Copenhagen Business School

MSc., Advanced Economics and Finance
Master's Thesis

Stress Testing the UPP - An Empirical Analysis

Jimmy Gårdebrink
&
Scott Egan

supervised by
Prof. Cédric Schneider

Pages: 77
Characters: 105,945

September 14, 2016

Abstract

This paper uses data from the United States airline industry to examine the efficacy of the upward pricing pressure (UPP) measurement in predicting price changes after a merger. The predictions of the UPP are directly compared to those of a logit merger simulation and the actual, ex-post price changes that occured. Our results indicate that when used under a particular set of assumptions, the UPP produces results similar to those of a merger simulation. Thus we find that the UPP is a reasonable alternative to more intensive structural estimations and simulations. In examining the observed price changes and comparing them to both the merger simulation and the UPP, our results indicate that neither tool has strong predictive accuracy, however we cannot draw further conclusions about the usefulness of such analytical tools.

Contents

List of Figures

List of Tables

1 Introduction

Practicing antitrust economists at competition authorities and other institutions have entered a new phase when it comes to merger analysis. The focus has shifted from looking at measures of market concentration to predicting price effects, and some might even say that they are steering into a new paradigm. The traditional approach in merger analysis was previously fundamentally dependent on defining markets in order to analyze market shares and market concentration. Over the last ten years, new tools have been derived and made available that offer pragmatic ways of analyzing the anticompetitive effects of mergers. This new emphasis is based on the unilateral pricing effects of mergers, with one of the most talked about new tools in the antitrust economists toolbox being the Upward Pricing Pressure (UPP) index.

The UPP was first proposed as a screening tool by Farrell and Shapiro (2010) in their influential paper, "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition". The default version of the UPP, as introduced by Farrell and Shapiro, identifies the incentive for firms to raise prices post-merger. It does this by comparing the incentive to raise prices from lost competition to the incentive to decrease prices due to cost efficiency gains. The UPP can therefore give both positive values (incentive to raise price) or negative value (incentive to decrease price).

In 2010, the United States Department of Justice and the Federal Trade Commission released the new Horizontal Merger Guidelines[1], and the Competition Commission and the Office of Fair Trading in the U.K. revised their Merger Assessment Guidelines[2], with the UPP included in both of these revised procedural guides. Competition authorities, consultancies and law firms are now also widely using the new screening tool. This is of course of great importance for business and authorities alike. However, even though

[1] https://www.ftc.gov/sites/default/files/attachments/merger-review/100819hmg.pdf

[2] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/284449/OFT1254.pdf

the literature regarding the UPP index is growing, there are many questions yet to be answered by economic research.

One of the main motivations for the use of UPP is that it overcomes some of the major issues with the old approaches for merger analysis. Firstly, the old approaches commonly led to the final decision being based on a simple measure of market concentration, such as the Herfindahl index. This is problematic since the overall welfare in a market can actually increase with concentration in some cases (Farrell and Shapiro, 1990). Secondly, if the market structure is endogenous one cannot infer the change in conduct and performance based on the existent structure. Thirdly, the definition of what products are included in the market is always somewhat subjective or even arbitrary. Theoretically, these approaches are based on cournot competition with homogenous goods. On the contrary, the UPP is usually based on Bertrand competition with differentiated products, which in many cases is a more appropriate assumption. However, the UPP is flexible as well in that it can accomodate various assumptions about market structure and can incorporate a wide range of products.

The UPP also has advantages over other tools of measuring unilateral pricing effects of mergers. It is argued that, compared to the more sophisticated merger simulations, the UPP offers a practical and simple tool that is easy to understand and that does not require the same burdensome amount of data. Nevertheless, the UPP also comes with some limitations. It has been criticised because the incentive to raise prices does not necessarily say anything about the magnitude of the price increase and therefore does not have the direct link to measures of consumer harm. The economics literature has recently investigated the ability of the UPP, in a purely theoretical setting, to predict price increases. There is, so far, indications that the UPP is robust to misspecifications, a downside of the more complex merger simulations (Miller et al., 2016). It is possible that the UPP is not only related to the magnitude of price increase (the stronger the incentive, the higher the price increase), but that it can serve in fact as a robust measure of it.

This thesis aims to compare price predictions based on the UPP to an alternative merger simulation. We have chosen to use a simple logit model for the comparison with data from a merger in the airline industry. The airline industry has been widely researched for merger analysis, which makes the results comparable and easy to connect to previous literature. We use publicly available data to get prices and market shares, while margins are estimated using a structural demand estimation. The logit model serves the purpose of being a viable alternative to more complex models of demand, since it is a simple, workhorse model for merger simulations with an abundance of documentation and programs available to the public. It has proven to also be robust to theoretical misspecifications and most importantly, it does not require much more detailed data than does the UPP (Miller et al., 2016). Our research directly relates to the debate about whether the UPP serves as a predictor of price change, as the simple merger simulation is assumed to do.

Furthermore, we look at the literature of ex-post evaluation of merger simulations in order to investigate the meaningfulness of comparing one tool to another. The models are, in the end, pre-merger predictions of what is likely to happen as a result of the merger. It is therefore of great importance that the research that will set the path for antitrust economists in the future connects also with empirical analysis of which tools give the most valuable information. We make a comparison of our predictions from the two models to actual data ex-post and highlight the difficulties of evaluating the tools for unilateral price effects.

Our results indicate that the UPP does indeed serve a purpose as a predictor of the magnitudes of price increase, at least to the same extent as the simple logit merger simulation does. Under the assumption of full pass-through rates, the UPP predictions of price changes are nearly identical to those of the logit merger simulation. This illustrates that if the information that goes into the UPP is from the same source as the data used for the merger simulation it produces similar predictions. However, we cannot draw strong conclusions about whether the two methods are equally good or equally bad at predicting the outcome. On the one hand, we analyze only one merger; on the other, we have over 250 markets where the two merging firms overlap. The results does not indicate that the tools predicts the post-merger outcome well. However, we are reluctant to draw bold conclusions about the appropriateness of these tools based on the ex-post analysis. Rather, this section highlights the need for more empirical evaluation of merger tools in order to not only rely on theoretical models and assumptions.

2 Background and Literature Review

2.1 Airline Industry

We have chosen to analyze the United States airline industry for several reasons. First, there is a wealth of publically available data published by the US Department of Transportation, which collects information on, among other things, itineraries, prices, and fare class. This surplus of data is matched by an abundance of mergers within the market. In the past 11 years alone there have been 13 mergers between commercial US airlines of varying size[3]. We will very briefly describe the industry and its operating environment before diving into the merger literature.

After the Airline Deregulation Act was passed in 1978, free entry, price setting, and capacity setting took hold in the sector for the first time and since then the organization of the industry has been extremely dynamic. Iconic legacy carriers such as PanAm Airways and TransWorld Airlines (TWA) have disappeared, and the hub-and-spoke

[3] *Source:* Zephyr Database, published by Bureau van Dijk. Accessed September 13, 2016

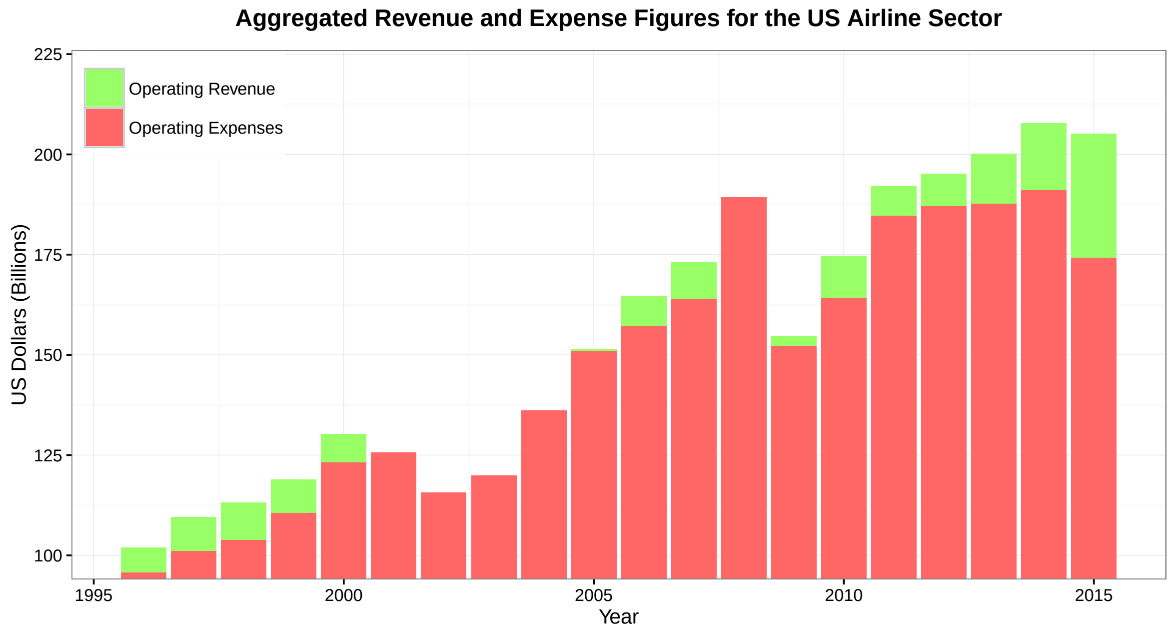

Figure 1: Aggregated operating revenue and operating expenses of the airline sector reveals an industry struggling to stay profitable [a].

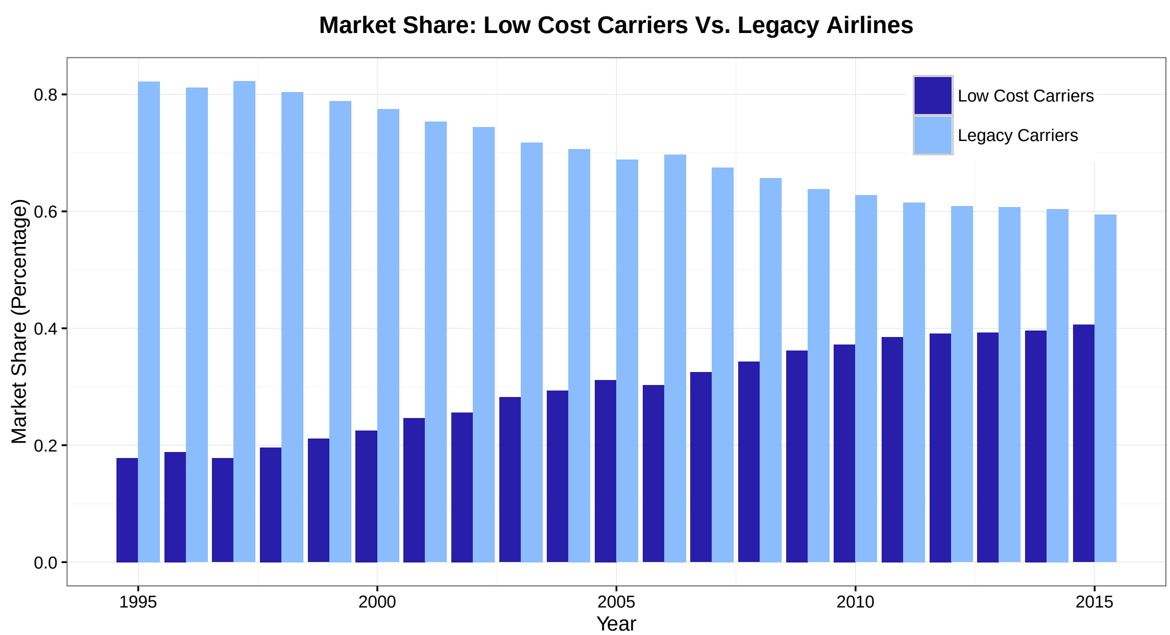

Figure 2: Low-cost carriers are capturing more market share than ever at the expense of the established players[b].

[a] *Source:* http://web.mit.edu/airlinedata/www/default.html

[b] *Source:* http://web.mit.edu/airlinedata/www/default.html

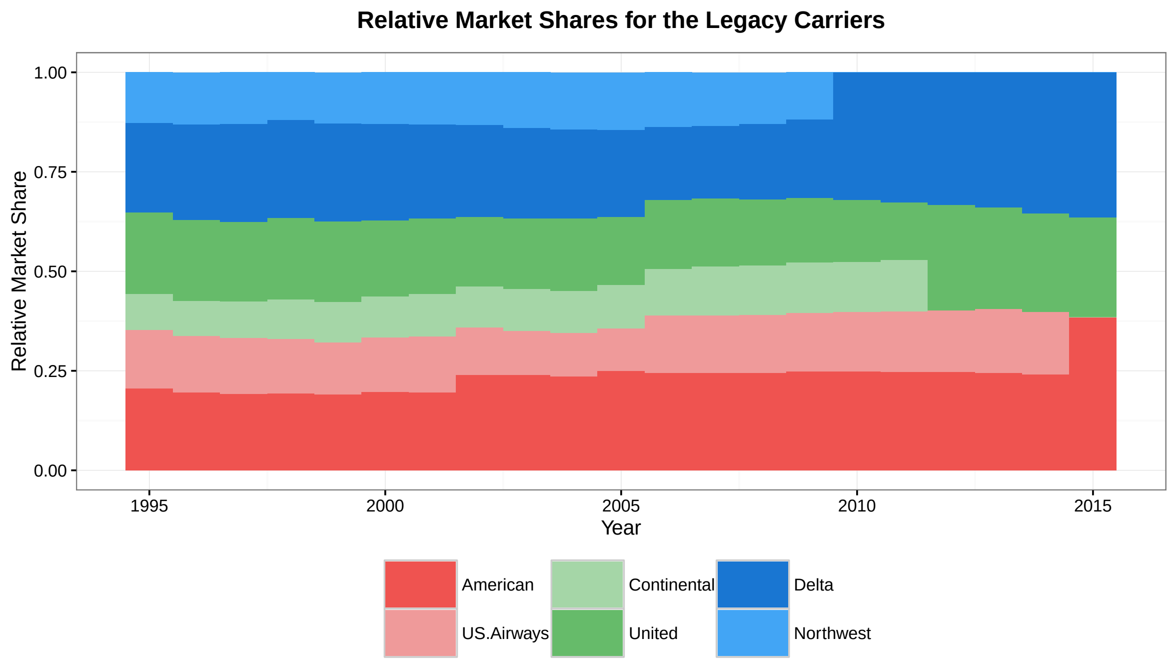

Figure 3: The 21st century has so far brought about major consolidation amongst some of the oldest American air travel firms[a].

[a]*Source:* http://web.mit.edu/airlinedata/www/default.html

system of route organization has emerged as the most prevalent strategy for getting consumers where they want to be. More recently the industry has seen numerous entries by low-cost carriers such as Jet Blue and Virgin America, and a sustained rise in fuel costs. These factors, coupled with the tremendous negative impact of the September 11th terrorist attacks in 2001, forced further consolidation on the industry to the point where today, only 3 of the original national legacy carriers exist.

Overall, the airline transportation industry in the United States has never had the reputation of being the most profitable, healthy industry. Despite overall growth of the sector, as well as historically high revenues, profitability throughout the past two decades has remained low, with the exception of 2015 (where fuel costs dramatically decreased). Such an environment made for plenty of opportunitites for airlines to merge operations in order to survive or be acquired at great values. One such case exists in the merger of America West Airlines and US Airways in 2005. At the time of the merger, they were America's 8th and 6th largest carriers by enplaned passengers, with the resulting combined firm operating as the 5th largest carrier by the same measure in 2006[4]. The merger between these airlines made economic sense, with both firms losing market share and struggling to remain competitive within their respective price class. The airlines also operated over fairly different geographic areas; America West maintained its fleet and focused its operations throughout (not surprisingly) the western half of the United States and US Airways had a strong presence in the northeast of the US (although they did operate as a national carrier).

2.2 Background of Merger Analysis

The role of the competition authority is to review mergers in order to draw conclusions about the overall welfare effects of a merger and block proposed mergers if they are predicted to be harmful. In practice this often means that a merger should be blocked if it would increase prices in the market. The majority of proposed mergers are allowed to pass without any remedies. The harmful effects that can arise when firms consolidate are important not only to competition authorities but also to the competing firms and consumers. In the years between 1988 through 2005, over 46,000 proposed mergers were filed and handed to the Federal Trade Commission and Department of Justice in the US in accordance with the Hart-Scott-Rodino Act. On average, 3.8% of these intended mergers were forced to submit more information to the authority because of concerns of lessening of competition each year from 1988-2005 (Weinberg, 2008).

In practice, the approach used by authorites when assessing mergers is based on economic models and is therefore closely linked to microeconomic theory (Werden, 2009). There are two main types of harmful effects: unilateral effects, and coordinated

[4] http://web.mit.edu/airlinedata/www/Traffic&Capacity.html

effects. The latter comes from the general belief that in a more concentrated market, firms are more likely to engage in and be able to sustain tacit collusion (Ivaldi et al., 2007). The former - unilateral effects - arise from the individual firm's incentive to raise prices after the merger imposes a new market structure (i.e. n firms now compete as $n-1$ firms) (Ivaldi et al., 2003). This thesis concerns the unilateral effects from a merger.

