -claim, i. e. a derivative with date of maturity  $T$ , which depends only on the outcome  $X(T)$ , and therefore is denoted by  $\mathcal{Y} = \Phi(X(T))$ , with the *contract function*  $\Phi$ . As already pointed out the price to

acquire this claim at  $t_0$  with  $t_0 < T$  is not unique, due to the lack of replication by the non-tradeable underlying Web service. However, the price of any additional claim  $\mathcal{Z} = \Gamma(X(T'))$ , depending on the same stochastic process  $X$  and arbitrary  $T'$ , must be priced in consistency to  $\mathcal{Y}$  so as not to introduce any arbitrage opportunities. We denote by  $F(t, X(t))$  and  $G(t, X(t))$  the market price processes of  $\mathcal{Y}$  and  $\mathcal{Z}$ , respectively. As these processes depend on the same single source of uncertainty  $W^P$ , we can, similar to Black-Scholes, form a riskless portfolio of these two derivatives, instead of one derivative and the respective underlying, which must return the rate of interest  $r$ . Using the *Itô formula* the price processes are given by

$$dF = \alpha_F F dt + \sigma_F F dW^P, \text{ with}$$

$$\alpha_F := \frac{\frac{\partial F}{\partial t} + \mu \frac{\partial F}{\partial x} + \frac{1}{2} \sigma^2 \frac{\partial^2 F}{\partial x^2}}{F}, \text{ and}$$

$$\sigma_F := \frac{\sigma \frac{\partial F}{\partial x}}{F}.$$

We have used a short notation and suppressed the arguments  $(t, X(t))$  for  $F$ ,  $\mu$ ,  $\sigma$ ,  $\alpha_F$  and  $\sigma_F$ , with  $x = X(t)$ . We define  $dG$ ,  $\alpha_G$  and  $\sigma_G$  analogously. In order to eliminate the Wiener process, and thus, the risk, we form a locally riskless portfolio of these two derivatives, i. e. it remains riskless only for a very short amount of time before it needs to get adjusted, which must therefore return the risk free interest rate  $r$ . Then, simple algebraic manipulation leads to

$$\frac{\alpha_F(t) - r}{\sigma_F(t)} = \frac{\alpha_G(t) - r}{\sigma_G(t)} =: \lambda(t),$$

with each side of the equation dependent only on the respective derivative, and  $\lambda$  the *market price of risk*, being equal for all derivatives of  $X$ , also known from the Capital Asset Pricing Model (CAPM) theory (cf. [34, 37]). This market price of risk already incorporates factors like expected resource outages, liquidity preferences and risk attitudes, aggregated over all market participants. Assuming  $\lambda(t)$  to be known for all  $t$ , we can express the unique price by a boundary valuation problem (for the readers convenience we omit the respective equations here) or equivalently use the *Feynman-Kač representation* to price the new derivative  $\mathcal{Z}$  in a risk neutral valuation under the martingale measure  $Q$  (cf. [20]):

$$G(t, X(t)) = e^{-r(T'-t)} E_{t, X(t)}^Q[\Gamma(X(T'))], \quad (2)$$

with the dynamics of  $X$  under  $Q$  given by

$$dX(t) = \{\mu(t, X(t)) - \lambda(t, X(t))\sigma(t, X(t))\} dt + \sigma(t, X(t)) dW^Q(t).$$

In this way, any additional derivative of the underlying price process  $X$  can be priced consistently to the already existing (traded) ones, thus, preserving the efficiency of the market.

In a complete or internally consistent market the martingale measure  $Q$  (or equivalently the market price of risk  $\lambda$ , depending whether we choose the solution of the price problem to be represented by (2) or a boundary value problem) is unique, and the assumption of this function to be known marks the crucial point in our scenario: While we empirically can observe  $X$  under  $P$  and thus, determine  $\mu$  and  $\sigma$ , this does not hold for  $\lambda$  as the function depends on the contingent claims of type  $\mathcal{Y}$ , which cannot be priced uniquely, but

