

Market Intelligence and Price Adaptation

Emad Eldeen Elakehal
The Bookdepository Ltd.
Gloucester GL2 5EB
United Kingdom
emad@bookdepository.co.uk

Julian Padget
Department of Computer Science,
University of Bath
BATH BA2 7AY
United Kingdom
jap@cs.bath.ac.uk

ABSTRACT

In the context of e-commerce, it is critical for a retailing company to be able to assess the market and respond quickly to changes in competition and/or service levels and availability of its products. If the company operates globally, and the geographical constraints on the supplier are different from those on the consumer, it is even more crucial to assess each market with respect to its characteristics and dynamically price every individual product accordingly. In this paper we report on an agent-based distributed system that was developed to support the retailing operations of one of the largest online book sellers in the UK, which has millions of books on offer and operates internationally in a number of market places. The system supports the collection of market data, processing of this data, and applying dynamically one of set of predefined pricing patterns. We describe the principles on which the price adaptation is based, the (long-running and proven successful) distributed scalable software architecture that supports it, illustrate how the mechanism copes with some typical situations system that arise and highlight some of the lessons learned from the development experience.

Keywords

Multiagent systems, marketplaces, e-commerce, real-time differential pricing

1. INTRODUCTION

In focussing on its own business processes, objectives and performance indicators, it can be easy to forget that each company is always part of a supply chain, even if most of the time the downstream element is an individual consumer. Commonly, most companies advertise their goods not only through their own direct to consumer websites, but also through third party mediators who process the transaction between the consumer and the seller company.

These third parties—also known as selling channels—marketplaces depend on suppliers with a physical position to be able to deliver the goods. Other competitors will often have

multiple presences in different geographical markets and normally in most of these marketplaces. Thus we characterize our scenario: a company that supplies books through multiple third parties in multiple markets—and thus also the question: how to work out the right price for good x on channel y in market z ?. The type of dynamic pricing described in this work is what is known as the ‘posted price mechanism’, where a good is sold at a take-it-or-leave-it price determined by the seller. Dynamic posted prices are also inter-temporal prices [6], where the seller changes prices dynamically over time, when there is change in one, some or all of the pricing factors, to achieve sales objectives such as maximise sales volume, minimise time from advertisement to sale, maximise profit or increase the company market share.

In [5], we set out the design of an agent-based system for the management of the Book Depository catalogue and stock-control—a system which has been live now since 2006. In contrast, here we focus on the next stage in the chain—selling—which in an international e-commerce context means multiple, diverse markets. Naturally, stock and selling are connected, in this case through the Listing Management Agent, so that a catalogue update affects what goods may be offered. Although the Book Depository does maintain its own websites, its primary channels are as a secondary seller through portals such as Amazon marketplaces. Consequently, this introduces two further factors into the pricing feedback loop: differentiation between geographical markets—although customers are not constrained to go to `amazon.de` just because they live in Germany, say, although it is presently more common—and responding to pricing changes by other parties with offerings on the same (selling) channel.

Thus, the main contribution of this paper is explain how a generalized linear model [12] can be embedded into a complex multi-channel e-commerce environment, to take account of a range of key performance indicators in setting and changing the price of a very large number of goods, in response to both internal (business controllable) and external (e.g. competitor controlled) factors.

The rest of the paper is laid out as follows: in section 2 we set out the properties of the various markets in which the pricing agent is operating to provide context for section 3, where we describe the various components that implement the market intelligence gathering process and both its technical and logistical integration with the rest of the business. We present results of the analysis of data from a normal – that is, not in the run-up to a holiday – one week period, to demonstrate the functionality of the system in section 4. We finish in section 5 with a discussion of related and future

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

International Conference on Electronic Commerce '12, August 6-8, 2012, Singapore Management University, Singapore.
Copyright 2012 ACM 978-1-4503-1197-7/12/08 ...\$15.00.

work and some conclusions.

2. MARKET CHARACTERISTICS

Over the last few years, Internet retailing has gained a significant share of the total retail sector. In the UK according to the ONS [7]¹ Internet average weekly value of Internet retail sales in September 2011 was £539.4 million, which was approximately 9.6 per cent of total retail sales² compared to the total value of all sales by non-financial sector businesses in 2007 which was 7.7 per cent. The Internet sales by UK businesses reached £1,8465.2M in 2009. According to [9] the most online purchased goods are books and CDs.