A merger between two (or more) firms can generate harmful unilateral price effects by relaxing a competitive constraint that, before the merger, each firm imposed on one another. When assessing and analysing unilateral effects from a merger, the most common way for antitrust economists is to look at this dynamic. The opportunity to profitably raise prices above cost depends on to which extent the consumers in the market consider products as substitutable. In the case where the products of the merging firms are close substitutes while the products of the competing non-merging firms in the same market are less close substitutes, the merging firms may have incentives to raise prices, as consumers find it more difficult to switch away from the merged entity's goods. Therefore the risk of adverse unilateral effects are typically considered to be higher in industries with differentiated products. With this in mind, firms often try to differentiate themselves and their product to their competitors because it gives them market power by lessening consumers willingness to switch to another product when prices are raised above marginal cost. After a merger, the loss of direct competition for sellers of differentiated products can be defined in two complementary ways. First, the merger allows one firms to recapture the loss in profit that previously would have been lost to the competing firm in the event of an increase in price. And secondly, it can also be the case that the merger relaxes the competitive constraint of the firm that is the response of a significant rival.

In markets with differentiated products, much of the foundation of antitrust analysis is built around the assumption of the Bertrand pricing game. In short, it means that firms are assumed to compete on prices. The game, as we consider it in this paper, assumes that firms can produce multiple differentiated products with constant marginal cost. Firms set their prices (instead of fixing a quantity) in order to maximize their profits. Between all products, prices acts as strategic complements. This means that if one firm increases the price of one of their own products, other products in the market face a higher demand (as some customers switch due to the price increase), which in turns gives the incentive for other firms to raise prices. Thus, firms face a trade-off between higher prices for each unit sold and the resulting drop in quantity demanded when it raises prices. The level of these losses in sales discourages firms from raising prices further.

The merger simulation is directly modelled by assuming that the two merging firms now operate as one firm with a new set of products to profit-maximize over under common ownership. If products are substitutes, by design, some of these lost sales will be recaptured because they now go to products that previously belonged to a competitor but now are part of the common ownership portfolio. The Bertrand model, at least for a set of common demand systems, predicts therefore that if there are no efficiency gains accounted for the prices of the merged firm's products will increase and prices of outstanding products in the market will at least not decrease. Of course, it could be that the merging entity restructures its product portfolio post-merger or uses other pricing strategies, such as price discrimination, to alter the strategic landscape. This is not investigated in this version of the Bertrand game and is not accounted for in the simulations.

Mathematically, the Bertrand game is set up in the following way: in a market with K total firms, each firm k chooses the profit maximizing price vector $[p_j]$, where each element represents the price for one of the firm's J total products. Thus, firm k solves

$$\max_{[p_j]_{j=1}^{J_k}} \pi_k = \max_{[p_j]_{j=1}^{J_k}} M_t \sum_{j=1}^{J_k} s_j(p)(p_j - c_j) \tag{1}$$

where M_t is the size of market t, $s_j(p)$ is the endogenous market share for product j, and c_j is product j's marginal cost.

The optimal pricing vector $[p_j^*]$ is found by differentiating the profit function with respect to price in order to obtain the first order condition:

$$s_j(p) + \sum_{j=1}^{J_k} (p_j - cj) \frac{\partial s_k(p)}{\partial p_j} = 0 \tag{2}$$

Whether a market is better characterized by Bertrand competition or by the main alternative of Cournot competition, (competing on quantities rather than prices) has to do with how products are produced and the flexibility of changing quantities sold in the market. Under cournot competition quantities are driving prices, not vice versa. This can be intuitively thought of as firms first deciding on the quantity produced, having to stick to that decision, and then prices will adapt to ensure market equilibrium. The airline industry is often modelled as Bertrand competition but there are reasons to at least consider why this is a good assumption (See for example Berry, Carnall and Spiller (1996) and Peters (2003)). When demand for air travelling in the US is relatively low the airlines are more likely to have many empty seats on each flight and prices are likely to be low. In such a situation where there is spare capacity firms are more likely to compete on prices according to Bertrand. If one airline unilaterally raises its price for a specific route, some customers will be lost to other options, some to another airline and perhaps some would choose a different means of travelling or choose to not travel at all. If demand is very high on the other hand and all flights are almost full the quantity sold (i.e. the number of passengers) is constrained by the number of seats, the number of aircraft in operation, and the number of landing slots at the airports. Such a scenario may be more likely to resemble Cournot competition. If one flight is taken out because the aircraft is more profitable on a different route, the customers on the original route will try to get on another flight, which is also likely to face capacity constraints. Thus, the firms will have an incentive to raise prices on the remaining routes. Unilateral analysis using both merger simulations and other tools such as the UPP can still be used for other industry structures than Bertrand (Moresi, 2009). However, in many cases, price competition á la Bertrand is assumed as a default.

2.3 Merger Analysis Before the UPP

Horizontal merger analysis has been in one way or another a government responsibility in the United States for over 100 years. The Clayton Antitrust Act of 1914 was among the first efforts by authorities in the United States to explicitly prohibit mergers and acquisitions that would substantially lessen competition, a mandate which would influence merger policy throughout the western world for nearly a century Mueller (1996). More recently though, competition policy in the United States has been dictated by the Horizontal Merger Guidelines as put forth by the United States Department of Justice (DOJ) and the Federal Trade Commission (FTC). Before 1992, focus was centered on empirical calculations of market share and market scope, and whether a given merger would create a market which was too highly concentrated; the guidelines were created under the presumption that there was a high correlation between market concentration and poor market performance. The 1962 Philadelphia National Bank case is on early example of a concentration based measure being used to set a legal precedent; a merger between Philadelphia's second and third largest bank was blocked by the Supreme Court on the grounds of the high market share the merged entity would have enjoyed[5].

[5] `https://www.law.cornell.edu/supremecourt/text/374/321`

The Supreme Court ruled that the proposed merger would have resulted in one firm controlling an "undue" proportion of the market share and was "inherantly likely to lessen competition substantially"(1963). Such a judgment was largely based on analysis conducted with the Herfindahl-Hirschman Index (HHI), a tool used to measure a particular markets concentration as the squared market shares of each firm. In market t with K firms

$$HHI_t = \sum_{k=1}^{K} s_k^2 \tag{3}$$

where s_k is the market share for firm k. Problems with such a practice and the corresponding measure exists though, as the use of merely one single measure to describe predicted merger consequences fails to account for other relevant economic factors. Particularly, cost synergies resulting from a merger are not accounted for, as mergers often create efficiencies that would be difficult or impossible to achieve within pre-merger competitive boundaries. Nor are there allowances for spectrums of differentiated products; a product or service is either within the relevant defined market or out-of-market, regardless of what an actual demand substitution pattern may look like (i.e. non-discrete substitution on an aggregate scale). The practice of defining the market itself also poses problems, as the boundaries of a market are often times ambiguous, both in the scope of the products in the market and the geographic scope of the market. These are all non-trivial objective (and often times subjective problems) that must be addressed in order to perform a merger analysis. An oft-quoted case illustrating the difficulties authorities face in accurately defining a relevant market is that of the FTC v. Whole Foods[6]. The FTC had moved to block a proposed merger between two premium natural/organic supermarkets, with the defense that the result would harm competition

[6] https://www.ftc.gov/system/files/documents/cases/080114ftcwholefoodsproofbrief.pdf

and decrease welfare. This indeed would likely have been the case, if the aforementioned relevant market was distinct and separate from that of more traditional supermarkets. The court however ruled that it was not, stating that the FTC had failed to prove that their market definition was the relevant definition. One can see how the use of a strict concentration measure in such a case would give widely varying results depending on the breadth of the market definition.

Recent research has however attempted to validate the usage of the Herfindahl Index as a tool in merger analysis via the microeconomic foundations of the tool. Miller, et. al (2016) note that in order to use the Herfindahl Index as a predictor of price changes, consumer diversion between the merging parties and outside parties must be proportional to market share. Under this condition, diversion of consumers from product i to product j becomes $s(j)/(1-s(i))$. However it is noted by these authors that the unilateral competitive effects and the change in the Hefindahl Index, noted as DHHI, may only be weakly correlated due to the interactions of this diversion definition and mark ups. Dafny, Duggan and Ramanarayanan (2012) corroborate this notion in their econometric study of health insurance premiums in the United States; they concluded that increased market concentration could explain roughly 12 percent of the increase in insurance premiums a significant but small amount. The arrival of new tools in the competition economics toolbox has not served to dismiss the HHI's analytical usefulness. Rather they have and will be used along side each other in both complementary and substituting roles.

2.4 Upward Pricing Pressure as a Tool For Merger Analysis

The need for development of an analytical merger device with strong microeconomic foundations has long been recognized in industrial organisation and competition analysis (see Werden (1996) and O'Brien and Salop (2000)). Building on such previous research, the upward pricing pressure as a measurement tool was first formally defined by Shapiro and Farrell in 2010. The authors, who represented the DOJ and FTC respectively at the time of publication, echoed the concerns discussed in the previous section and sought to create an analytical tool that could be used quickly and practically by authorities with a given level of diagnostic accuracy. They specified a measure based on differentiated product Bertrand competition with which the loss of direct competition is compared to the estimates of efficiency gains in marginal-costs.

As a simple example, consider a market where two firms, each producing a single product, decide to merge, and that the market is characterized by Bertrand competition. The post-merger price for the combined firm can be written as:

$$p^m = mc_1 + \frac{1}{[\partial q_1/\partial p_1]/q_1} + (p_2^m - mc_2)\frac{\partial q_2}{\partial p_1}/\frac{\partial q_1}{\partial p_1} \tag{4}$$

The difference between the pre- and post-merger price is the pricing pressure index:

$$UPP_1 \approx p_1^m - p_1 = (p_2 - mc_2)\frac{\partial q_2}{\partial p_1}/\frac{\partial q_1}{\partial p_1} - E_1 mc_1 \tag{5}$$

The diversion ratio D_{12} is defined as:

$$D_{12} \equiv \frac{\partial q_2}{\partial p_1}/\frac{\partial q_1}{\partial p_1} \tag{6}$$

The Upward Pricing Pressure index uses a second term in order to incorporates the potential efficiency gains that can come from consolidation. The term E_1 reflects an efficiency credit, i.e. a parameter that accounts for the reduction in marginal cost for the combined firm because of the synergies resulting from the merger. In practice, the efficiency credit is not a measurement nor a prediction. Rather, it is an assumption that is used in the tool for screening purposes (Werden and Froeb, 2011). How the efficiency should be determined is sometimes debated and can vary from case to case. Farrell and

Shapiro (2010) suggest using 10% reduction in marginal cost which is sometimes referred to as a default efficiency credit. The added feature of efficiency gains is beneficial, since significant synergies do occur in some mergers. However, it can be argued that it also adds a level of arbitrariness to the UPP.

The UPP has recently been given much attention and is appreciated for its simplicity and intuitive interpretation. However, the challenge in using the UPP lies in the application of it. While it for example does not require a large amount of data, the information needed is more extensive. At the initial screening stage, which is when using the UPP is typically proposed for, information is normally very scarce and limited. When they proposed the measure, Farrell and Shapiro (2010) argued that the information needed for utilizing the UPP often is readily available at an initial screening stage; for example, profit margins for the merging firms must often be submitted by the merging parties regardless of the analytical approach taken by the authorities. Diversion ratios on the other hand, are seldom available and must be estimated before the tool can be used. Such a procedure usually takes considerable time and considerable resources may have to be devoted to the task of estimating them, depending on the data available at the time.

2.4.1 Where Does Data On Margins Come From?

The importance of margins with regards to unilateral price effects is both intuitive and well-studied. It is directly stated in the US Horizontal Merger Guidelines (2010) that "high pre-merger margins may also make significant price increases more likely", with the reason being that the diverted sales that can be recaptured after a price increase are more valuable if margins are high and the merged entities are able to make a price rise profitable. It is straightforward therefore to say that higher margins leads to a higher upward pricing pressure.

In many cases, margins are submitted by the merging parties to the competition authority before the screening process, and thus the UPP can often be calculated based on these. It is not perfectly clear however, how to make sure that the reported margins represent the correct definition of margins as the model interprets them; the cost that is subtracted from the price of each product in a firm's first order condition is supposed to be the marginal cost. This can be troublesome for a couple of reasons. First, as is the case in the airline industry, defining a product is not always as straightforward as it can be in some retail industries. Secondly, when measures of variable costs are extracted from accounting data they are often in the form average cost rather than marginal cost(Werden and Froeb, 2011). Even though accounting measures can often act as a reasonable proxy for the actual marginal costs for branded consumer products it is difficult to verify that this is the case. Therefore the measure of marginal costs can be more complex, or arbitrary, in a real-world setting than in the theoretical model. The correct margins to use for calculations of UPP are the margins that, in the best possible way, reflects the profits that are made on the recaptured sales. They should thus be whatever measure most accurately reflects the profit made on this marginal sale. A firms accounting data often provides practitioners with the average margins over the total sales of a specifically defined product. This will differ from the marginal profit in several cases - for example, if the firms discriminate in their price setting between loyal customers and the more unattached customers, the margins on the loyal customers are only infra-marginal and the average will be incorrect. The airline industry is often taken as a schoolbook example of price discrimination (loyalty bonuses and first class seating are two examples), and thus, accounting data might not do a sufficient job of serving as a proxy for marginal profit or costs even if we have strictly defined the products appropriately.

Finally, it should be noted that the measures supplied by accounting data most likely do not contain the information that allows us to investigate the incremental profit of additional sales in the long run. This is however only important if the authority wants to understand how any price effects from the merger will persist in the future. The incremental profit in the long run is likely to be lower than the immediate incremental profit since the merged entity will have to consider in the long run how to maintain the needed capacity in order to continue supplying the diverted sales. As we have mentioned, in the airline industry capacity constraints can differ between periods of high and low demand, which may make this issue more complex, and thus it is beyond the scope of this paper.

2.4.2 Where Does Data On Diversion Ratios Come From?

The most theoretically accurate method of determining diversion ratios as needed to estimate the upward pricing pressure is by directly calculating them from the the own- and cross-price elasticities of products in each market. The system of price elasticities however, needs then to be estimated in a credible way if they are not readily availible (previous competition cases or economic literature can offer estimations, as two examples). Own-price elasticity, ϵ_{jj}, is the percentage change in demand for product j when its price is changed by 1% (or an even smaller marginal increase), whereas the cross-price elasticity of product j with respect to $i's$ price, ϵ_{ji}, measures the percentage change in demand for product j that accompanies a marginal increase in the price of product i. The own-price elasticity is negative for normal goods; substitute products have a positive cross-price elasticity.

There are two commonly used types of diversion ratio: the unit sales (or customer) diversion ratio, and the revenue diversion ratio. We only focus on sales diversion ratios. Sales or customer diversion ratio can be written as:

$$D_{ji} = \frac{\partial q_i / \partial p_j}{\partial q_j / \partial p_j} = \frac{\epsilon_{ij} q_i}{\epsilon_{jj} q_j} \tag{7}$$

In practice, there are at least three ways to calculate or estimate diversion ratios. First, in the case where detailed price and quantity data can be collected (eg, store scanner data), a demand estimation can be modelled and own- and cross-price elasticities can be obtained. For example this approach were adopted in cases such as the Volvo/Scania[7] and Kimberley-Clark/Scott mergers[8]. However, a few key issues arise with regards to estimating diversion ratios using demand estimations. While it may be the most precise and advanced methodology for calculating such ratios, it is not likely to be possible at an initial screening stage where time constraints are tight (precisely why the UPP was developed). Moreover, if all the price and quantity data; time and resources are available to estimate demand, why not use a more flexible and sophisticated merger simulation to estimate the unilateral merger effect? Indeed, this trade-off between precision and user-friendliness is a common theme across the merger analysis literature.

Another source for the diversion ratios necessary for UPP analysis can be data collected by the firms themselves during the course of business. Firms may conduct surveys with existing customers regarding where or who they would shift their business to if the present option was unavailable. New customers may be surveyed on where they previously shopped. In some cases it may be possible to estimate diversion ratios based on this information, although in practice such information is rarely complete or available in an appropriate form for the purposes of calculating diversion ratios reliably. Because the motivation of a firm in collecting such information will likely be different

[7] Case No COMP/M.1672, paragraph 178, `http://ec.europa.eu/competition/mergers/cases/decisions/m1672_en.pdf`

[8] Case No IV/M.623, paragraph 172, `http://ec.europa.eu/competition/mergers/cases/decisions/m623_en.pdf`

from that of competition authorities, survey designs may be prone to biases that present favorable results for the firm. Framing, non-responses, and coverage completeness are all biases that must be accounted for in creating an accurate consumer survey, biases that firms may neglect in survey design.

Similarly, diversion ratios can be based on consumer surveys orchestrated by the competition authority themselves. Given the complications associated with the above two approaches, customer surveys tend to be the most common and practical method for obtaining diversion ratios in actual competition cases. This involves asking consumers directly which products or firms they would substitute to if they were to switch away from the currently chosen product following a small price increase. This methodology is often undertaken by competition authorities themselves as part of the due diligence in analyzing the pre-merger market conditions, and has the advantage that biases in survey design can be minimized and precise information can be collected.

2.4.3 UPP As a Price Predictor

In order to explain how the UPP relates to actual price changes, rather than just pricing pressure, we first need to go through the concept of pass-through rates. Pass-through rates refer to the extent of which an increase in marginal cost is passed through to consumer prices. Or put differently, the rate at which a decrease in the firms cost corresponds to a decrease of the firm's price.