are assumed to be traded in the market anyway. Therefore,  $\lambda$  needs to be observed on and derived from the market itself by applying numerical methods like least squares optimization to theoretical and observed prices. One begins with a parameterized family of functions  $\lambda(t, X(t), \beta)$ ,  $\beta \in \mathbb{R}^k$ , a number of observed actual prices  $\tilde{F}_i(t, X(t))$ ,  $i = 1, \dots, n$ , of the traded derivative  $\mathcal{Y}$ , computes the theoretical pricing functions  $F_i(t, X(t), \beta)$ , and solves the least squares minimization problem for the actual time slot  $t = t_0$  and known outcomes  $X(t_0)$

$$\min_{\beta \in \mathbb{R}^k} \left[ \sum_{i=1}^n \left\{ F_i(t_0, X(t_0), \beta) - \tilde{F}_i(t_0, X(t_0)) \right\}^2 \right]$$

in order to minimize the error between the two prices and determine the optimal parameter vector  $\beta$ .

Thus, the critical point lies in the ansatz of the chosen function template and the respective parameters. This choice, of course, is not unique and is similar to financial short rate models, with the parameter determination closely related to what is known as the inversion of the yield curve. As the estimation of optimal models and parameters is complex and requires specialized abilities, like access to market data and analysis, it seems appropriate to implement this task as a Web service itself, as is discussed in our architecture proposal.

## 4.2 Architecture

Environments in which distributed units provide functionality in a loosely-coupled manner require some sort of process or set of rules to align activities in order to generate a desired outcome, i.e. they require *coordination*. The objective of coordination is to make a set of entities – either by providing incentives or establishing constraints upon them – pursue a common goal, e.g. producing a defined outcome. Coordination can be formalized by designing adequate *mechanisms*, i.e. sets of rules that govern the interaction between the various entities. Coordination is the key instrument to organize multiple activities especially in distributed environments. In the context of Web services two specifications provide frameworks to implement coordination scenarios, WS-Coordination [27] and WS-CF [8]. This work focuses on WS-Coordination as it is a finalized standard in contrary to WS-CF, which is still a public review draft. WS-Coordination is based on concepts and roles that are represented by Web services. *Initiator*, *coordinator* and *participants* communicate using a common *context* that glues their interaction to the coordinated activity. The framework allows for different coordination protocols to be plugged in to coordinate domain-specific work between clients, services and participants.

The architecture for the acquisition of Web service derivatives involves five parties and is depicted in Figure 2. The Web service provider continuously updates his future service availability to a Web service derivative intermediary, which is responsible for the calculation of a fair derivative valuation. This task can be carried out by the market maker introduced in [29] which is responsible for running the Web service market. As soon as a prospective Web service consumer issues a request along with the required data, like the derivative type and time of maturity, the intermediary requests via another Web service a specialized derivative parameter provider. Determination of these parameters, in-

cluding the drift and volatility of the underlying price process as well as the market price of risk, may not always be an easy task and requires specific skills. In order to minimize the inevitable numerical errors in this calculation, historic and most actual market data are needed. In the case the derivative parameter provider has not the means to access this critical data, he issues another request to a market information provider.

Having received the necessary market data, the market price of risk is calculated and delivered to the derivative intermediary, who in turn determines now the unique consistent derivative value. Afterwards, he sends a notification together with a unique request ID to some chosen Web service provider enabling him to hold the reservation, and to the derivative requester itself.

In case the derivative contains only the right but not the obligation to exercise it, before or at the date of maturity the service consumer, being the holder of the derivative, notifies the respective provider, which has written (i.e. sold) the claim, of its decision.

## 4.3 Discussion

As markets grow in respect to participants as well as to traded contract amounts and volumes, the general importance of this market gains weight. Fast changing supply and demand, even fostered by electronic markets and a result from their accessibility, will drive the market further away from fixed price models, resulting in a competition among market participants, i.e. Web service providers as well as requesters, towards dynamic price processes. As this happened before in many industrial commodity markets, and also in markets trading intangible goods or alike services, e.g. electricity or Google's Adwords auction (cf. [13]), it is reasonable to assume, that matters will be the same with future Web services markets. As in these businesses the volume of involved monetary means increases, the focus on risk management becomes more important.