E-commerce systems are typically built to establish business to consumer (B2C) relationships in the form of either (i) direct-selling web sites, or (ii) marketplaces where a significant level of competition exists and results in the need to change price continuously. Both insurance and airlines industries are good examples of dynamic pricing where the same goods/services may be sold to different people for different prices, even at different times on the same day. Such practices are common in on-line selling of such products through call-centres, where different products are given different priorities (in order to meet sales targets within periods as short as a few hours), which is reflected in the dynamic (manual) updating of the call-handling logic. A more detailed description can be found in [20]. Regardless of ethical and fairness arguments, dynamic pricing has become an integral part of most successful online businesses and without this option, a company would likely suffer serious loss of market share or presence to competitors that do adopt this tactic.

Consequently, we view dynamic pricing both as a necessity for a company with global e-commerce operations and as a solution to the challenges arising from (i) the very different marketplaces, in the form of competition levels and currency fluctuations in real time, as well as (ii) the management of large numbers of goods with very fast changing stock levels.

According to Matook and Vessey [15] such marketplaces³ “are virtual, technology-enabled trading spaces that facilitate the exchange of information, goods, services, and payments among multiple buyers and sellers across companies” (see also [1, 22]), and are normally Internet-based environments. The main functions of such marketplaces are: (i) to match buyers and sellers by allowing for product determination either by feature or search, as well as allowing price discovery and comparison of all available offers (ii) to facilitate all transactions, by taking over the payment process as well as offering delivery of information, goods, or services and maintaining high trust level through reputations, and rating-like systems (iii) to offer institutional infrastructure, by setting regulation rules and monitoring/enforcing these rules on all marketplace members.

Various factors differentiate the marketplaces in general are:

1. **Different currencies:** some markets offer prices in only one currency such as **Amazon.co.uk**, where marketplace goods are offered only in GBP, while **Ama-**

¹Office for National Statistics <http://www.statistics.gov.uk>, Retail Sales Report, September 2011

²Excluding automotive fuel

³Also known as E-marketplaces, marketspaces or business-to-business electronic marketplaces

zon.de marketplace operates in EUR and **Amazon.com** in USD. In contrast **playtrade.com** offers buyers the opportunity to select their preferred currency. Given the continual change in exchange rates and fluctuations due to goods coming from different suppliers, costs can change all the time. We could easily have the situation of a good that is supplied from Germany (cost in EUR) and shipped from the UK (operations cost in GBP) to a customer in the USA who paid in USD.

2. **Seller entry/exit barriers:** some of the marketplaces are open to anyone to sell; these are known as low-entry/exit barrier marketplaces. In these markets there is typically a large number of sellers of homogeneous products, such as Amazon marketplaces in general, while other marketplaces control the number of sellers by the selection of well-known long-established sellers such as the **playtrade** marketplace or Barnes and Noble marketplace. The first type of marketplace represents a perfect competition situation where prices would be driven down all time, which leads to low profit in the long term and this situation has the potential that real or threatened entry could force economic profits to zero[8].
3. **Operation commission and postage contribution:** most marketplaces charge sellers a percentage of the total sales value. This is normally 7–15 percent, while other marketplaces charge a flat fee per item regardless of its value. This puts an overhead on low-value items. On the other hand, some marketplaces levy a flat fee on the customer for shipping and handling and part of this is passed on to the seller, while the rest is normally to cover the cost of payment transaction. Yet other marketplaces, like **playtrade** require the seller to post inclusive prices and then they do not charge buyers further fees. Flat-rate shipping fees affect the very low-value items as all sellers probably have to sell with relatively large profit margins.
4. **Pricing policy and territory regulations:** Some marketplaces set minimum and maximum price values that the seller may not exceed. At the same time, some goods cannot be priced lower than the suggested retail price by their manufacturer; a clear example of this is the fixed book price system in Germany, called “Sammelrevers”, so that all book sellers—even those who operate online—must sell German books at the same price. This policy extends to include sellers that operate outside Germany, so that they must also comply with this system when they directly sell cross-border German books. The Book Depository, therefore, must price many of its offers in the **amazon.de** marketplace exactly like the other sellers.