In the theoretical models that lay the foundation for merger simulations, pass-through rates comes straight from the demand system, which is estimated. There exists a large body of literature that looks at how functional form restrictions on the demand system can affect the accuracy of the simulation (e.g. Froeb, Tschantz and Werden (2005); Miller, Remer, Ryan, Sheu et al. (2016). A more general solution to quantify price changes caused by a merger was introduced by Jaffe and Weyl (2013), which shares the same base principals as the merger simulation. If one can observe pass-through rates, or if one can accurately estimated them from data, then a first order approximation to the price change can be calculated. The first order approximation equals the UPP multiplied by the merger pass-through matrix. Upon inspection, they show that merger pass-through rates depend only on the first and second derivatives of demand, while higher order terms can be omitted. Miller, Remer, Ryan and Sheu (2015) provide Monte Carlo evidence that the first order approximation is an accurate predictor of true price effects. Considering this insight, it then becomes clear that the UPP itself may possibly provide a useful prediction of the price changes on its own, if it is the case that the identity matrix can accurately proxy for the merger pass-through matrix $[G]$:

$$\Delta P_j = [G] * UPPji \tag{8}$$

In some cases, the UPP nearly equals the first order approximation for the merging firms. This happens because the diagonal elements of the merger pass-through matrix are somewhat below one, while the off-diagonal elements are positive. Thus, using the identify matrix to proxy merger pass-through overstates some effects and understates others. This results in a measure that is close in magnitude to the first order approximation because of these counterbalancing effects.

Miller, Remer, Ryan, Sheu et al. (2016) gives a simple example that illustrates how the UPP can work as a first order approximation. We will revisit this exercise for the same purpose. We consider a market where demand is given by a logit demand system. The market is populated by three firms. Each firm has a margin of 50% and a market share each of 30%, with the last 10% of the potential market belonging to an "outside good". If the first two firms merge and become one, the first order approximation described above in (8) would become

$$\begin{bmatrix} 0.204 \\ 0.204 \\ 0.052 \end{bmatrix} = \begin{bmatrix} 0.771 & 0.180 & 0.297 \\ 0.180 & 0.771 & 0.297 \\ 0.122 & 0.122 & 0.776 \end{bmatrix} \begin{bmatrix} 0.214 \\ 0.214 \\ 0 \end{bmatrix} \tag{9}$$

where 0.214 is the UPP for firm one and two. The first order approximation (0.204) is close to the true price effect from the re-equilibration, which is 0.190. Interestingly, the UPP is also relatively close to the true price change. By assuming a passthrough rate of one (or more accurately the identity matrix), the UPP can serve as a useful predictor of such a price change. This is the idea that we consider in this thesis. When the actual passthrough matrix has the feature as above, with the diagonal elements below one and off-diagonal elements above zero, the countervailing biases makes the UPP a slightly overestimating approximation. This is the case when the true demand system exhibits log-concavity, as it occurs in the logit demand model.

2.5 Merger Analysis Using Simulation Techniques

Merger analysis using simulation techniques is based on a combination of theoretical models of market conduct and structure as well as information and assumptions about the overall parameters of the model. Such a model is then used to predict the effect of the merger on post-merger equilibrium, with the key output being predicted price changes. Typically, this starts with estimating demand from pre-merger data that models how quantities for all products in the market depends on different parameters and characteristics. The merger simulation is then made by simply stating that the competitive choices previously made separately by two firms are now modelled as being one, allowing the outcome to re-equilibrate.

Due to both improved technology and evolving methodology, the 1990s proved to be a turning point with regards to the econometric techniques available to practitioners and authorities. Specifically, techniques for specifying and estimating structural models of consumer demand were postulated by a range of authors; such advances made merger simulations in markets with many differentiated products possible for the first time. Berry, Levinsohn and Pakes (1995) introduced a random-coefficients, discrete-choice logit specification as a way to estimate demand in the automobile market. Nevo (2001) extended the literature, providing a study of ready-to-eat cereal brands as well as comprehensive supporting documents describing the technique in full. Their methodology allowed for characteristics of the products (color of cars, or mushiness of cereals, as two relevant examples) to be quantified and analyzed, with consumer preferences for said characteristics modelled accurately as a result. As these techniques have evolved, the computational simulations of the effects of real-world mergers has been an increasingly important tool for competition policy, both in the US and in Europe.

2.6 Comparison Of UPP And Full Merger Simulation

We have stated that the UPP has received much of its popularity due to its pragmatic simplicity. In order to illustrate this simplicity of the UPP we here go through the theoretical differences between the UPP and the full merger simulation. Firstly, let us view the modelled behavior that maximize the firms profit. In a market t, there are $k = 1, 2, ..., K$ firms each producing a total of $j = 1, 2, ..., J$ differentiated products. Each firm is assumed to maximize its profit to the best of their capabilities. When the equilibrium is of Bertrand form, the profit function of the firm is given by (1). The firm's profit in each market depends on the markups and market shares for each product offered in the market. Note that this is the profit of a multiproduct firm. Maximizing profit with respect to price yields the first order condition given by (2). The firm takes the effects of all products in its ownership into consideration when setting its prices. A merger simulation, the way we refer to it in this paper, is essentially using this first order condition to set a new equilibrium under a proposed scenario. When looking at the unilateral price effects of a merger, we change the ownership structure of the market such that the firm now takes a different set of products into consideration in the profit maximization formula. One needs therefore to know how demand depends on the prices of each product available in the market and the complete pattern of substitution between said products. Thus, a demand estimation is at the heart of the merger simulation. It should be noted that the firms are not considered to be multi-market maximizers. However, recent research by (among others) Benkard, Bodoh-Creed and Lazarev (2010) suggests that airlines consider their entire hub-and-spoke network when considering profitability and mergers. Following this line of thought it may be more appropriate to consider airline firms as multi-market maximizers in order to cover a more comprehensive scope.

The UPP can intuitively be thought of as a single product merger simulation. It ignores the re-equilibration of the other endogenous variables in the merger simulation. By holding quantities, partial derivatives and all competitors prices constant, and taking the difference between the post- and pre-merger first order conditions, we get the UPP formula. The intention of the UPP measure was to do away with the need for a structural demand estimation by assuming that only the price of the merging firm is re-equilibrating. While gaining a practical and pragmatic tool with theoretical and computational convenience, as well as fewer data requirements, economists who choose to use the UPP are faced with new challenges associated with the tools simplicity. This is an issue that academia and practitioners are well aware of and has therefore led to various recommendations on how the UPP should be used. Farrell and Shapiro (2010) suggest only using the sign of the UPP, since it gives a robust measure for the incentive for price change. On the other hand Pakes (2010) advocates using the magnitude as a rough approximation of the merger simulation price change and Jaffe and Weyl

(2010) suggest using the UPP in combination with its first derivative to approximate the simulated price changes in a another manner. In order to make more assertive claims about the usefulness of the UPP as a tool to assess the anticompetitive effects or a merger it is important to fully understand under what circumstances the UPP makes accurate predictions about the effects and when it makes worse predictions than another available alternative. Therefore it is important for the economic literature to better investigate which tools are better at accurately predicting what actually happens after the merger, not just what performs best in theory.

2.6.1 The Debate About The UPP Versus Merger Simulations

The effort to compare and contrast the UPP with merger simulation techniques have been ongoing since the breakthrough of the UPP. Part of the debate among prominent antitrust academics is summarized by Werden & Froeb (2011). We here give a short resum of some of the main arguments in order to clarify the economic intuition and practical meaning of the theoretical comparison.

It is first important to note, as we have already made clear, that the UPP was intended as a screening tool and was brought forward for its simplicity and robustness. Farrell and Shapiro (2010) made the case that the UPP is robust as a screening tool because it does not have the same problems as the more complex merger simulations. They state that the merger simulations “risk miss-specification” when the merger takes place in “complex industries” because of the ambitious task of fitting a structural model to the data. The UPP, in contrast, gauges the fundamental information in the market and quantifies the change in incentives since it “nets out such complexities that are present both before and after the merger”. Moreover, they state that the merger simulations “tends to be opaque to non-specialists” such as lawyers and economists outside of antitrust analysis and “can be demanding in terms of data requirements”.

This is not to say that rigor and complexity should be passed over in favor of a single simple measure; rather, the UPP is a practical tool that can quickly and easily be presented in many settings.

Further critique towards the merger simulations came from Joseph Simons and Malcolm Coate (2010) They state that the methods are not "very successful" because of the fact that they generally requires "reasonably precise estimates of a demand system" which is proven to be a non-trivial task indeed. Simons and Coate go on to criticize the use of merger simulations because of the lack of empirical evidence that the predictions by these models are accurate and reliable. They say that the lack of empirical confirmation makes it difficult to use as viable evidence in a court. Although they direct similar critique towards the UPP, they propose using a version of the UPP as a final assessment of unilateral effects as it "reduces the black box nature of merger simulation and minimises the data requirements". Overall, however, they oppose the intention of Farrell and Shapiro to use the UPP as a screening tool because of lack of empirical evidence.

In another response to Farrell and Shapiro, Epstein and Rubinfeld (2004) argued that the data that is needed to calculate the UPP can be put to better use with another method, they propose a method similar to compensating marginal cost reduction (CMCR). In a response to this, Farrell and Shapiro answered that the theoretical derivation of the UPP is far more general than that of CMCR since it is more flexible in its assumptions about market structure. Furthermore, both Richard Schmalensee and Simons and Coate have brought forward their arguments for using different calculations and assumptions than the UPP uses to measure price effects of horizontal mergers. Schmalensee argues that the UPP does not measure what is actually causing consumer harm saying that the UPP is a less directly relevant quantity. He proposes the use of simple merger simulations instead for this reason because it measures concrete price increases which are closely linked to measures of consumer welfare.

In summary, critiques have been directed towards the UPP for being too simple and for not being a measure of what actually matters, i.e. changes in price. It has been proposed that with the level of data needed for UPP analyses, a simple merger simulation produces much richer results with only marginally more effort required. On the other hand, there has been criticism directed towards the use of merger simulations because their complexity clouds the interpretation of the results for non-economists. Furthermore, there has been critique directed to both methods because there is not much empirical evidence to support the predictive accuracy of either these methods. This thesis addresses directly the discussion about the UPPs relation to price changes. We also address the discussion about the lack of empirical evaluation of the UPP and an alternative merger simulation tool.

2.7 Ex-Post Evaluation of Merger Analysis Techniques

This thesis also concerns the literature surrounding ex-post testing of merger screening tools. We relate in particular to two types of research: the literature of predicting effects of mergers by the use of structural models á la Werden & Froeb (1994), and the much more scarce literature on ex-post evaluation of the predictive power of the UPP as a merger screening tool (which is not completely unexpected, given the recent timeline of this tools development). The predicted price changes from both merger simulations and the UPP can be viewed as a forecast of the post merger outcome which in practice can be used to evaluate the appropriateness of the merger. It is therefore important to empirically assess how accurate these predictions are. While there can be a variety of reasons for why predictions differ to the actual outcome it is important to point out that the merger simulation and the UPP aims to quantify the same effect. Therefore, the accuracy of the two predictions can be compared. Since the simulated prices unavoidably rely on assumptions about demand and market structure, a comparison between simulated and actual post-merger prices can also be viewed as a test of these assumptions. This is where the UPP and the full merger simulations differ and largely where we hypothesize the differences in predictive power arise.

Our research closely relates to Peters (2003) in that he evaluates merger simulation techniques in the airline industry. Peters looks at three different demand models for his simulations, all of which are discrete choice models. He states that the random coefficient model of Berry et al. (1995) allows for more flexibility and more precise assumptions than the models he analyzes, but that the computational costs associated with the models complexity prevents him from doing so. The different techniques are then applied to five different mergers in the airline industry between 1986 and 1987, and he finds that the merger simulations do have predictive power. However, Peters stresses that the post merger prices depend not only on the changes in market structure but also on supply side factors, such as changes in marginal cost and firm conduct. This is of course a factor that usually neither the UPP nor merger simulation methods take into account. In practice, this typically involves a more comprehensive analysis such as that of Berry and Jia (2010), one combining qualitative and quantitative methods and an aggregate assessment of entire industrial sectors. If the main source of difference comes from supply side mechanisms however, one could speculate that the UPP does not have to be worse in this regard than the full merger simulation, since neither of them take it into account.

Bass, Huang, and Rojas (2008) examine the severity of structural misspecification of the demand system on merger simulations using a series of Monte Carlo experiments. Using Monte Carlo experiments to test empirical models is one way of linking the theory behind the empirical models to their performance. They find that the discrete choice model generates better predictions of the unilateral effects of a hypothetical merger compared to the other models tested. This was the case even when the true demand model specification was not logit; the logit still gives reasonable predictions of the price effect of a merger. They also point out that the market potential, i.e. the market size, is a key component for estimating the own and cross price elasticities. Overestimating the market size may lead to biases in the estimated elasticities, which may be a significant problem when the aggregated data comes from a sample of the market.

Most of the studies that have looked at the actual outcome of mergers ex-post have indicated that price increases are indeed common as a result of mergers Werden and Froeb (1994). Some of these studies also made notes on the timing of these price increases. This can be an important observation for several reasons. While there exists an academic consensus that acknowledges that mergers can result in both incentives to raise prices and reductions in marginal costs, there is little evidence or consensus about when this is realized. Cost savings might very well take time after the merger is approved and operationalized. We have to consider that many studies on merger retrospectives have only covered a short term effect which might not include the realization of cost savings due to efficiency gains. Furthermore, there are at least two studies that indicate that the merging firms in the airline industry may be raising prices before they are legally allowed to coordinate their operations (see Kim and Singal (1993) and Borenstein (1990)). The phenomenon of price increases before the end of the waiting period (or even before the merger is publically announced) is something that our standard models of unilateral effects cannot explain since they are not yet one firm. Such cases of coordination between firms indeed prove difficult to investigate, as proving collusive behavior requires significant legal support. Either way, it is a situation that we do not investigate in this paper.

One alternative explanation lies in fact that the airline industry is known to have significant switching costs through, as one example, the use of frequent flyer programs

(Carlsson and Löfgren, 2006). In an industry where consumers face switching costs there is an incentive for the firm to initially charge a lower price in order to win the customers over and gain market shares (Klemperer, 1995). Once they are won over, the switching cost would bind them to stay and the surplus for the firm is extracted at a later stage by raising prices. The reason prices might be raised earlier in these industries in the case of a merger is due to mechanisms of corporate governance. If the manager of the firm knows that a merger will happen, the previous incentive to keep prices low in order to gain market shares might be lost. If management is in risk of losing their jobs due to restructuring the incentives for long term investment in market share can be gone. Prices would then rise as soon as management knows about the merger and fall as soon as new management comes in (Weinberg et al., 2007).

The literature on evaluating the UPP as a screening tool is scarce. In one example, Miller, Remer, Ryan, Sheu et al. (2016) evaluate the accuracy of the UPP in predicting post merger price changes by conducting Monte Carlo simulations of artificial mergers. Their results are overall encouraging and they find that the UPP is quite accurate for many demand specifications. In some cases, when the full merger simulation was done with incorrect functional form assumptions, the UPP outperformed the simulations. They do stress however, that while on the one hand the study has few implications for empirical discussions, it is nevertheless important for antitrust practitioners due to its theoretical clarifications and re-confirmation of the expediency and usefulness of the UPP. Judging by the results of research with monte carlo experiments, where the true underlying demand system can be controlled for, the simpler models show a certain degree of robustness to misspecifications.

An attempt to examine this difference with empirical methods was done in a working paper by Cheung (2013), which seeks to compare the price predictions based on UPP to a more complex merger simulation. Drawing conclusions from this kind of comparison however has to be based on assumptions about the accuracy of the granular merger simulation. By referring to the demand substitution patterns from the structural demand estimations as true ratios it explicitly assumes that the there are no errors in the demand estimation.

2.8 Airline Specific Literature

This has subsequently led to a wealth of research on the industry. Joskow, Werden and Johnson (1994) explicitly showed that a new airline entrant in a market increases output and reduces fares, while the opposite is also true for exit. Berry, Carnall and Spiller (1996) have recently written on the overall demand for airline products and changes in the industry over the past 15 years while specifying a model which attempts to capture heterogeneity amongst consumers. Their slight modification of the random-coefficients logit model of demand estimation used by Berry, Levinsohn and Pakes (1995) to describe the automobile industry allows them to model the airline industry according to their assumption that there exists a discrete, bimodal distribution of consumer types. They claim that airline passengers largely fall into one of two distinct groups: those that are business travelers, with relatively low price sensitivity and preference for direct flights, and those who are more price sensitive and less concerned with other characteristics (the so-called leisure traveler). This view of the industry and consumer preferences was validated by the US Department of Justice in a speech given by McDonald (2005), where he notes that both airlines and antitrust analysts take the leisure-business consumer distinction seriously when setting prices and evaluating the effects of mergers, respectively. The specifics of the random-coefficient logit model and the modifications made to it will be introduced and describe in detail in a later section.