Derivatives are a natural instrument to hedge potential risks that all participants are exposed to. Besides economic issues like escalating prices, consumers bear the risk of non-availability of Web services, especially in environments where resources or specialized skills are rare. As this fact could lead to severe consequences being hard to model, e.g. failing business processes because of an unavailable mission-critical Web service, we restrict our point of view to economic risks, where rare Web services are represented by high prices. Contingent advance reservation via derivatives helps both service providers and consumers to hedge those risks, additionally rendering other positive effects: Though we defined derivatives on a set of equivalent Web services, in reality users may face serious adaption costs if the service acquisition happens on the spot market only. Reservation allows for long-term planning and process adaptation, diminishing these costs. This also holds for the service provider who can reduce his operational costs, considering the expected amount of requests for his planning horizon. Additionally, even in times of unexpectedly high prices, the provider maintains an advantageous position, as he can be assured that the consumer will request the service with him and not with any other competitor, thus, compensating potential losses.

In equation (1) we assumed the stochastic price process  $X(t)$  to be determined by a single source of uncertainty  $W^P$ , which is a Wiener process. With  $\mu$  and  $\sigma$  being constant

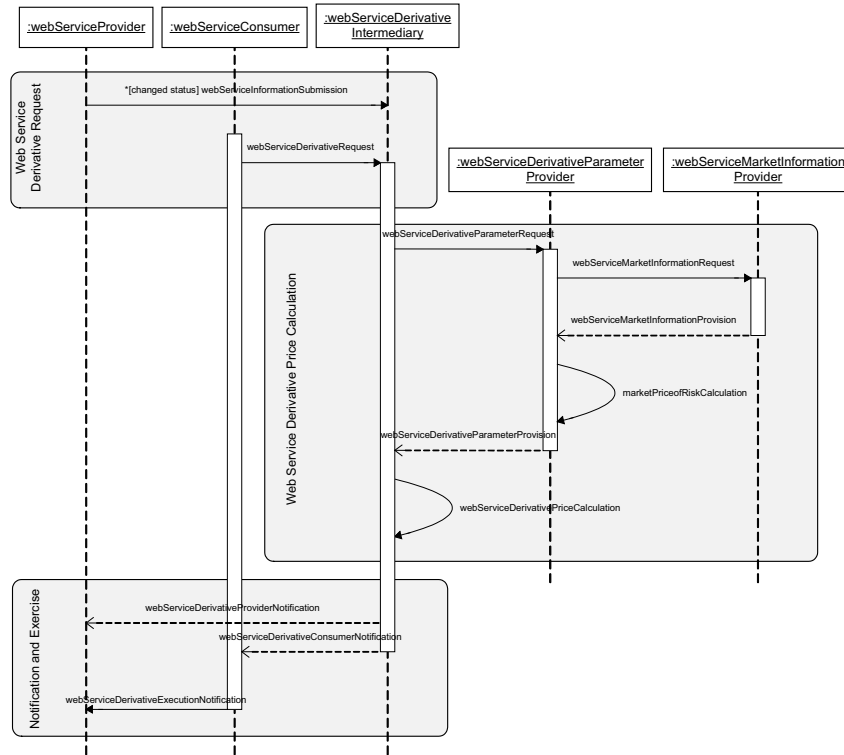


Figure 2: Web Service Derivative Acquisition UML Diagram

this yields a geometric Brownian motion. Though this is a reasonable and well accepted approximation for stock price processes, it is long known that this does not hold for other underlyings like physical commodities or currencies, where additional seasonal factors affect the process and price jumps occur. In [36] the authors discuss several stochastic processes for modeling the prices of day ahead contracts traded at the EEX, and calibrated the models to empirical market data prior to a statistical analysis. In their scenario a process, consisting of a deterministic part  $D(t)$ , describing the seasonal patterns and linear long-term trend, and a stochastic component  $S(t)$ , best describes the peculiarities like seasonalities, spikes and mean reversion observed in electricity spot prices. This stochastic part is characterized by four independent stochastic processes, namely two Wiener processes describing long- and short-term fluctuations, and two exponential Poisson processes adding jump components. Though it is not possible to establish an obvious correlation between electricity and Web service supply, demand and related price fluctuation patterns, similar characteristics let us assume that Web service price processes most likely will reveal a similar behavior including several stochastic factors.