The above summary makes clear that there many variable factors in (e-)marketplaces, all of which must be taken into account—along with other factors to be discussed in the next section—in determining the post price of a good. Furthermore, while the above serves to identify the static factors affecting pricing, there is a significant challenge to address in actually collecting the data about those factors, which is the topic of the next section.

3. SYSTEM ARCHITECTURE

We start by outlining the system architecture and the components within it, as shown in Figure 1, that has been implemented, and now runs continuously, gathering information about the markets in which the company participates and how this is linked to (i) the inventory that was described in earlier work [5] and (ii) the delivery supply chain.

A particular novelty of the system we have developed is that both the listings manager and the catalogue manager are implemented as multiagent systems (MAS). The motivation for the adoption of agents aims to strike a balance between conventional software engineering and the potential advantages of agents: a multiagent approach helps to reduce coupling—in the software engineering sense—and so relax the constraints deriving from conventional centralized, planned, and sequential control. This is achieved through the intrinsic autonomy of agents and the distribution of system intelligence, so that the requirements for a given task are embedded in individual agents. The utilisation of a MAS approach allows for the replacement of centralized control by a distributed network of agents, where each of them can be focused on its local view of its environment. The system logic then emerges from the dynamic interaction between the agents in real-time, while the system schedule emerges from the concurrent independent decisions of the local agents on each node. The agents are able to adjust their behaviour automatically to respond to changes during the execution of their plans, which helps in optimizing the performance of the system.

The sub-system that we focus on here is responsible for deciding the price at which to offer a particular good (book) through a given channel (e.g. Amazon marketplace, Abebooks, etc.) in a certain market (Europe, USA, UK, etc.). This decision depends on upstream and downstream information. *Downstream* factors primarily mean for what price are competitors offering the same good, while *upstream* factors include the current stock situation and how soon can the good be delivered. Determination of the offer price is not predicated solely on the current market state, as it might be in a purely economic analysis, but has also to take into account practical issues such as the capacity to deliver the goods in a timely and cost-efficient manner.

The Listings Management System (LMS) controls for what price a good is offered on a given channel. The high-level requirements are:

- R1** The LMS should be able to update/add/remove book offers on all selling channels
- R2** The LMS should respond to the market change in timely manner to maintain high level of competitiveness
- R3** The LMS should respond rapidly to the stock changes either in quantity or supply route to reduce the occurrence of offers based on stock or delivery information that is no longer correct and the acceptance impossible-to-fulfil sales.
- R4** The LMS should validate the updates for book offers before publishing
- R5** The LMS should report any error occurring during any step in the update cycle

Given these requirements, we now look in detail at the key actors in the pricing sub-system: (i) The Listings Manager

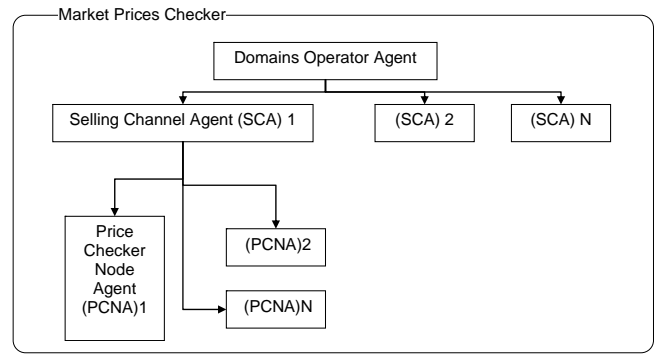


Figure 2: Market prices checker agent

Agent (LMA) (ii) Market Prices Checker Agent (iii) Stock Level Manager Agent (iv) Pricing Agent

3.1 Listings Manager Agent (LMA)

The LMA is responsible for managing The Book Depository listings (offers) over number of selling channels; all of these channels currently are online marketplaces. Listing management normally addresses three main classes of change: (i) Adding new offers to the marketplace (ii) Modifying current offers—balancing the combination of price, delivery promise and quantity (iii) Removing some offers from the marketplace.