Berry and Jia (2010) further on the aforementioned structural demand model of the industry by performing a comprehensive analysis of the airline industry, one in which they model both the cost and demand functions of the industry in order to explain many of the changes which occurred in the industry between the late 90s and 2006. Their estimations allow them to classify how consumers preferences for flight characteristics changed over the studied period: passengers shifted towards preferring direct flights and demonstrated overall an increase in price sensitivity. Again, the model will be described in detail later, but one of the main advantages of using such a demand specification is that it allows for the quantification of both observable characteristics of products (such as direct vs. connecting flights) and unobservable characteristics (day of the week the flight occurs on, as one example. The fact that there has been so much consolidation in the airline industry in the United States over the past couple of decades has given academics plenty of events to study with regards to mergers in the industry. Werden, Joskow and Johnson (1991) were among the first to empirically examine the direct impact of a merger in any industry when they analyzed price and quantity effects surrounding the merger of TWA and Ozark Airlines in 1986. They concluded that in the year following the merger there was in fact a significant increase in fares and a reduction in the overall number of flights being offered in certain markets. This theoretically sound outcome was reinforced by (among others) Brueckner and Pels (2005) in a similar paper focused on the European market and the merger between Air France and KLM in the mid-2000s. More recently, Cheung (2013) performed a comparison of merger analysis tools and their outputs with regards to the same merger we have chosen to study. Without looking at any actual price changes in the aviation sector, Cheung presents a theoretical comparison of merger analysis tools in the context of the America West/US Airways merger of 2005. She states that, among other things, the pass-through conditions of the UPP as outlined by its original authors are not trivial and can have serious implications for the resulting calculations. Finally, Peters (2003) revisited the concept of ex-post analysis of price changes following a merger. Peters however placed the price changes side-by-side with the results from different methods of simulating the merger and compared the predicted changes with the actual changes. His main findings indicate that the merger simulations did not provide accurate predictions of the realized post-merger prices.

3 Methodology

While the inputs needed to run a basic merger simulation are straight forward enough prices, quantities, and margins obtaining these in a coherent, structured manner is unfortunately not. Airline fares, as well as economically sound estimations of market size and thereafter market share, are extracted through some simple data manipulation. However obtaining information on margins of airlines on a product by product basis requires advanced econometric techniques. Thus, in order to perform the simulated merger of America West Airlines and US Airways, we must 1) extract margins from an econometrically sound procedure, 2) input the margins, prices, and market shares into a merger simulation and 3) interpret these results and compare them to a UPP measure.

3.1 Choice Of Demand Estimation Model

In order to simulate a merger, information about mark-ups on products must be provided as an input. In order to get consistent approximations of actual air travel mark-ups, we begin by modelling demand of air travel in the United States. Specific demand models for industries with differentiated products have been developed within the field of industrial organisation, with most of the modern, more advanced models falling into one of two main categories: functional form models, where a parametric functional form is chosen to describe the relationship between quantities and prices, and discrete choice models, where the econometrician explicitly specifies the utility function of the consumer in the market, from which the demand relationship is derived. The functional form models includes the Almost Ideal Demand System (AIDS) and constant-elasticity models. The discrete choice models on the other hand includes the logit, nested logit and random coefficient models, all to some extent being extensions of one another. Both types have been extensively used by competition economists and antitrust authorities for merger analysis. For the airline industry, and for our purpose, the discrete choice models are likely to be the better choice for a number of reasons. First, our data is quarterly and not more frequent as may sometimes be the case in for example retail industries where there can be weekly or daily time series of data. For this reason, we would need years of data to collect enough observations to reliably estimate a separate demand model for each market, and this is problematic for obvious reasons. Second, one feature of the airline industry is that the same products are not available in different geographical locations. We cannot compare tickets with the same airline but with different routes as we can with for example breakfast cereals of the same brand sold in different cities; a flight to New York is much different than a flight to New Mexico, whereas branded cereal in both locations will be exactly the same. Essentially, a functional form model makes the assumption that the cross-price elasticity between two brands are similar in all cities, while the discrete choice methods model demand based on the specific characteristics of a particular product. Our model builds on well-established demand estimation techniques such as those of Berry, Levinsohn and Pakes (1995) (hereafter referred to as BLP). Fundamentally, the BLP (1995) procedure is in the style of structural modelling, in which empirical analysis starts with a rigorous theoretical model in which consumers are utility maximizers, firms are profit maximizers, and everything, including the error terms, has an economic interpretation. That is in contrast to the reduced form approach, in which the economist essentially looks for conditional correlations consistent with his theory. What we get with the structural approach is the assurance that we do have a self-consistent theory and the ability to test much finer hypotheses about economic behavior and consumer preferences. What we lose by choosing this procedure is simplicity and robustness to specification error. The random coefficient models, arguably the more flexible and advanced of the discrete choice methods also maintains the advantage of the other models in handling a large number of products. It is superior to prior methods because 1) the model can be estimated using only market-level price and quantity data, 2) it deals with the endogeneity of prices, and 3) it produces demand elasticities that are more realisticin particular, cross-price elasticities are larger for products that are closer together in terms of their characteristics. This is of course a desired feature for a merger simulation, since the outcome largely relies on the estimated price elasticities. In the ideal world of the econometrician you would, when estimating price elasticities, conduct a controlled experiment as the first choice. Ideally you would want to expose customers to randomly assigned prices and record the consumer behavior. Of course, in the real world prices are not randomly assigned. In fact, quite the opposite they are assigned by profit-maximizing firms that take all the available information into account. Some of this information is known to the firms but not to us, as econometricians; thus, we have to include it in our error term. These unobservables are a feature that our model acknowledges and includes. One general solution to this problem is the use of instrumental variables. We also acknowledge the fact that different market spaces have different products available and that the data from one market therefore should not be seen as one observation of consumption for a given list of prices. Instead, in each market it is the relative likelihood of purchasing products, viewed as different bundles of characteristics. Our model then connects these probabilities to a utility function in order to estimate price elasticities. This model therefore increases its power by increasing the number of markets with variation in different bundles of characteristics on offer. It should be noted that if some relevant characteristics are not observed, and these are correlated with our observed characteristics we will face another problem of endogeneity. For example, departure time and how far in advance the ticket was purchased is not included in available data. This might however, not be of vital importance since these attributes are more likely to differentiate the price of products within one airline, rather than systematic differences in price across airlines. Another issue that needs to be addressed is time. In markets where dynamic effect plays an important role it can be difficult to appropriately account for short and long term consumer behavior. In some retail markets promotions and households inventories can create short-run responses to price changes that do not reflect the appropriate demand response that is relevant to antitrust analysis (Baker & Reitman (2009)). In our case this might not be of great importance since the airline industry perhaps does not exhibit the same characteristics, especially not consumer stockpiling. Further, we use quarterly data and not a more frequent series like those sometimes used for retail demand estimations, which may bring the risk of inconsistent coefficients. However, our use of quarterly data and the choice of only using one quarter pre-merger comes with a trade-off. On the one hand we have to be aware that seasonal patterns exist in the airline industry, which we do not investigate, ann the other hand we do not have to worry about the negative impact of averaging these effects out.

3.2 The Model

In following others who have used a random-coefficients discrete choice model (Berry et al. (1996) and Berry and Jia (2010)), we have made a few modifications to the well-known BLP procedure in order to better accommodate some features of the airline industry. It is still a random coefficient logit model, epmloying a discrete choice framework that assumes that consumers in each market choose from a menu of 'inside' good products (air travel from a variety of airlines) and an 'outside' good (the choice to not participate in the airline market other means of travelling than flying). The first step in constructing such an estimation is to model the utility for consumer i, in market t, buying product j, which takes the following form:

$$u_{ijt} = x_{jt}\beta_r - \alpha_r p_{jt} + \xi_{jt} + v_{it}(\lambda) + \lambda\epsilon_{ijt} \tag{10}$$

The variables of the utility function are defined as follows:

- x_{jt} is a vector of observable product characteristics, such as the number of connecting flights and distance
- β_r is a vector of preferences for the observable characteristics, specific to consumers of type r
- α_r is the price sensitivity for consumers of type r
- p_{jt} is the price of the flight
- ξ_{jt} is the unobservable (to researchers) characteristics of product j,
- v_{it} is the nested logit random taste that is constant across airline products and differentiates air travel from the outside good,
- λ is the nested logit parameter that varies between 0 and 1,
- $\epsilon_{ijt}$is an independently and identically distributed (across products and consumers) logit error.

An important difference from earlier versions of discrete choice models is that the taste parameters in the random coefficient logit model are allowed to vary between consumers. This is arguably the biggest contribution of BLP (1995); assumptions are made about the population, and, with the use of demographic information and other known characteristics about the consumers, random "consumers" are sampled and connected to the product which most suits their set of randomly-assigned, normally distributed preferences. However, in following Berry and Jia (2010), we make a strong assumption about the consumers heterogeneity. We assume there exists in the population two distinct types of air travel consumers: those with high price sensitivity and those with low price sensitivity - the so-called "business" and "leisure" traveler distinction as discussed in the background literature. Because of this assumption, our taste parameters can take on only two values: (β_1, α_1) and (β_2, α_2). Furthermore, we assume that the taste parameters take on a bimodal distribution with a certain probability γ and $(1 - \gamma)$ that consumer i is of type 1 or 2 respectively. Intuitively, this means that leisure (business) travelers self-select their "type" by purchasing cheaper (more expensive) products (along with other characteristics that may be correlated across consumer types) (this assumption also proves to offer a huge computational simplification which will be discussed below). Our data has information on fare class, which one would initially think would capture this heterogeneity quite well; first class fares are often prohibitively expensive for the family travelling on vacation, and provide perks such as early boarding and disembark that would attract the time-sensitive business traveller. This measure is not necessarily, however, a good indication of the consumer type, as the data on fare class is self-reported by airlines, where some choose to use the business class distinction for all itineraries sold (Jet Blue, as one example). This fact makes the indicator unreliable. Thus, we need a way to verify our assumption that two distinct groups of price sensitivity exist other than our intuition. We also need a way of identifying the probability of being in each group, which is our parameter gamma. If we believe that leisure travelers purchase products that, on average, have different characteristics from those purchased by business travelers, this should be identifiable in the data.

After specifying the utility function, the next step in constructing the demand model is to create a way to represent the probability of each consumer selecting each of the products in the market. Conditional on actually choosing to travel by air, the probability of a consumer of type r purchasing product j in market t is:

$$\frac{e^{(x_{jt}\beta_r - \alpha_r p_{jt} + \xi_{jt})/\lambda}}{Z_{rt}} \tag{11}$$

Where Z_{rt} is equal to:

$$Z_{rt} = \sum_{k=1}^{J} e^{(x_{kt}\beta_r - \alpha_r p_{kt} + \xi_{kt})/\lambda} \tag{12}$$

This is essentially the utility for consumer type r of choosing product j over the sum of utility of choosing any product of air travelling. The probability of a type r consumer choosing to fly (versus driving, or making a phone call, as two examples of substitutes for participating in the airline market) is:

$$s_t^r(x_t, p_t, \xi_t, \theta_d) \equiv \frac{Z_{rt}^{\lambda}}{1 + Z_{rt}^{\lambda}} \tag{13}$$

Where θ_d contains the parameters to be estimated. This term $s_t^r(\bullet)$ can intuitively be thought of as the sum of the utility that one gets from any inside product over the sum of both inside and outside product utility. Since the utility of the outside good is normalized to zero, $e^0 = 1$. If we let γ_r be the percentage of type r consumers in the population then the overall market share of product j in market t is

$$s_{jt}(x_t, p_t, \xi_t, \theta_d) \equiv \sum_r \frac{e^{(x_{kt}\beta_r - \alpha_r p_{kt} + \xi_{kt})/\lambda}}{Z_{rt}} s_t^r(x_t, p_t, \xi_t, \theta_d) \tag{14}$$

The simplicity of assuming discrete types of consumers in the population is revealed above. One will notice that the model's theoretically predicted market share is a summation of each consumer types respective weighted market share. In a full random-coefficients specification (i.e. those under assumptions of normally distributed characteristics) the spectrum of consumer types is continuous, and thus requires integration to analytically define the market share function. This also negates the need for Monte Carlo draws of consumers from the population's characteristics distribution. As noted by Berry, Carnall and Spiller (1996), this proves to be extremely valuable when working with such a large, detailed data set. Such an exact functional form also allows for the value ξ to be solved for analytically via inversion of the market share.

In order to construct the moment conditions used to estimate the parameters we need to define an expression for the vector of unobservable characteristics (i.e. day of the week of the flight, or an airlines service presence at the given airport). Since the e can use the our specification of predicted market share and, for a given set of parameters, solve for each ξ_t as a function of product characteristics, prices, observed market shares and given parameters:

$$\xi_t = s^{-1}(x_t, p_t, st, \theta_d) \tag{15}$$

This vector of unobserved characteristics is extracted by means of a recursive equation; the iterative scheme is initially laid out by BLP (1995) and is known as a contraction mapping method. It is the so-called "inner loop process" in the BLP (1995) methodology that calculates the mean utility. We follow Berry & Jia (2010) and Berry, Carnall and Spiller (1996) in using a contraction map over ξ instead of computing the mean utility. We can make this simplification because, as mentioned above, the BLP (1995) procedure does not allow the inversion of the theoretically predicted market share to be solved analytically, while Berry and Jia (2010) does. The Berry and Jia (2010) modified version of the BLP (1995) contraction mapping procedure uses the following iterative scheme: for each value of our parameters, we compute the subsequent value for ξ as:

$$\xi_{jt}^{M} = \xi_{jt}^{M-1} + \lambda[ln(s)_{jt} - ln(s_{jt}(x_t, p_t, \xi_t, \theta_d)] \tag{16}$$

Where M is the iteration number and $M = 0, 1, ...M_{max}$. M_{max} is the last iteration which stops the minimization procedure once the following condition is fulfilled:

$$max[|\xi_1^M - \xi_1^{M-1}|, ..., |\xi_K^M - \xi_K^{M-1}|] < 10^{-12} \tag{17}$$

Dubé, Fox and Su (2009) numerically go through the convergence of this inner loop (the process by which the market shares are inverted) and stress the need for a stringent convergence tolerance (which we believe we have defined properly) in order to insure that the subsequent outer loop optimization (the actual estimation of the demand parameters) converges appropriately, hence the nearly-zero maximum difference.

3.3 Supply Information: Markups

Because we assume that firms set their prices according to a Bertrand game and that prices are in a static Nash equilibrium, markups can be isolated from the assumed first order conditions. We continue to follow the procedure as put forth by Berry and Jia (2010) when estimating markups, given the demand parameters θ_t. The markups for product j in market t are defined as $b_{jt}(\theta_d, s_t, x_t, p_t)$. The estimated markups are in monetary units and are subsequently used to calculate margins as a percentage of the price. The margins serve as one of the key inputs for our UPP calculations and for our simple logit calibration. Berry and Jia (2010) directly estimate parameters effecting marginal cost as part of there analysis in the form of:

$$\omega_{jt} = p_{jt} - b_{jt}(\theta_d, s_t, x_t, p_t) - w_{jt}\psi \tag{18}$$

As the parameters governing marginal costs are not necessarily of interest to us, we can disregard the marginal-cost shifters w_{jt} and the cost side unobservables term ω_{jt} and extract only the mark-up term $b_{jt}(\theta_d, s_t, x_t, p_t)$. The result is in absolute form, so in order to have it in a form usable for the merger simulation we simply divide mark ups by the fare

$$margins_{jt} = \frac{b_{jt}(\theta_d, s_t, x_t, p_t)}{p_{jt}} \tag{19}$$

3.4 Demand Calibration With A Simple Logit Model

Once the markups from the demand estimation are extracted, we re-calibrate the demand parameters with simpler methods. We use our previously determined market prices, market shares and estimated margins as inputs to simulate the merger with a simple logit merger simulation. Then, using the same elasticities calculated by the logit demand system, we derive our diversion ratios needed to calculate upward pricing pressure. The logit model used by the merger simulation is a less accommodating, less complex version of the discrete-choice model we adopt in order to extract our supply information. It follows the same overall line of thought, but does not use the same amount of information about each product. An important difference to remember is the consumer-specific deviation from the mean utility is independent from product characteristics in the logit model. Subsequently, the patterns of substitutions between products are as well independent of product characteristics.

The mean indirect utility that a consumer in market t obtains from product j is now simplified to

$$\overline{u}_{jt} = \xi_{jt} + \alpha p_{jt} \tag{20}$$

Subsequently, in a given market t with n products, the probability that a consumer purchases product j is given by

$$s_{jt} = \frac{e^{(\xi_{jt}+\alpha p_{jt})}}{\sum_{j=1}^{n} e^{(\xi_{jt}+\alpha p_{jt})}} \tag{21}$$

By again assuming that firms compete á la Bertrand, the logit model yields the own and cross-price elasticities:

$$\varepsilon_{jj} = \alpha(1 - s_j)p_j \tag{22}$$

$$\varepsilon_{ji} = -\alpha s_i p_i \tag{23}$$

In order to estimate the $n+1$ parameters (n number of ξ - one for each products in a market - as well as th price sensitivity of demand, α) our logit model use a system of $2n$ equations, n first order conditions and n choice probabilities. Also in the simple logit we have the opportunity to allow our theoretical consumers to not only have the choice of an inside market product but also to choose an outside product. Our market shares for each firm and product in each market are up to this point based on the methodology of Berry and Jia (2010) by assuming a market size that corresponds to the populations of the two end cities. We note that when we use elasticities to perform the post-merger re-equilibriation, the choice of how we conceptually define the market size and share becomes of new importance. When we estimate demand responses we have no information about this conceptual outside good beside its market share, which is directly linked to the assumed constant market size. If the market size is assumed to be large, which is often the case in the popular estimation techniques as well as in our random-coefficients demand estimation, the outside good is modelled as a popular choice. This popular choice is then likely to capture a fair part of the customers that wants to switch if a product becomes more expensive. Therefore, our own- and cross-product elasticities are more likely to be reasonable if we redefine the market to not include an outside good. Thus, for the simple logit demand calibration, the market shares of the products will sum to one.