As with this more complex price process we have more than one independent source of uncertainty, our model in Section 4.1 needs to be generalized to the multidimensional case, resulting in a market price vector of risk and a volatility matrix. This is straightforward as is shown in [4]. Since in this models we deal with probably  $m \geq 1$  stochastic processes, we also have to specify  $m$  different benchmark derivatives, prior to be able to establish a consistent Web service derivative market.

#### 4.4 Wavelet Preprocessing Analysis

As we have stressed, the underlying stochastic process  $X(t) = S(t) + D(t)$  plays a key role in determining fair derivative prices, hence most accurate modeling should be intended. In this section we focus on how to adapt and improve established methods from other fields (here: electricity markets), being the closest to our scenario, to model the deterministic part  $D(t)$ . In [36] this part of daily measured prices was assumed to be of the form

$$D(t) = \sum_{i=1}^k s_{i,1} \sin\left(\frac{2\pi kt}{365.25}\right) + s_{i,2} \cos\left(\frac{2\pi kt}{365.25}\right) + \sum_{j \in N} \mathbf{1}_j s_{j,3} + t\mu$$

with  $k$  to be chosen the first  $k$  summands of the Fourier series modeling (partial-) annual seasonalities,  $N$  being the set of week- and holidays representing (in conjunction with the indicator functions  $\mathbf{1}_j$ ) more fine granular seasonal patterns,  $\mu$  a general linear trend, and  $t$ , as usual, the time. The parameter vector  $s$  was then estimated by a least squares approximation of the original process taken over several observed periods, thus, separating the deterministic from the stochastic part.

While this method may work well for electricity markets, which are subject to local restrictions given the physical characteristics of the traded good, matters will be different in Web service markets. Though some services may be restricted to a geographical context (e.g. a weather information service, restricted to a certain area itself), this does not hold for the general case. Given a global accessibility of these services, annual, weekly or daily patterns are likely

to intensify or mitigate, depending on the degree of possible global usage. To illustrate this case we exemplarily analyzed server usage statistics of the English and German Wikipedia sites, being the most frequented ones, from April to August 2008<sup>11</sup>, shown in Figure 3(a), with the hours counted from April the 7th, 2008, 00:00 UTC. Though being a service of a very similar nature, the German usage statistics reveal stronger seasonal fluctuations due to local dependencies as e. g. around Easter. While the English site, being accessed from a much more global community, still shows a relatively sharp decline during the summer months, local events affect only a small percentage of the service requesters, and thus, have less effects on the whole usage statistics.

Thus, we are confronted with the problem of choosing a reasonable  $k$  and finding the proper set  $N$ . While for  $k = 1$  we still have the reasonable interpretation of detecting seasonalities of a one year period, an analogous argumentation for  $k > 1$  is difficult to justify, since periods of half or quarter of a year cannot be reasoned, yet these intervals are predetermined by the yet substantiable choice of the basic period. By choosing an overly large  $k$ , undesired effects from smaller scales like weeks or days will be captured by the approximation, too, while picking  $k$  too small, will result in an *oversmoothed* solution, i. e. the approximation does not capture the essential seasonalities contained in the original signal. The dilemma of choosing the best  $k$ , which does not only depend on the period but also on the not yet analyzed signal, can be tried to be captured by the extensive use of indicator functions, as depicted above. However, these functions have to be adjusted after a rigorous examination of the relevant elements in  $N$  and the respective affected user percentage separately for every Web service. As this can result in a very costly analysis (if feasible at all) simply transferring the methods employed in electricity markets seems to be inconvenient and requires an adapted different approach.