The selling channels in our scenario are either marketplaces where large numbers of sellers are operating, or are self-run websites (by publishers). The LMA should balance the prices between these channels as well where they operate in similar geographical spaces. Thus, the LMA needs up-to-date knowledge of the following: (i) The company current offers in each separate channel (ii) Inventory status (Stock levels - Supplier Delivery Promise - Cost) (iii) Market status (Products availability - Price - Popularity of the product - Competition level) (iv) Business logic that defines the response to the overall status of market and stock measures.

The LMA in turn uses the (i) Market Prices Checker Agent (ii) Stock Level Manager Agent, and (iii) Pricing Agent to achieve its goals and we now describe each of these agents in more detail.

3.2 Market Prices Checker

The main purpose of this sub-system is to collect knowledge about the different markets in which the company operates and keep the market price and availability information current. Figure 2 presents a schematic of this sub-system, showing that in fact, there are many pricing checking agents. Furthermore, as we discuss below, the pricing checking agents themselves form a sub-system of several agents. The market price checker agent carries out the following tasks:

- **Continuously build a list of ISBNs to check their market prices** from all ISBNs that are available for sale (in stock). A list of in stock items are defined based on availability for sale from the company virtual stock as well as observing the properties of each marketplace such as selling restrictions, territory rights, censorship rules and the marketplace catalogue.