When estimating elasticities for use in a merger simulation the econometrician face a trade-off between variance and bias. If there is no assumed or imposed structure on demand, the number of elasticities to be estimated increases exponentially with the number of products in the market. Thus, if no assumptions on structure are made the challenge quickly becomes too great for the data to handle. This is a case where the resulting estimator will have a high variance. The variance is reduced by asking less of the data and imposing a structure (i.e. assumptions) on demand. However, the assumptions may be imposing unrealistic patterns of substitution. In this case, the resulting estimator will be biased. The logit model asks very little of the data it uses, compared to other more complex models. It therefore prioritizes limiting the variance by making bold assumptions.

The simple logit, as we model it, imposes on the patterns of substitution the property of "Independence of Irrelevant Alternatives" (IIA). This implies that the substitution from one brand to all other brands is in proportion to their relative market share. As an example of this property, consider a market with three products: A, B and C, with respective market shares of 60%, 30% and 10%. The assumption of IIA means that if product C raises its price, the substitution that goes from product C to product A must be twice of that to product B. This is because product A has twice the market share of product B. If no other evidence is present, properties of independence of irrelevant alternatives is often viewed as a natural assumption and is therefore often the default (Werden and Froeb, 2002). One consequence of this assumption, and part of the justification of making it, is that any cross price elasticity of demand, with respect to the same price, will be exactly the same. This is a direct result of the fact that cross price elasticities are being modelled to be proportionate to relative market shares.

Even though it is clear to most economists that this is not an assumption that is by any means reflecting the real world perfectly, it is considered a good first step. Unless other reliable evidence exist to show otherwise it is practical and reasonable to assume that the products of the merging firms are neither especially close nor especially distant substitutes. If this is the case, it means that the IIA property approximately holds. Our merger simulations, using the logit model, is relying on this line of thought.

3.5 Price Changes

3.5.1 Predicted Price Changes

Our results attempt to compare three separately calculated price changes: that which is predicted from the logit simulation, that which is approximated by the UPP, and the actual observed post-merger price change. The UPP price changes comes from simply inserting the relevant estimated margins and diversion ratios, which are derived from the logit calibration described above and are defined by equation (7).

We assume no efficiency gains for this exercise for neither the logit predictions nor the UPP predictions. In order to compare the UPP to the merger simulation predicted price changes we assume a pass-through rate of one and we simply divide by the price pre-merger to get a percentage unit. The predicted price increase implied by the UPP for a merger between firm i and firm j can thus be written as:

$$\%\Delta p_j = UPP_j = \frac{(p_j - mc_j)D_{j,i}}{p_j} \tag{24}$$

After recovering supply and demand estimates, the last step predicting price changes is to execute the merger simulation. From from our simple logit demand calibration, we estimate the system of elasticities that determines the substitution patterns between all products within a given market. The firms are assumed to compete á la Bertrand

and we can therefore back out the first order conditions. For each market, the first-order conditions, together with the estimated demand functions, define a system of nonlinear equations in price. The first order condition for product j is defined by (2). This equation defines each firm's behaviour, as it is the structure that lead to the pre-merger equilibrium. We can then manipulate the ownership structure that governs the set of products belonging to each firm. By simply stating to the model that the product in each market that was previously maximized over by America West now falls under the portfolio of US Airways, and then letting the market re-equilibrate, we can extract the new post-merger price. It should be noted, especially for the comparison to the observed ex-post price increases, that a merger simulation in this manner is only trying to forecast the effect from a loss of competition. A merger could potentially affect prices through other mechanisms (efficiency gains through economies of scale, or increased ability to tacitly collude with competitors, as two examples), however our predicted price change measures do not take such factors into account.

3.5.2 Actual Prices Changes

Measuring a price change for firm j in market t from time $t = i$ until time $t = x$ is straight forward enough:

$$\%\Delta p_{jt,x} = \frac{p_{jt,x} - p_{jt,i}}{p_{jt,i}} \tag{25}$$

In measuring these actual observed price changes we follow the likes of (as one example) Peters (2003), in that we only attempt to measure the direct price effects of the merger. We do not take into account any dynamic changes in prices or long term effects on competition or industry structure (besides of course the initial merger). Ignoring these makes sense with regards to the fact that the two predictive models also make the same assumption (or at least do not incorporate it). It is common knowledge however, that seasonal demand is extremely prevalent in the airline industry, so prices in quarter 1 are not directly comparable to quarter 2. Air travel in general is also characterized by its own measure of inflation, which is weighted towards the cost of fuel and other industry-specific cost drivers. Thus, in order to account for seasonality and longer term inflation in the industry, we define a measure of price changes called the abnormal price change, where the abnormal price change for airline i is

$$\Delta p_{jt,x}^{abn} = \%\Delta(p_{jt,x} - \overline{p_t}) \tag{26}$$

where $\overline{p_t}$ is the average fare in market t. We are essentially taking the difference between the merging firms price change over time and the price change of each of the markets average fare. In doing this we hope to net out those factors that shift the aggregate airline supply curves, leaving only the merger price effects. For example, if fuel prices go up, all firms should be effected equally and prices will adjust upward accordingly. However, under the assumption that no single airline firm can internalize the increased fuel costs more efficiently than any other, we should see the effects of this increase netted out when comparing the abnormal merger increase to the market average.

4 Data

4.1 Data Sources

Our primary data source is the Airline Origin and Destination Survey (the DB1B dataset) as published by the US Department of Transportation's (DOT) Bureau of Transportation Statistics[9]. This publically available data set is published quarterly, and

[9] `http://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=125`

consists of a ten percent random sample of all airline tickets purchased by reporting US-based airline carriers. Within each entry, many characteristics of the tickets purchased can be retrieved, including itinerary fare, flight distance, airline flown with and whether or not the consumer purchased a round trip ticket. In addition, we use data collected by the US Census to provide estimates of the 2005 population levels for each of the relevant Metropolitan Statistical Areas (MSAs) in our estimation procedure[10].

4.2 Sample Selection

The raw data from the DB1B comes in 3 separate sets, each of a differing level of aggregation. Thus the data must be filtered, condensed, and combined based on itinerary identification numbers that persist across the data sets. This procedure is tedious, but allows us to form rich descriptions of each flight purchased; specifically, we can identify the firm providing the tickets, the price paid by the customer, and the market of the airports that are flown between. Our initial estimation procedure is based on data from the first quarter of 2005, as this was the last complete quarter before the merger between America West Airlines and US Airways was initially announced. Because of the significant seasonal effects which exist in the airline industry, our ex-post analysis is also directed at quarter 1 of the subsequent year.

Since our estimation procedure follows Berry and Jia (2010) nearly to a t, our data treatment also closely tracks theirs. We begin our analysis by establishing the notion of a geographic market, as is commonly used in the airline industry literature. A market is a direction-dependent pair of cities. The direction dependency is important because it allows for characteristics of consumers and products in the origin city to be accounted for in the estimation procedure. For example, travelers flying from New York City to Las Vegas and back will likely be travelling for different reasons, and therefore have

[10] `https://www.census.gov/popest/data/historical/2000s/vintage_2005/metro.html`

different preferences, than those flying from Las Vegas to New York City and back. We follow Benkard, Bodoh-Creed and Lazarev (2010) in limiting our analysis to only the top 75 busiest airports in the United States[11]with the addition of Long Beach airport (LGB) in the Los Angeles region; Berry and Jia support such treatment by saying that the benefit of adding more airports and thus more data is outweighed by the additional computation time and the limited products offered within smaller markets. Long Beach is included due to personal knowledge of the Los Angeles market and its role as a substitute for Los Angeles International Airport (LAX). In fact, Los Angeles is one of nine MSAs in our study containing more than one airport.[12] These groups of airports can be considered as offering competing products, as it is likely an individual flying from Los Angeles (as one example) will search for the best product from all surrounding airports rather than just one. The degree of substitutability betThese 76 airports thus fall into a total of 60 unique metropolitan statistical areas, which altogether form 3355 unique origin-destination markets for analysis.

We continue by removing products which do not represent round trip flights; one-way flights are fairly infrequent in the data (cite infrequency) and would possibly represent a wholly separate consumer type and preference set. Products with itinerary fares less than 40 dollars are dropped. Products where the ticket-issuing airline carrier changes mid-itinerary are dropped, along with products where any portion of the journey does not take place via air travel (a bus connection, for example). Any journey with more than four flights (i.e. two stopovers) is dropped, and any itinerary where the DB1Bs Dollar credibility indicator is false is dropped. This short list of conditions, as well as the restriction in airports used, results in the removal of about 1.5

[11]A comprehensive list of these airports and their corresponding MSA can be found in Appendix I

[12]The metropolitan statistical areas containing more than one airport are: Los Angeles (LAX, SNA, ONT, BUR, LGB), San Francisco (OAK, SFO), New York City (LGA, JFK, EWR), Chicago (MDW, ORD), Houston (HOU, IAH), Dallas-Fort Worth (DFW, DAL), Washington DC (BWI, DCA, IAD), Miami (MIA, FLL), and Boston (BOS, MHT, PVD)

million itineraries, or approximately 60 percent of the original sample. As noted by Berry and Jia (2010), however, this still leaves one with a computational problem; the computational requirements needed for the nonlinear optimization estimation over the remaining 963000 airline products is simply too much. Thus, we follow their lead in developing fare ranges for each product. If a group of ticket fares fall within a specified range, the prices are averaged out and the passengers are summed together. We believe that Berry and Jias justification for doing so is appropriate: consumers will not view purchases with slightly different prices ($428 versus $436) as distinguishable products. Working with this assumption allows us to aggregate our data down to 289,761 distinct products representing itineraries for nearly 3,000,000 passengers. The size of the market is calculated as the geometric mean of the origin MSA population and the destination MSA population. While this does not provide necessarily an precise estimate of the market size in absolute terms, it does allow for the size of the destination and the demand for high-population destinations to enter the market share equation. This is in comparison to simply using the origin population as the market size, which would result in all destinations having the same market size, regardless of destination characteristics. A simple example to illustrate the reasoning for this is to consider a small city such as Albuquerque (ABQ), New Mexico. There will most likely be much higher demand for flights from ABQ to large cities such as Chicago than to smaller cities such as Albany, New York. Using the geometric mean of the origin and destination populations allows us to proxy for these differences. Finally, to produce a market share term, the total passengers within each products fare bin is divided by the market size.

4.3 Data Treatment, Beyond The Demand Estimation

Once cost-side estimates are recovered from the Berry and Jia (2010) logit estimation procedure, we construct a data matrix containing once again the market and carrier

	Total Sample	US Airways	America West
Number of Observations	288876	21885	18802
Number of Unique Markets	3355	1519	1123
Number of Unique Origins	60	52	46
Number of Unique Destinations	60	58	47
Average Market Size, Millions	1.682	2.816	3.432
	(0.793)	(1.624)	(2.340)
Average Number of Carriers Per Market	6.10	7.59	8.55
	(1.91)	(2.34)	(2.96)
Average Fare ($)	435.29	397.81	430.69
	(293.85)	(281.45)	(226.58)
Average Carrier Share, Origin (%)	-	13.21	5.15
		(6.13)	(2.44)
Average Carrier Share, Destination (%)	-	13.27	10.46
		(6.13)	(8.24)

Table 1: Here we present a summary of some of the relevant variables used in the random-coefficients demand estimation. The second and third columns are the subsets of the data representing only US Airways and America West

indicators, the fare information, the margin estimations (which are in absolute dollar terms), the margins as a percentage of the fare, and the market share of the products. We identify markets which our carriers of interest (America West and US airways) both serve, resulting in an abbreviated data set of 296 markets with such an overlap, containing in all about 65000 products (by the aformentioned definition). Things get a bit tricky here however, since our products are defined by the fare ranges we have described above. If we were to conduct a merger simulation based on these product definitions, we would have to adjust our ranges to account for the fact that some products would be increasing to a point where they would no longer fall in their initial fare range (i.e. a $418 dollar fare increases by 1 percent, putting it into a completely seperate $420 fare bin). In order to avoid this problem and the obvious data complications it creates, we resort to aggregating the products into one entry, where margins and fares for each carrier are passenger-weighted averages over all their products in the given market. This simplification takes away our ability to determine how the entire distribution of fares in a market changes, however we feel the results are still valid and valuable on this aggregated level. The market shares can be simply summed together, since they are already passenger-weighted by definition. The data at this point is ready for a merger simulation.

4.4 Ex-Post Data Treatment

Our final step in working with the raw data is to describe the actual price changes in the markets of interest over a period of time. We look at quarter 2 in 2005, quarter 1 in 2006, and quarter 1 in 2007; we expect an incremental price increase in the first quarter after the merger (2005, Q2) and a further, more measurable price jump in the first quarter of 2006, with prices increasing slightly more in the final quarter of our study. Data from the DB1B data source cited in section 4.1 is again collected for the 3 time periods mentioned above; the data does not have to be as descriptive as that which goes into our demand estimation, so only information on the geographic market, fare, carrier, and exclusion criteria as described in 4.2 is collected. This data undergoes the same treatment as the 2005, quarter 1 data with regards to filtering and aggregating.

5 Results

5.1 Demand Estimation Parameters And Markups

The results of the random coefficients logit demand estimation are presented in table 2. While the estimated parameters are not crucial to our merger simulation procedures per say, it is important that our estimated parameters make economic sense and properly reflect what we expect is occurring in the airline industry. With this in mind, our margin estimation procedure is validated, as our results are all coherent with the expected economic intuition. Consumers decrease quantity with increases in price, and the two fare coefficients indicate that there are indeed two distinct consumer types with differing price sensitivities; type one represents the leisure travelers and type two represents the less-price-sensitive business type of consumers. The negative sign on this parameter confirms that the customers indeed dislike higher prices. Similarly, and as expected, the utility of the consumers decrease with the number of connections. As Berry and Jia (2010) discuss in their analysis of the airline industry, demand for air travel is believed to have a U-shaped relation with distance. The short flights within the continent compete with cars and trains which is captured by the outside good. These alternative options becomes worse substitutes to flying however as distance grow, which means that demand increases with distance initially. At a certain point, further distance doesnt exclude anymore outside options and longer distance only becomes a burden for the traveller, hence demand starts to decrease with further increase in distance. This relationship is also captured in our demand estimation which can be seen by the positive sign for distance combined with a negative sign for the square of distance. Overall, our estimation produces parameters which are similar in magnitude to those of Berry and Jia (2010), who perform the same procedure but in 2006 rather than 2005. It also corroborates with Cheung (2013) who estimates demand in a similar fashion over a two-year period up to the same point in time as us. Our estimated price sensitivity for the leisure traveller is -0.798 which is combarable with -0.77 for Cheung (2013) -1.05 for Berry and Jia (2010) The more sensitive business travellers have an estimated parameter of -0.057 which is to compare with -0.12 for Cheung (2013) and -0.10 for Berry and Jia (2010) Our parameters are very close in magnitude to both of these studies. All the other parameters have the intuitive sign and are statistically significant. We estimate the share of type 1 consumers, i.e. leisure travellers, to be

	Type I	Type II
Fare	-0.798	-0.050
	(.0156	(0.0014)
Connection	-0.576	-0.502
	(.0156)	(0.0236)
Constant	-4.382	-7.087
	(0.1648)	(0.2973)
Distance	.0914	
	(.0235)	
Distance Squared	-.0021	
	(.0025)	
Lambda	0.558	
	(0.011)	
Gamma	0.695	0.315
	(0.0967)	(0.0967)

Table 2: Demand parameters as estimated by the random-coefficients logit model. Type I consumers are the leisure type, while type II consumers are the business type.

70% of the consumers. It does not seem unreasonable that the share of the more price sensitive consumers make up a larger part.