Wavelets have long become a popular tool in time series analysis [31]. In our scenario they allow for a nonparametric regression analysis instead of assuming and specifying a priori a parametric model by sinusoidal and indicator functions like the one we depicted above. In contrast to the Fourier expansion, which allows only for a global spectral analysis, Wavelets are bases functions localized in time and frequency, enabling us to analyze signals and time series locally on different scales. Roughly speaking, we are capable of separating long term (e.g. annual) seasonalities from short-term (e.g. weekly and daily) patterns without assuming a specific underlying function. This is realized through iterative smoothing on different scales: Applying a Wavelet decomposition of level  $J$  to a signal implies  $J$  times an iterative partition of the signal into a smooth part  $\tilde{S}_j$ , and a detailed part  $\tilde{D}_j$ ,  $j = 1, \dots, J$ , on a dyadic grid. At every level  $\tilde{S}_j$  (with  $\tilde{S}_0$  equal to the original signal) is decomposed into  $\tilde{S}_{j+1} = \tilde{S}_j + \tilde{D}_j$ , with  $\tilde{S}_j(t)$  being the average mean of the  $2^j$  interval centered about  $t$ . Therefore, while we still need to take care of choosing a reasonable level  $J$  in order not to include too many details in our approximation or even oversmooth it, we have to some extent direct control over the approximation procedure, as we can select the interval over which the average mean is to be calculated. E. g. if we choose in our scenario  $J = 9$ , we smooth the signal by computing the average mean at every point centered in an

interval  $2^9 = 512$  hours, which equals approximately three weeks.

Analyzing the observed signals in such a way also eliminates one possible problem one encounters when assuming a specific deterministic model: Though the parameters of the model are best fit to the data through regression methods, it cannot be ensured that the assumed model really captures all deterministic information present in the signal. By analyzing the data on different scales via Wavelet techniques one does not encounter these problems, as the approximation relies only on the observed data itself. This is illustrated in Figure 3(c) where we did not assume an explicit parameterized trend function (due to the reasons pointed out above), but compared the weekly mean and the Wavelet approximation. The approximation was derived by reconstructing the scaling coefficients obtained via a level 9 discrete Wavelet decomposition using the Daubechies-10 (D10) Wavelets (for technical details see [10]).

Since we are mainly interested in the analysis and synthesis of periodic signals, the Wavelet approximation reveals a flaw at this point, since even if we analyze a signal over several periods, the final approximation will not be periodic (as we did not even choose a period, but rather the scale of the decomposition). The idea we propose is then, to use Wavelets as a preprocessing technique prior to the Fourier expansion: Having computed  $\tilde{S}_J$ , we perform a periodic approximation, with period length  $P$ , solving the least squares problem

$$\sum_t \left[ \tilde{S}_J(t) - \sum_{i=1}^k s_{i,1} \sin\left(\frac{2\pi kt}{P}\right) + s_{i,2} \cos\left(\frac{2\pi kt}{P}\right) \right]^2 < \epsilon$$

and choose  $k$  sufficiently large, until the preselected error limit  $\epsilon$  is reached. The original data can then be adjusted by such detected deterministic seasonalities in order to extract the pure stochastic part.

Based on the assumption that Web service request time series are correlated with the their dependent business processes and their integration into human working patterns, potential weekly and daily patterns should also be separated from the remaining stochastic part(cf. [1]). Figure 3(b) depicts in detail the first three weeks of the analyzed data.

Observable daily and weekly patterns encourage to repeat the above procedure, only changing the approximation level and Wavelet base according to the desired scope of analysis. A derived seasonality for one week (D4, level 6) is illustrated in Figure 3(d). The daily case is analogous.

By removing the deterministic components from the original data we receive the stochastic part. For long-term synthesis a linear trend  $t\mu$  would have to be estimated additionally to the annual, weekly and daily patterns. The remaining stochastic signal can then be stated by a parameterized model. The initialization of the model and estimation of the parameters can be done by established methods, like *maximum likelihood estimation* or *Markov Chain Monte Carlo*, applied to the remaining processes (cf. [36]).

The analysis of usage or demand statistics instead of real price processes may not be done without questioning. However, our approach is justifiable for two reasons. First, as there is actually no observable dynamic price process of electronic services, being a focus of research itself, analyzing accessible usage statistics seems more promising than heuristically establishing correlations to observable price processes

<sup>11</sup>accessible at <http://dammit.lt/wikistats/>