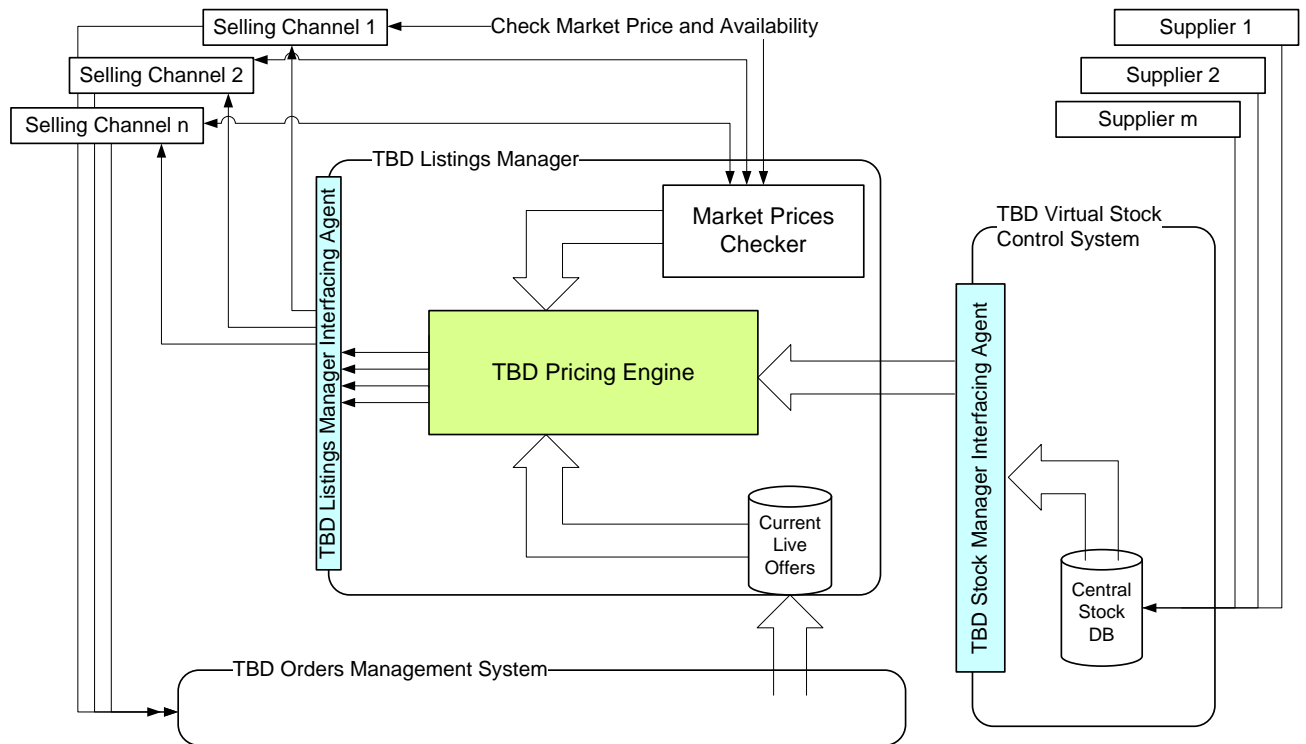


Figure 1: System Architecture

- **Prioritize the ISBNs and prepare batches of ISBNs** that are due for checking, in order to cope with the time constraints and make the best of the available checking resources.
- **Collect and filter the market prices** by checking on the other sellers' profiles, so that if the seller meets some pre-defined criteria to be considered as competitor in this specific market then its offer would be considered. This is because we are only interested in the lowest market price of that product from competitor sellers.
- **Tag each ISBN** with one of the following states (i) Unknown: we have not checked the market of this item (ii) Market status known and lowest price is known, or (iii) Market status known and the product is not available for sale in the market.

Further observables include, popularity of the product in this market and the best service level available, that is the delivery time promise from competing sellers.

The price checker sub-system includes number of proactive agents, as follows:

- **Domains Operator Agent (DOA):** The DOA prepares the full range of ISBNs that need to be checked from each selling channel over the next 24 hours, and assigns priorities for checking. This range of ISBNs will then be available to each Selling Channel Agent to handle over the specified period of time. Also the

DOA is responsible for continuously updating the market prices, using the data coming back from different SCAs. Actual updates are currently done twice a day.

- **Selling Channel Agent (SCA):** SCA is the dispatcher and task organiser for all the available checking nodes (PCNAs). It maintains a record of the number of available checking nodes, as well as their ability (how many ISBNs they can check per hour based on history or first registration) and also the status of the checking nodes, that is checking in progress, or checking done or checking on hold due to error... etc and finally the available parsers per node that would define which methods the node is able to take for checking.
- **Price Checker Node Agent (PCNA):** These are the real workers of the system, they are simple agents which have a number of methods for checking each or all channels and algorithms to parse the incoming data as well as filter these data based on the business rules.

3.3 Stock Level Manager

The stock level manager agent acts as an interface between the Stock Control System described in [5] and the pricing system described here. It provides a list of the in-stock ISBNs, their cost price and the available supply routes. The stock level manager does not have any rules on ISBN selection, it just provides a uniform query interface for other agents to access stock information.

3.4 Pricing Agent

The pricing agent is responsible for analyzing all the collected information and producing the best price for each item on each selling channel according to the pricing strategy the company would adopt at each particular pricing moment.

Traditionally, retail companies would price their goods based on one or more of the following pricing tactics: (i) **Premium pricing** where a high price is set due to the uniqueness of the product or service (ii) **Penetration pricing** where prices are set artificially low to gain market share (iii) **Economy Pricing** all costs kept to the minimum to allow for basic service/goods offered at low prices (iv) **Price skimming** charging high prices over a short period of time to benefit from a competitive advantage (v) **Psychological Pricing** that aims to get the customer to respond emotionally rather than rationally, such as the 99p mark. (vi) **Product Line Pricing** where price is set to reflect the benefits of parts of the line, such as pricing a book only for £8 and book plus CD for £10. (vii) **Optional Product Pricing** by offering more options to increase the total spent amount, such as offering choose your seat for extra charge. (viii) **Captive Product Pricing** where the main product such as a mop is sold for very low price, while the refills are priced with big profit margins. (ix) **Product Bundle Pricing** where slow moving stock is combined with fast moving goods. (x) **Promotional Pricing** such as buy one get one free, or 4 for price of 3 etc. (xi) **Geographical Pricing** this is more obvious on the internet due to the variation of shipping cost. (xii) **Value Pricing** when there are economic issues, the company might start a value line of its products to retain sales. Pricing for the internet, while largely amenable to the same approaches, does, however due to the lack of physical presence of goods and insecurity perception, face new challenges. Prices are normally motivating for consumers online; however it has been found that consumers are capable of making a trade-off between convenience and price [17]. Other factors are quite important to the consumers such as reputation of the retailer and the quality and richness of product information. Additionally it is affected by one or all of the following factors:

- It is easier for internet businesses to obtain access to pricing information in straightforward formats that allows for quick comparison. Also the ease of obtaining more information about consumers means internet retailers can tailor the price online to suit the consumer profile.
- It is a very rapid process to change prices across the whole range of products, with very low cost in both advertising as well as changes. It has been found that menu costs are lower in internet markets [4], for example.
- The customers have greater access to almost all online stores and there are many price comparison services that would help the customer to compare and choose. The lower search costs lead to lower profits and the cost of search is normally matched against the benefit from any improvement of the price [19].
- The scope of pricing is widened to be international rather than local or national, with all international sellers became competitors, so a customer in Australia

may have access to goods produced and sold in US with huge price reductions over his local seller.

The internet seller however can limit its pricing tactics to only one of few based on the selling channel, thanks to externally defined business environment constraints. For example, even if the seller is able to source and sell a German book cheaper than its regular retail price, the seller would not be able to use Penetration Pricing Tactics in `amazon.de` due to the fixed price regulations, however he can do that on `amazon.com`. Another example: if the channel states that the seller should ship the sold product within 48 hours, then the seller would not be able to offer a basic service in the form of slower shipping options i.e. 2 weeks delivery time for lower price.

Due to the sensitivity of the business information in this area we do not include specific rules that are applied or discuss in great details the Book Depository's pricing policies of each domain or product, however the prices produced would normally follow one of, but not limited to, the following patterns: (i) Match the market price (ii) Ignore the market price and sell at the cost price (iii) Ignore the market price and sell at lowest possible price (iv) Track the market price with a specified window that can be positive (overpricing) or negative (undercutting).

The pricing agent uses a simplified version of the statistical technique of Generalised Linear Models (GLM) [12], where the selling price of each good can be tailored for each selling channel based on a number of variables. The GLM regression method assumes that the response variable is a linear combination of the explanatory variables and is normally distributed.

The inverse function of a linear model that specifies the (linear) relationship between a dependent (or response) variable Y , and a set of predictor variables (X_i):

$$Y = b_0 + b_1 X_1 + b_2 X_2 + \dots + b_k X_k$$

A pricing agent could use any number of variables representing the significant variables of the current situation (Y) of the target selling channel, for example the competition level, availability of the product, shipping cost and time, ...etc. The X_i s define which response variables are connected to which elements of Y . Solving the system for b , matches the product with one of the pricing policies set by business and hence one of the pricing patterns is applied. The current model, however, does require collecting all market data for the full set of products, which is time consuming process. An enhancement could be done by adopting a similar approach to the one presented by Bouveyron and Jacques [2].

Adopting one of these patterns does not mean necessarily a profit for the seller. The seller can make a profit even if the selling price is the lowest possible price, but equally it can be making a loss, even while selling at a price that is higher than the market price.

3.5 Implementation and Deployment

We have adopted the Agentscape middleware [16] for system implementation because of its support for: (i) scalability: we needed to be able to expand the system by adding as many nodes as required, without the need to change the overall architecture; this is straightforward in Agentscape because new hosts (physical nodes) can be combined with existing locations (logical nodes) at will, and (ii) mobile agents: so the system is ready for any possible change that

Row	Description	Amazon UK	Amazon US	Amazon DE	Amazon FR
1	Total number of records uploaded	5221927	5724741	5174090	5682187
2	Number of items with no price change	3824156	5272266	3461411	3893557
3	Remaining items	1397771	452475	1712679	1788630
4	Number of items deleted due to stock level	74001	92007	135060	248989
5	Remaining items	1323770	360468	1577619	1539641
6	Items with unknown market price	2	99453	27845	40102
7	Remaining items	1323768	261015	1549774	1499539
8	Number of items unavailable from other competitors (Out of market)	6953	5630	8767	7549
9	Remaining data	1316815	255385	1541007	1491990
10	Items with price INCREASE over 10%	68183	10704	6608	12246
11	Items with price DECREASE over 10%	14791	14579	19013	53318
12	Items with price INCREASE over 5%	12865	7000	3868	4645
13	Items with price DECREASE over 5%	11165	12914	24407	14129