The critical output from the structural demand estimation procedure is the vector of estimated markups and subsequently calculated margins of the airline products. Margins are calculated as the estimated markup as a percentage of the fare, and then a market specific, passenger-weighted average margin is presented for each carrier. In order to better fit the methodology used for the simulations we collapse the data into only one product per airline. It also makes the comparison with actual price changes more intuitive and informative. The margins that are extracted from the structural demand estimation are derived from a model of Bertrand competition. This is an advantage because Bertrand competition is a fundamental assumption behind the simple logit model used for the merger simulation, and while data on margins could in practice

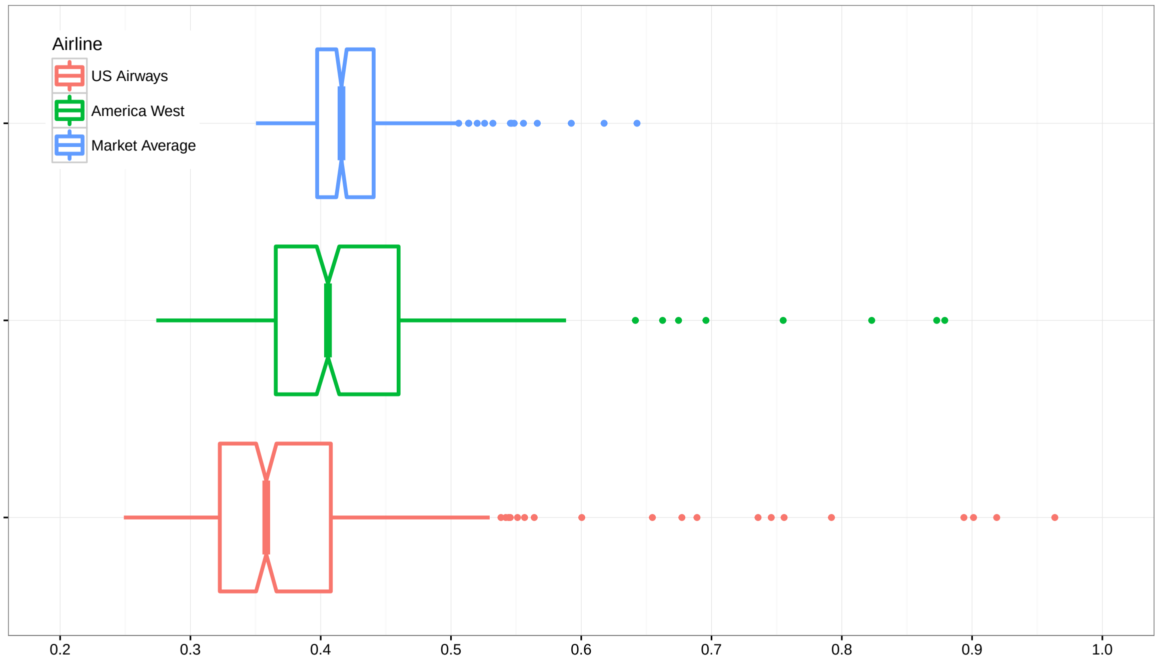

Figure 4: Summaries of the margins (calculated as a percentage of the price) used as inputs in the merger simulations.

be supplied by firms themselves or collected from accounting data, such information would not necessarily be in a form that is coherent with the academic definitions of the markets or the assumptions about preferences we have made. The summary statistics for the estimated margins used in the logit calibration is shown in table 3. The fact that these margins derived from the structural demand estimation are seemingly high for an industry which is believed to have been in financial distress are initially concerning. However, our results do corroborate with previous empirical estimations of margins in the airline industry such as, for example, estimations of market conduct by Fischer and Kamerschen (2003) and a study on endogenous costs by Neven, RÖLLER and Zhang (2006). Moreover, since our estimation procedure produces demand parameters of reasonable and expected magnitudes, and the margins are calculated as intermediate steps within the procedure, we have no reason not to trust the derived margins.

Airline	Market Count	Margins	Standard Deviation (Margins)
America West	296	0.421	0.086
US Airways	296	0.383	0.107
United	295	0.452	0.085
American	293	0.438	0.071
Delta	289	0.473	0.096
Northwest	282	0.432	0.093
Continental	262	0.417	0.090
Southwest	223	0.425	0.099
ATA	163	0.378	0.076
Frontier	114	0.365	0.079

Table 3: Margin summaries for the top ten most frequent airlines in our initial 296 markets in which both America West and US Airways operated.

5.2 Merger Simulations And UPP

In this section we compare the predicted prices changes as computed from the UPP measurement against the ones predicted by the logit merger simulations. For each market observation, we use the margins as derived by the random-coefficients logit demand estimation, the simple logit demand elasticity matrix, and the subsequently calculated diversion ratios. We can then simulate the merger, extract the predicted price changes for US Airways and America West Airways products and calculate the UPP equivalent. We note that our UPP calculations resemble the logit predictions per design, as they use the same margins as the merger simulation, as well as the same substitution patterns. However, the difference between the methodologies lies in the fact that the merger simulations also takes the other firms' competitive re-equilibration into account, while the UPP ignores this.

As discussed in this thesis, diversion ratios are in practice often estimated by surveys and not by estimating demand via readily available data. Similarly, margins can be collected simply by using public accounting data as opposed to marginal cost estimations derived from econometric techniques. If we were to compare UPP statistics to merger simulation results using information that is not derived from the same demand estimation, it would in almost any case result in the UPP and merger simulation deviating further from each other. Any such comparison and results would be very difficult to generalize since the data collection and estimation of diversion ratios differ on a case-by-case basis.

Our scenario does in fact have practical implications because the logit model used to do the simulations does not require much more information than the UPP does. Unlike the most advanced demand models, the simple logit serves as a reasonable alternative to the UPP. In order to contribute to the ongoing debate of the usefulness of the UPP measure in serving as a predictor of the magnitude of price increases we compare the UPP predictions to a viable alternative from the set of merger simulations available.

With no cost efficiency gains for the merging firms considered, the UPP and the logit predictions are almost identical. The correlation between the two variables is close to one. The UPP predicts slightly higher price increases compared to the Logit model, which is to be expected due to the pass-through approximation. This result has several important implications. It implies that the assumption of full pass-through is not a severe source of error and that indeed it does not hinder the UPP from being used to indicate the magnitude of predicted price changes resulting from a merger. This subsequently implies that the critique saying that the UPP is not actually related to price changes and thus is detached from measures of consumer harm is to some extent rejected by this result. If a merger simulation based on a simple logit demand model serves as a measure of price changes, then the UPP can produce almost the same result. As we have demonstrated, the UPP builds on the same fundamental theory as a merger simulation, and while new sophisticated econometric techniques can make the

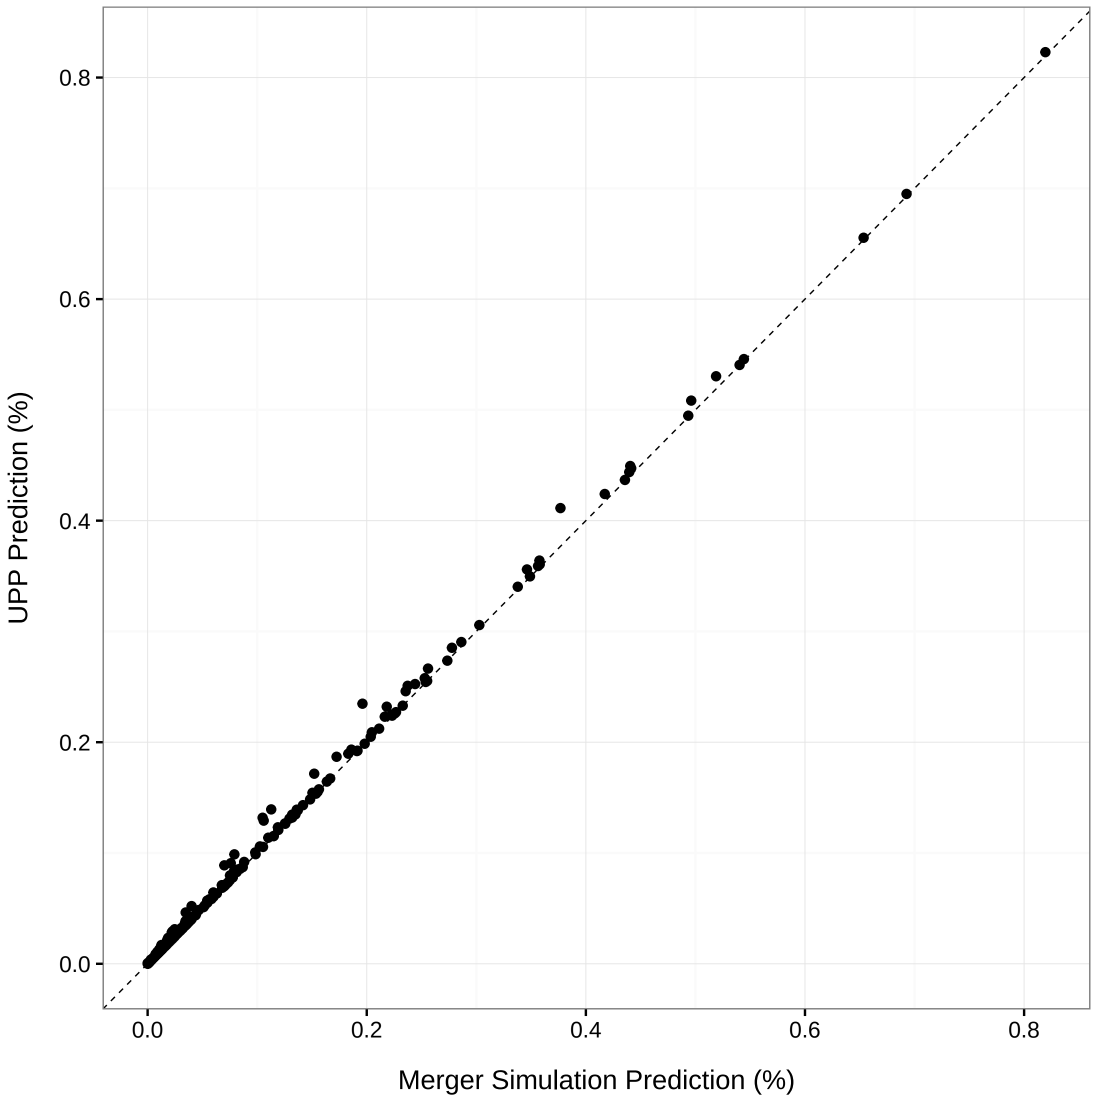

Figure 5: When plotted against each other, one can see that the relationship between the two predictions is nearly perfectly correlated ($\rho = .999$).

	Average Logit Prediction (%)	Average UPP Prediction (%)	Average Difference	Correlation, Logit & UPP
US Airways	8.49 (12.87)	8.71 (13.03)	0.22 (0.49)	.999

Table 4: Summary of the price change predictions and the differences between the two, with standard deviations indicated in parentheses.

merger simulations more flexible, complex and detailed, the UPP output is still similar to the results of a simple merger simulation, in cases like this. Further, one important implication is that the UPPs ability to predict price change lies almost entirely on the input data that it uses. The debate should therefore be more focused on how, in practice, authorities can collect data and estimate margins and diversion ratios in a credible, robust manner that mimics the theoretical assumptions as closely as possible.

In order to give a more detailed picture of how the margins and fare inputs are used, as well as how the results are formatted, we can zoom in on a few specific markets which exhibit particular results. First, we take the market for round-trip flights from Cleveland to Las Vegas, one which has presented us with fairly average predicted price increases for both firms; 7.16% for US Airways via the merger simulation and 7.19% via the UPP measure (compared with aggregate mean price increases of 8.49% and 8.70% respectively).

The simple logit calibration and subsequent merger simulation input data is presented in Table 5, with the final column showing the merger simulation predicted price increase for the given product.

We can see that in this market, US Airways was seemingly distressed. With the lowest margins of all operators in this market and a small market share, US Airways was under pressure. America West had a somewhat larger presence in this market, with a 14% market share. This was in fact a home market for Continental Airways, who

Carrier	Market Price	Market Share	Margins (as % of price)	Predicted Price Change (%)
US Airways	$260.7	0.007	33.8	7.166
America West	$443.72	0.142	40.4	0.185
American	$329.75	0.008	35.4	0.001
Continental	$584.13	0.637	51.6	0.076
Independence Air	$168.00	0.001	43.8	0.125
Delta	$276.96	0.010	41.1	0.002
Northwest	$397.08	0.006	44.6	0.001
United	$353.33	0.005	40.7	0.001
Southwest	$393.07	0.184	42.2	0.030

Table 5: Logit/merger simulation input values for the Cleveland-Las Vegas route within our merger simulation

showed a market share of almost 64%. This means, for our predicted merger outcomes, that US Airways (post-merger) have some room to raise prices, since America West would recover some of these lost profits. However most of the shifting customers would likely find their ways to the big player in this market, Continental, and the other firms are left fighting for the left-overs. This is also reflected in the diversion ratios, shown in Table 7, where Continental recovers more sales than any other firm.

We can compare the Las Vegas-Cleveland market with a market presenting more statistically outlying results. The route between Philidelphia and Phoenix is one such case, with merger simulation inputs and results presented in Table 6. The UPP estimate for this market is 3.45 percentage points higher than the logit merger simulation at 41.1% - economically and statistically a very significant difference.

Both the UPP and the Logit gives high predictions of price changes. Examining the diversion ratios and elasticities for US Airways and America West shows why. Diversion from US Airways to America West is much higher in the Philidelphia-Phoenix market than in the Las Vegas-Cleveland market. Similar trends exist in the elasticities; ratios

Carrier	Market Price	Market Share	Margins (as % of price)	Predicted Price Change (%)
US Airways	489.07	0.219	45.4	37.768
America West	538.75	0.507	55.6	13.032
American	477.79	0.016	53.4	0.276
Continental	698.75	0.005	82.1	0.058
Delta	183.00	0.007	49.5	0.332
Frontier	213.00	0.002	35.3	0.094
Northwest	376.57	0.025	34.7	0.571
ATA	214.00	0.001	35.1	0.045
United	459.12	0.108	44.9	2.170
Southwest	395.33	0.110	41.4	2.551

Table 6: Logit/merger simulation input values for the Phoenix-Philidelphia route within our merger simulation

of the cross-price elasticity to the own price elasticity show that the two airline products have a higher degree of substitutability, and therefore capture a greater share of the market participants after a merger in the latter case compared to the former. In fact, this was one of America Wests home markets with the highest market share of 50%. This means of course, with the IIA properties of the Logit, that America West is expected to capture large parts of those consumer who wants to switch after a price increase from US Airways.

The difference between the two predictions for this market can have various explanations, but the most important factor is likely to be the theoretical differences that underlies the them. It is noticeable that the Logit model predicts substantial price increases for other airlines, suggesting that the re-equilibration of the competing firms play a larger role to hold down the price raise from US Airways, while the UPP does not consider this.

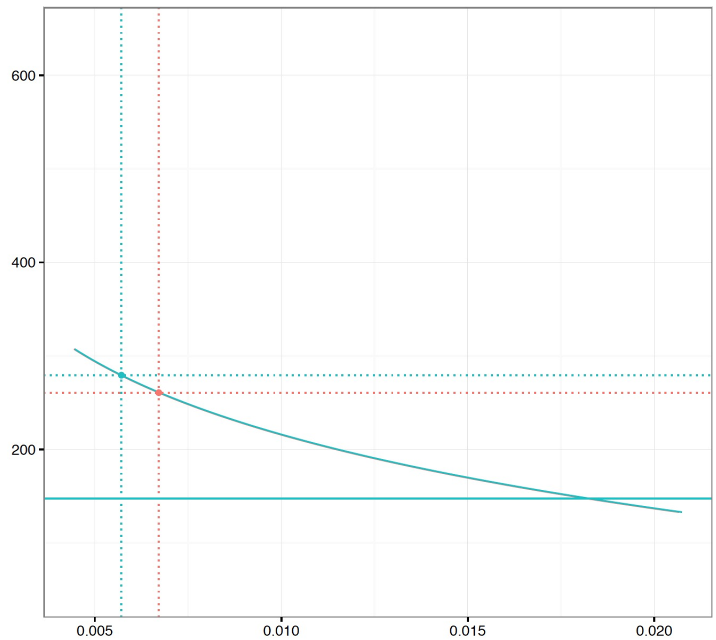

Figure 6: Demand for the Cleveland-Las Vegas US Airways product as modeled by the simple logit model, both pre-merger (red) and post-merger (blue) demand curves.

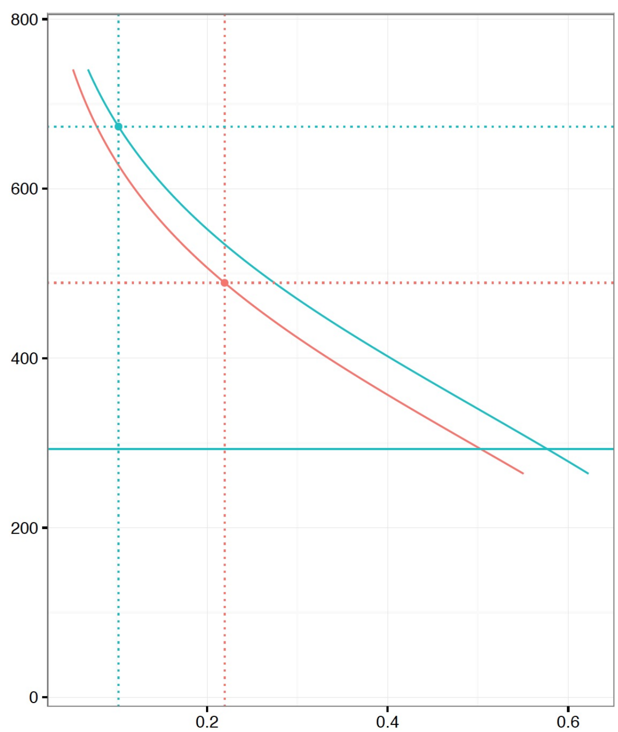

Figure 7: Demand for the Phoenix-Philadelphia US Airways product as modeled by the simple logit model, both pre-merger (red) and post-merger (blue) demand curves.

	US Airways	America West	American	Continental	Independance Air	Delta	Northwest	United	Southwest
US Airways	-1	0.1433	0.0081	0.6409	3e-04	0.0104	0.006	0.0055	0.1856
America West	0.0078	-1	0.0093	0.7422	3e-04	0.012	0.0069	0.0063	0.215
American	0.0068	0.1435	-1	0.6418	3e-04	0.0104	0.006	0.0055	0.1859
Continental	0.0185	0.3916	0.022	-1	7e-04	0.0284	0.0163	0.0149	0.5075
Independence Air	0.0067	0.1423	0.008	0.6368	-1	0.0103	0.0059	0.0054	0.1844
Delta	0.0068	0.1438	0.0081	0.6433	3e-04	-1	0.006	0.0055	0.1863
Northwest	0.0068	0.1432	0.0081	0.6404	3e-04	0.0104	-1	0.0055	0.1855
United	0.0068	0.1431	0.008	0.6401	3e-04	0.0104	0.006	-1	0.1854
Southwest	0.0082	0.1745	0.0098	0.7806	3e-04	0.0127	0.0073	0.0066	-1

Table 7: Diversion ratios for each product in the Cleveland-Las Vegas market space

	US Airways	America West	American	Continental	Independance Air	Delta	Northwest	United	Southwest
US Airways	-2.308	0.563	0.024	3.314	0.001	0.026	0.021	0.017	0.646
America West	0.016	-3.392	0.024	3.314	0.001	0.026	0.021	0.017	0.646
American	0.016	0.563	-2.915	3.314	0.001	0.026	0.021	0.017	0.646
Continental	0.016	0.563	0.024	-1.892	0.001	0.026	0.021	0.017	0.646
Independence Air	0.016	0.563	0.024	3.314	-1.497	0.026	0.021	0.017	0.646
Delta	0.016	0.563	0.024	3.314	0.001	-2.443	0.021	0.017	0.646
Northwest	0.016	0.563	0.024	3.314	0.001	0.026	-3.518	0.017	0.646
United	0.016	0.563	0.024	3.314	0.001	0.026	0.021	-3.132	0.646
Southwest	0.016	0.563	0.024	3.314	0.001	0.026	0.021	0.017	-2.857

Table 8: Own and cross-price elasticities as estimated by the simple logit model for the Cleveland-Las Vegas market space

	US Airways	America West	American	Continental	Delta	Frontier	Northwest	ATA	United	Southwest
US Airways	-1	0.6489	0.0198	0.0061	0.0092	0.0031	0.0321	0.0015	0.1389	0.1405
America West	0.4445	-1	0.0314	0.0097	0.0145	0.0048	0.0507	0.0024	0.2198	0.2222
American	0.2228	0.5145	-1	0.0048	0.0073	0.0024	0.0254	0.0012	0.1101	0.1114
Continental	0.2204	0.509	0.0156	-1	0.0072	0.0024	0.0251	0.0012	0.109	0.1102
Delta	0.2209	0.5102	0.0156	0.0048	-1	0.0024	0.0252	0.0012	0.1092	0.1104
Frontier	0.2198	0.5078	0.0155	0.0048	0.0072	-1	0.0251	0.0012	0.1087	0.1099
Northwest	0.2249	0.5196	0.0159	0.0049	0.0073	0.0024	-1	0.0012	0.1112	0.1125
ATA	0.2196	0.5072	0.0155	0.0048	0.0072	0.0024	0.0251	-1	0.1086	0.1098
United	0.246	0.5682	0.0174	0.0053	0.008	0.0027	0.0281	0.0013	-1	0.123
Southwest	0.2463	0.569	0.0174	0.0054	0.008	0.0027	0.0281	0.0013	0.1218	-1

Table 9: Diversion ratios for each product in the Phoenix to Philadelphia market space

	US Airways	America West	American	Continental	Delta	Frontier	Northwest	ATA	United	Southwest
US Airways	-2.496	1.784	0.048	0.022	0.009	0.003	0.062	0.002	0.325	0.283
America West	0.701	-1.738	0.048	0.022	0.009	0.003	0.062	0.002	0.325	0.283
American	0.701	1.784	-3.075	0.022	0.009	0.003	0.062	0.002	0.325	0.283
Continental	0.701	1.784	0.048	-4.546	0.009	0.003	0.062	0.002	0.325	0.283
Delta	0.701	1.784	0.048	0.022	-1.188	0.003	0.062	0.002	0.325	0.283
Frontier	0.701	1.784	0.048	0.022	0.009	-1.389	0.062	0.002	0.325	0.283
Northwest	0.701	1.784	0.048	0.022	0.009	0.003	-2.4	0.002	0.325	0.283
ATA	0.701	1.784	0.048	0.022	0.009	0.003	0.062	-1.397	0.325	0.283
United	0.701	1.784	0.048	0.022	0.009	0.003	0.062	0.002	-2.676	0.283
Southwest	0.701	1.784	0.048	0.022	0.009	0.003	0.062	0.002	0.325	-2.301

Table 10: Own and cross-price elasticities as estimated by the simple logit model for the Phoenix-Philidelphia market space

5.3 Comparison To Observed Ex-Post Price Changes

In this section we compare the predicted price changes to the patterns of actual price changes in each market. We acknowledge that isolating the effect from the merger is by no means a trivial task and our simple approach might not give the full picture.

One strategy for measuring the unilateral price effects of mergers is to produce a controlled quantification of an average, or aggregated effect from a merger over many distinct markets. However, due to the structure of the airline industry and the assumptions that the prominent literature makes, as a starting point, we want to compare each of our predictions to their equivalent real world outcome and not by looking at aggregated results. The reason is simple: our predictions treat each market as individually independent and the theory behind them say nothing of how firms treat profit maximization of multiple market spaces. Under the assumption that each firm is profit maximizing each market independantly of the other ones, we can view our predictions as a single merger in every overlapping market. This means, however, that we only have one observation per overlapping market and our comparison boils down to a simple comparison of differences, without much possibility to isolate the merger-effect from any other explanatory factors that might have an effect on the price from one period to the next. In an attempt to better reflect the price change that results from the merger in each market we take the abnormal price change. We calculate the difference between US Airways price changes and the average change in each market. This method attempts to cancel out any factors that affects all airlines in a given market in a similar way, such as inflation or a change in airport costs.

As we demonstrated in the previous section, the UPP and the simple logit merger simulation produce almost identical predictions when the inputs are derived from the same source. Thus, we focus on comparing the predictions from only the merger simulation to the actual price changes. It follows then that the UPP is either equally good or equally bad as the merger simulation predictions.

Figures 8, 9 and 10 plot the predictions of the merger simulations against the actual abnormal price changes measured as the passenger weighted average of itinerary fare for US Airways minus the similarly calculated market average. The figures include fitted trendlines as well as dashed lines from the origin with a slope of one (i.e. a line representing zero variation between the merger simulation predictions and the observed price changes). We can see that for many markets the merger simulation does not predict well the actual outcome. From one quarter to the next there is almost no statistical relationship between the predictions and the actual outcome. After one year there is a positive relationship, however it is still weak and it is difficult to say if it has any economic interpretation or not. Two years after the merger, we see some correlation between the predictions and the observed outcomes, however the deviations are substantial and the magnitude of the difference between the observations and their predicted counterparts is large.

Table 11 shows the aggregate results for all markets and compares the merger simulation and the UPP to the ex-post data. The average abnormal price change between overlapping markets from the first quarter of 2005 to the second quarter of 2005 was 13.96 percent, a figure which is comparable in magnitude to the average predicted price increase from the logit merger simulation (8.49 percent). For multi-market mergers, predictions or measurements that are used in antitrust cases are often an aggregation measure of some sort, and with that in mind it is tempting to think that the predictions of the merger simulations were in fact fairly accurate. However, the standard deviation of the abnormal price change in the second quarter of 2005 was above 40 percent, while the standard deviation of the merger simulation predictions was just under 13 percent. Coupled with an average difference between the abnormal price increases and the observed price changes of 11.96 percentage points for the same period, and one can

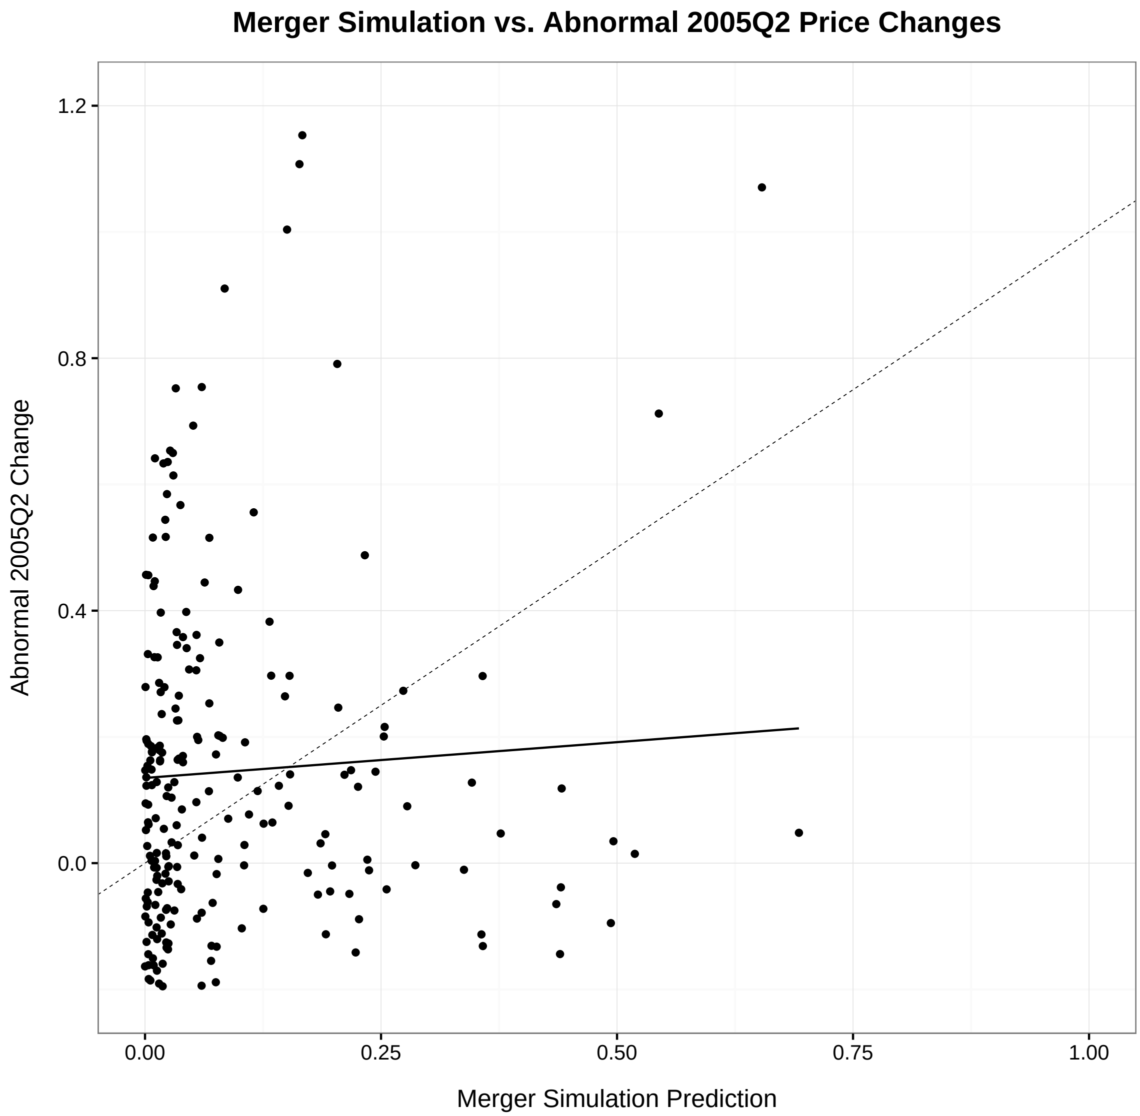

Figure 8: There is nearly no relationship between the merger simulation predictions and the abnormal price changes in the second quarter of 2005 ($\rho = 0.043$)

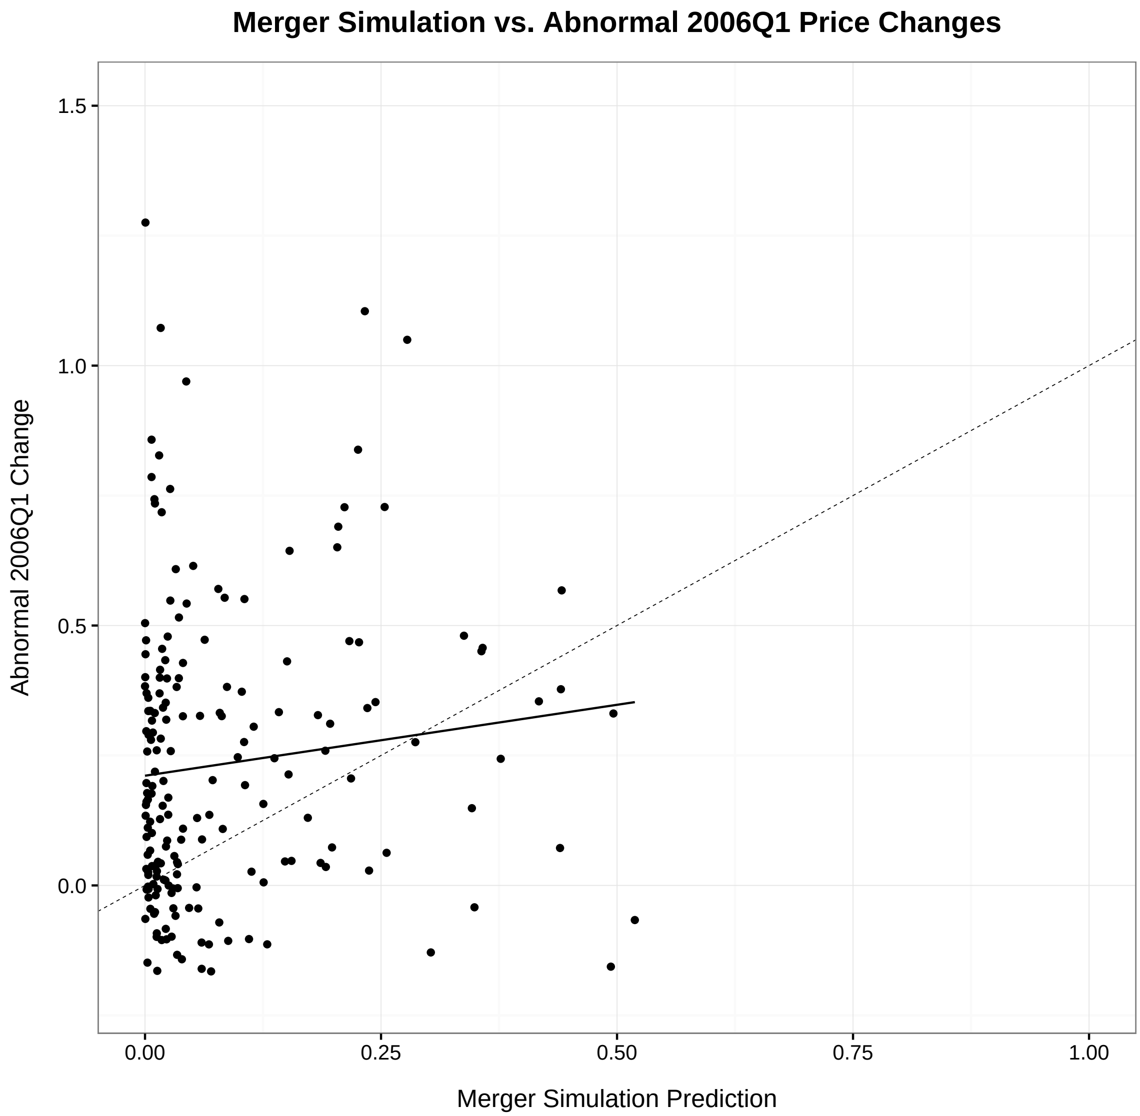

Figure 9: The relationship between the merger simulation predictions and the abnormal price changes is stronger than it was in the second quarter of 2005, however it is still very weak ($\rho = 0.101$)

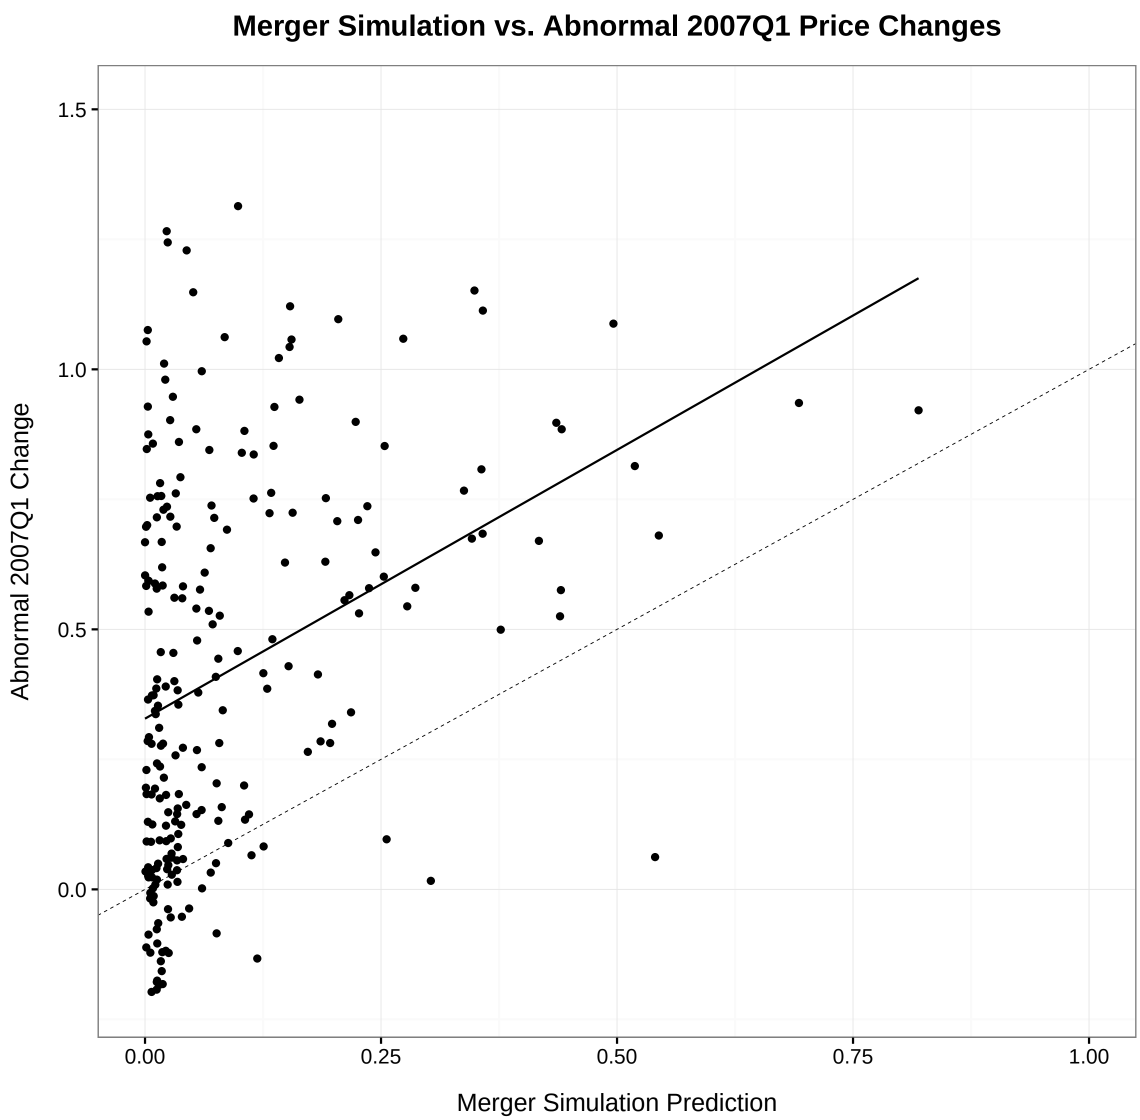

Figure 10: The merger simulation predictions and abnormal price changes finally exhibit a stronger relationship in the first quarter of 2007 ($\rho = 0.337$). However we can only speculate as to the cause of this.