Table 1: Selling Domains: Amazon UK, Amazon US, Amazon DE Amazon FR

might require migrating some of the agents between the physical hosts. In addition, we are familiar with Agentscape from the implementation of the catalogue manager, so its continued use makes sense from a practical point of view. The price-checking system is deployed on 12 machines (at the time of writing), and is able to check the prices of millions of ISBNs everyday and generates 19 listings updates per day to support 15 different selling channels. A further point regarding distribution of the computation is that Amazon limits the number of queries per minute per IP address, thus for the Book Depository to be able to process its whole catalogue, the task is in effect striped across multiple servers in order to be able to deliver the necessary throughput in a given 24hr period. This aspect is, again, easily delivered as a consequence of the design of Agentscape, but also supported by the level of abstraction afforded by taking an agent-oriented approach.

4. OUTPUTS

The purpose of this section is summarize the main (business) results achieved through the design and deployment of this architecture. These are:

1. The capacity to process pricing requests over multiple channels for the complete Book Depository catalogue
2. The means to apply different pricing strategies for different goods across a variety of channels
3. The facility to respond rapidly to price changes from competing sellers in those channels in an intelligent way that observes appropriate business rules

The price checker system has been in operation now since mid-2008 and can be regarded as a mature piece of business software that delivers effective function day after day. We now look at the results in a little more detail.

The Book Depository has wide range of titles in its catalogue, and the volume of sales data through which to sift is quite significant. To put this in concrete terms, the catalogue holds a total of 17,940,335 unique ISBNs, of which 10,833,410 are available in stock to sell. The Book Depository sells an average of 30,000 books daily, and this number reaches 60,000 in the period running up to Christmas. In

Table 1 we show a summary of the raw data for 4 different selling domains, `Amazon.co.uk`, `amazon.com`, `amazon.de`, and `amazon.fr`, we then highlight some cases that have been picked out to illustrate aspects of the price tracking behaviour.

Over one week the total number of offers exceeded 5 million titles on each selling channel. During that week we can calculate (from rows 1 and 2) that 27 percent of UK listings, 8 percent of US listings, 33 percent of DE listings and 31.4 percent of FR listings changed in price; this indicates that Amazon.com (US) marketplace is more stable in terms of competition dynamics. Checking the price change level (row 10) we can see that a small proportion of the listings have had significant price increase of over 10 percent, in particular on the DE and FR listings, where only 0.42 and 0.82 per cent respectively, increased by over 10 percent of their original listing price on the beginning of that week. The same has happened concerning price decreases (row 11): we can likewise observe only few listings decreasing in price by 10 per cent across all the selling channels considered here, except in the US listings where 5.7 per cent of remaining listings decreased.

As mentioned in section 3.4, the pricing agent follows one of several identified patterns, so we now focus on four examples that illustrate these behavioural patterns. In example 1 shown in Figure 4(a) the prices reached are tracking the market price by undercutting other competitors' announced price by 1 penny, it is obvious that other competitors are then cutting The Book Depository price by 1 penny in return (time steps 5–6 and 10–11). In the long run, the price would keep dropping to reach zero profit value, or even less than cost price. This is a typical representation of the perfect competition market as described in section 2.

Example 2 in Figure 4(b) is an anomalous case of aggressive overpricing by another competitor, that drags the price up to 19 times that of the TBD price. Adaptation initially leads to a wrong pricing, but then a correction function is applied to restore a price that is correct according to the company's business rules.

Two other examples are shown in Figure 5(c) and (d), where the pricing tool ignores the market prices, so that in (c) the offer price is higher than the market price and does not change even if the item becomes unavailable from other

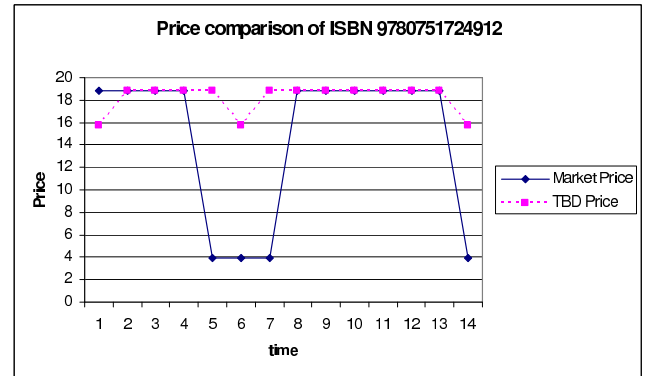
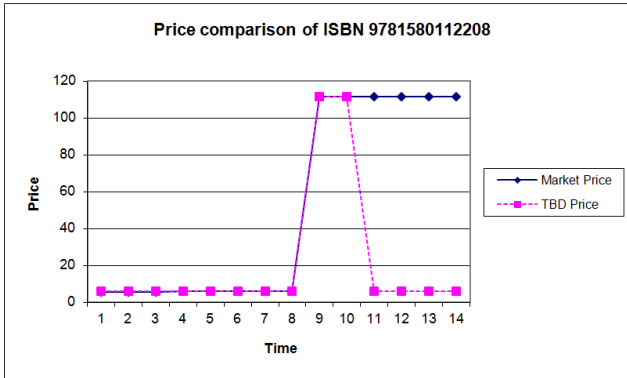
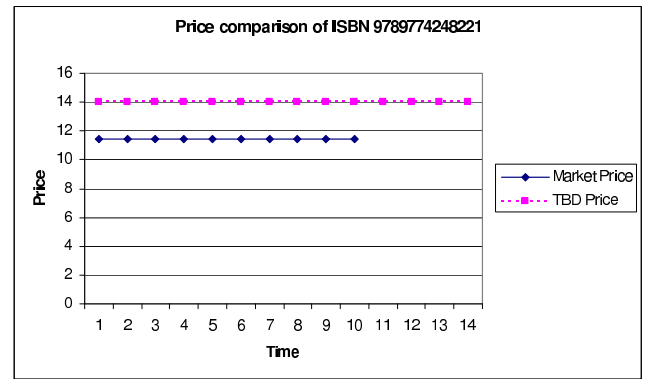


Figure 3: (a) An example of tracking market price with fixed undercutting window of 1p, (b) Static tracking can lead to wrong pricing, but a correction applied to restore the optimal price

Figure 4: (c) Tracking market price with a positive window (d) Independent pricing that is ranging between cost prices from different suppliers

competitors later on that week (time period 10 onwards). While in (d) the tool ignores the market price and the change in price results from a change in the supply route, leading to a change in cost.

5. DISCUSSION

At its simplest, the problem has similarities with the case of ZI [11] and ZIP [3] traders, where the channel provides the market environment from where a trader can collect data about current shouts and subsequently use this information to adapt its shout. Since there is no price negotiation as in the continuous double auction of ZI(P), the situation is slightly simplified in this respect, but not altered fundamentally: it just means that the shout price may also be the sale price.

Product price automation is an important issue for businesses today, especially for those operating in the electronic space, thus there has been much research in this area over the past few years. However, like much multi-agent work, most of these studies tended to focus on marketplace modelling, such as the early work of Kutschinski et al. [14] on using reinforcement learning techniques to improve dynamic pricing agents' performance, where they show the effect on the marketplaces and the effect of different adaptive pricing strategies. Closer to the seller's interest, is the work of [21], where they address the problem by presenting an algorithm based on the similarity of agents strategies in coordination

games. Other research has focused on simulating buyer-seller behaviour in the context of dynamic pricing in marketplaces such as [13]. It is less easy to find work reporting on actual implemented systems, apart from a presentation of system architecture [10] that was aimed at supermarket chains, and [18] where they introduced argumentation-based approach for automating the decision making for product pricing. This latter is a proof of concept rather than a functioning system and in addition, it presents a solution for automated decision making in very simple scenario where there is a static decision can be made over the whole range of products regardless of the fine differences between the products' market conditions and the competition level for each individual product. And above all, there is apparently no consideration of how the application can cope with the needs of a large company that has millions of offers on large number of markets, which is what differentiates the work presented here.

Dynamic pricing is a complex task, that requires the company to know not only its own operating costs and availability of supply but also the product status in each marketplace, the market behaviour, how much the current customer values the product and what future demand will be. In this paper we have presented part of a fully functional system (in continuous use since July 2008) that enables the company to dynamically price a huge range of products, and to respond quickly to changes in any of the pricing factors. The