see that the variance of the results is far too high to be able to draw conclusions from. With seemingly no statistically significant relationship between the predictions and the outcome across markets, the results prove to be very weak and do not give strong support for the predictive accuracy of these tools. Similar trends emerge from the analysis of 2006 observed price changes. While there is a slightly higher correlation coefficient, the overall spread of the abnormal price changes and the differences between the predictions and the actual outcome is too wide to draw conclusions from. Interestingly, the abnormal price changes in 2007 and the merger-simulated price changes become increasingly correlated, albeit to a point that is statistically still considered fairly weak. 2007 also happens to be the first year where America West products completely disappeared from the market after operations were integrated with US Airways, but whether or not the increase in predictive accuracy occurs due to the absence of America West or just coincidentally is unknown.

Our results do not provide a strong indication that the merger simulation method accurately predicts the price changes over many markets. However, this result has to be interpreted with great care. The fact that the merger simulation does not, as a general rule, accurately predict the actual outcome is an indicator that these simple tools are not working very well. It is, at first sight, quite straightforward to recognize that the antitrust tools are trying to forecast what would happen, and the actual outcome was simply not the case. Even if the aggregated results can give indications that the tools are right, on average, the differences from market to market are too large to draw definitive conclusions from. However, several factors hamper us from drawing any strong conclusions about the accuracy, or lack thereof, for the predictions.

We want to, before dismissing the use of the screening tools for antitrust purposes, humbly recognize that we might not be comparing the equivalent price changes to each other. The merger simulation and the UPP only takes the changes in market

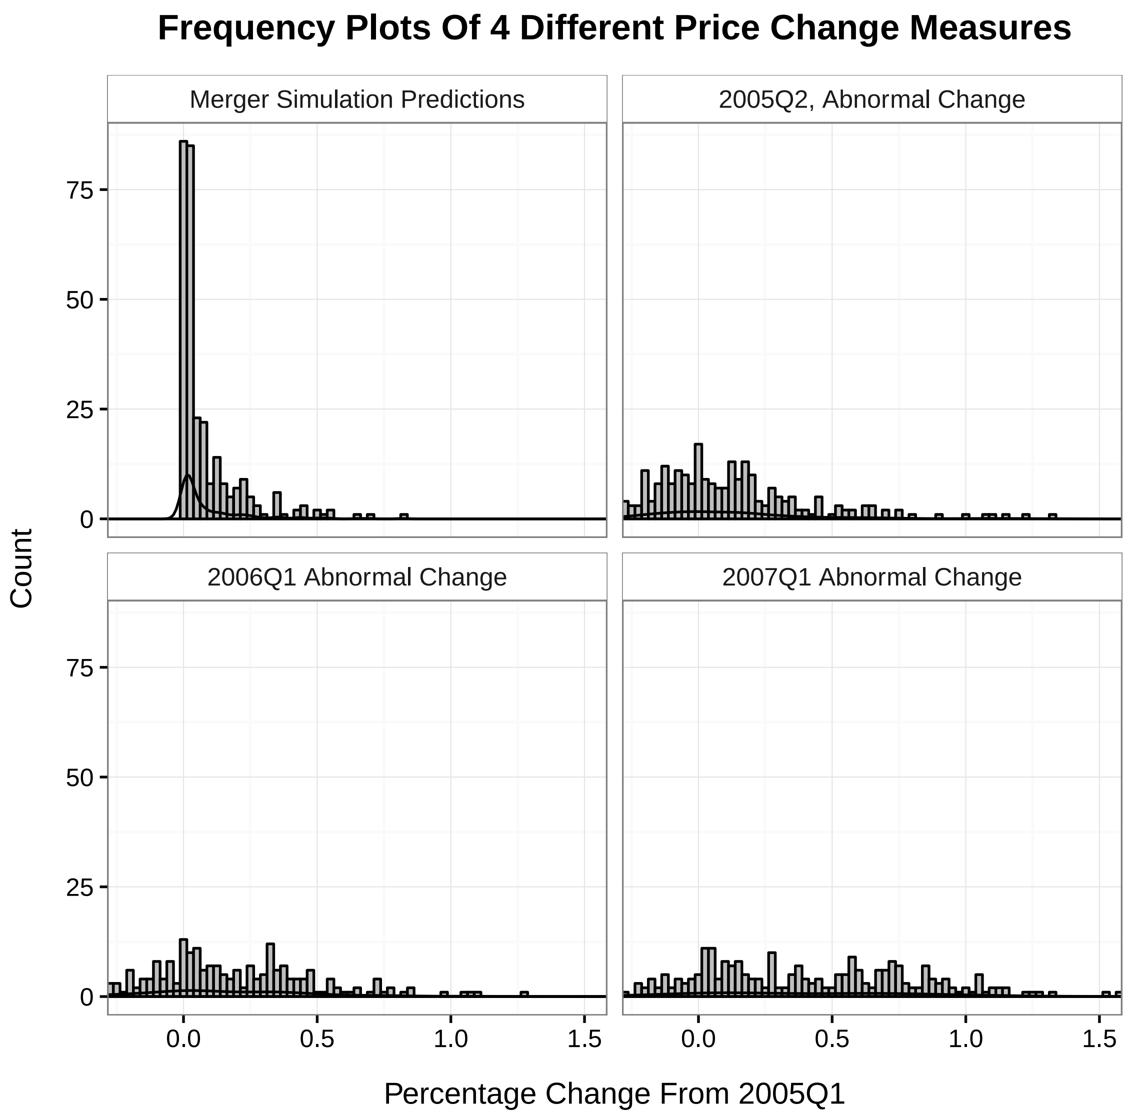

Figure 11: One can clearly see that the spread of the observed abnormal price changes is much greater than the spread of the logit merger simulation. Also of note is that, due to the assumption of zero cost efficiency gains, the merger simulation will by design not predict negative price changes.

	2005, Q1	2005, Q2	2006, Q1	2007, Q2
Average Fare, US Airways ($)	350.80	392.60	441.80	495.90
	(90.49	(142.56)	(139.93)	(132.05)
Average Market Fare ($)	413.30	419.10	438.70	448.60
	(138.04)	(49.77	(52.56)	(72.37)
Average Price Change (%)	-	15.96	28.06	55.03
		(44.13)	(50.48	(59.48)
Average Market Price Change (%)	-	1.99	6.44	8.85
		(8.92)	(12.59)	(13.03)
Average Abnormal Price Change (%)	-	13.96	21.62	46.17
		(41.90)	(46.58)	(53.61)
Average Difference (Abnormal - Market)	-	11.96	15.18	37.32
		(41.51)	(45.91)	(50.49)
Average Difference (Merger Sim - Abnormal)	-	4.80	13.80	36.77
		(43.27)	(46.72)	(50.69)
Average Difference (UPP - Abnormal)	-	4.54	13.52	36.52
		(43.38)	(46.81)	(50.78)
Correlation, Merger Sim and Abnormal	-	.0433	.1055	.3365
Correlation, Merger Sim and UPP	-	.0380	.1006	.3291
Market Count	296	248	227	260

Table 11: Aggregate summary statistics, with standard deviations in parenthesis, of the observed price changes in the original set of overlapping markets. Some of the initial markets were not offered in subsequent quarters, so the overall set and number of overlapping markets changes from quarter to quarter, with the total number indicated in the last row.

structure into consideration and simply assumes (or ignore) any other changes. The predicted changes are therefore theoretically supposed to take place immediately and with no frictions or complications. Our ex-post comparison is a second-best alternative measure to use as an empirical example. First, there are other factors that change from one moment in time to the next. These can be factors unrelated to the merger, such as seasonal volatility in demand or changes in input costs, such as fuel cost. This is why we make a simple attempt of isolating the effect from these kinds of factors by calculating the abnormal price increase. By assuming that most of these outside factors affect each airline in the same market the same way, we simply assume that whatever price change is left when this is taken away is closer to the effect from the merger. We recognize that this might not do the full job in isolating the effect from the merger.

Furthermore, there can be other effects relating to the merger that the model does not take into account. The most prominent example we have in our model is that the firms act as market-independent profit maximizers. In reality, US-wide operating firms maximize their profits over all markets they operate in; some may even be operating globally. Thus, it is not necessarily the case that the pricing incentives for a firm that treats every market as completely independent is the same as for a firm that recognizes that its behavior in one market can strategically influence the outcome in other geographically distinct markets. The standard merger simulation models ignore this fact, and it is an assumption that most academic models of airline behavior are built on. The complexity of modelling such hub effects proves to be prohibitive in merger evaluation techniques at the moment, and the research on such propositions is sparse, however it must be noted that such behavior could be contradictive to one of our fundamental assumptions.

We have to also consider the fact that our data source, although widely recognized and used in research within the airline industry and also within the field of antitrust and merger simulations, can have flaws that makes it impossible for us to fully capture the true pricing pattern. Prices in the airline industry depend on many factors which are in retrospect unobservable to the economist, such as when and how the ticket is booked. Although our random-coefficients demand estimation procedure attempts to account for as many of the characteristics not documented in our dataset as possible, we simply cannot be sure that our assumptions about consumer purchasing behavior are correct. On top of this uncertainty is the fact that our dataset is only a ten percent sample of all tickets purchased in the US; we can say that because of this there is a positive probability of statistical error.

6 Conclusion

This thesis has investigated the difference between the UPP and an alternative simple merger simulation as tools to predict unilateral effect from mergers, as well one possible procedure of how to obtain the necessary inputs for these measures. We use publicly available data from the airline industry and estimate margins using a structural demand estimation. By making assumptions about the pass-through rates we can use the result from the UPP as an approximate price change prediction. This makes them directly comparable to the predicted price changes from a logit merger simulation. Furthermore, we have used post merger data to explore the patterns in actual price changes across markets in the US, and have compared these observations with the different predictions.

Our results indicate that the UPP provides very similar results to a merger simulation if the input information is similar. Therefore it is reasonable to think of the UPP as a rough approximation of predicted price changes (at least to the same extent as a merger simulation is) and our result directly responds to part of the criticism that the UPP is not directly related to price changes. Thus, in a theoretical framework, where prices, margins, and market shares are known with certainty, the decision to use a logit merger simulation or the UPP as a price change indicator should be one of little impact. In fact the simplicity, and therefore advantage, of using the UPP measure comes from the fact that estimations of all products elasticities do not have to be derived; rather, the only substitution statistic which must be formulated is the diversion ratio between the two merging products/firms. In a setting where complete market information is easy to obtain, or the structural assumptions governing the market are not complex, the upwards pricing pressure index may not be of much interest to the practitioner, as it may be the case that a merger simulation is just as easy to implement. However where we believe the UPP has the potential to add the most value is in the analysis of complex differentiated products market spaces. In industries where the structural assumptions are complex, or the precise theoretical definitions of the input variables are difficult to formulate (as we have demonstrated is the case with the US airline industry), the estimation of one simple diversion ratio between two products versus a full elasticity matrix can save practitioners a significant amount of time. Another important implication of our result is that the UPP does indeed serve a purpose as a predictor of the magnitude of price changes. The complexity and controversy lies in how its inputs are collected or estimated; for the airline industry this proves to not be as trivial as one would initially think.

Unlike some other similar investigations we cannot draw bold conclusions about the accuracy of the predictions of UPP by only comparing it to predictions from a merger simulations. Further, after having compared the predictions to the actual outcomes in a multitude of markets, our research does not provide any indications that the predicted price changes were accurate overall, with the exception that the predicted sign of the change is generally accurate. Even this result though, is based on multiple assumptions - such as full pass-through rates and no supply efficiency gains - which may or may not manifest in reality. Because our model, like most models of unilateral effects of mergers, treats markets as individually separate mergers it is intuitive to compare predictions to the actual outcome market by market. However, our results highlights not only a gap in the literature of thorough empirical testing of merger tools and the predictive capabilities but also the acknowledgement of multimarket maximizing firms. Specifically, when working with the airline industry, we find the "traditional" assumption that firms treat different geographic markets as independent of each other to be dubious at best. Future research may benefit from investigating how airlines actually profit maximize and whether or not they consider the interaction between geographically distinct markets into account.

References

Baker, J. B. and Reitman, D. (2009), 'Research topics in unilateral effects analysis', *American University, WCL Research Paper* (09-37).

Benkard, C. L., Bodoh-Creed, A. and Lazarev, J. (2010), 'Simulating the dynamic effects of horizontal mergers: Us airlines', *Manuscript, Yale University* .

Berry, S., Carnall, M. and Spiller, P. T. (1996), Airline hubs: costs, markups and the implications of customer heterogeneity, Technical report, National bureau of economic research.

Berry, S. and Jia, P. (2010), 'Tracing the woes: An empirical analysis of the airline industry', *American Economic Journal: Microeconomics* **2**(3), 1–43.

Berry, S., Levinsohn, J. and Pakes, A. (1995), 'Automobile prices in market equilibrium', *Econometrica: Journal of the Econometric Society* pp. 841–890.

Borenstein, S. (1990), 'Airline mergers, airport dominance, and market power', *The American Economic Review* **80**(2), 400–404.

Brueckner, J. K. and Pels, E. (2005), 'European airline mergers, alliance consolidation, and consumer welfare', *Journal of Air Transport Management* **11**(1), 27–41.

Carlsson, F. and Löfgren, Å. (2006), 'Airline choice, switching costs and frequent flyer programmes', *Applied Economics* **38**(13), 1469–1475.

Cheung, L. (2013), 'The upward pricing pressure test for merger analysis: An empirical examination'.

Dafny, L., Duggan, M. and Ramanarayanan, S. (2012), 'Paying a premium on your premium? consolidation in the us health insurance industry', *The American Economic Review* **102**(2), 1161–1185.

Dubé, J.-P. H., Fox, J. T. and Su, C.-L. (2009), Improving the numerical performance of blp static and dynamic discrete choice random coefficients demand estimation, Technical report, National Bureau of Economic Research.

Epstein, R. J. and Rubinfeld, D. L. (2004), *Technical Report: Effects of Mergers Involving Differentiated Products*, European Commission.

Farrell, J. and Shapiro, C. (1990), 'Horizontal mergers: an equilibrium analysis', *The American Economic Review* pp. 107–126.

Farrell, J. and Shapiro, C. (2010), 'Antitrust evaluation of horizontal mergers: An economic alternative to market definition', *The BE Journal of Theoretical Economics* **10**(1).

Fischer, T. and Kamerschen, D. R. (2003), 'Price-cost margins in the us airline industry using a conjectural variation approach', *Journal of Transport Economics and Policy (JTEP)* **37**(2), 227–259.

Froeb, L., Tschantz, S. and Werden, G. J. (2005), 'Pass-through rates and the price effects of mergers', *International Journal of Industrial Organization* **23**(9), 703–715.

Huang, D., Rojas, C. and Bass, F. (2008), 'What happens when demand is estimated with a misspecified model?', *The Journal of Industrial Economics* **56**(4), 809–839.

Ivaldi, M., Jullien, B., Rey, P., Seabright, P. and Tirole, J. (2007), 'The economics of tacit collusion: implications for merger control', *The Political Economy of Antitrust (Contributions to Economic Analysis* **282**, 217–239.

Ivaldi, M., Jullien, B., Rey, P., Seabright, P., Tirole, J. et al. (2003), 'The economics of unilateral effects', *Interim report for DG competition, European Commission* .

Jaffe, S. and Weyl, E. G. (2013), 'The first-order approach to merger analysis', *American Economic Journal: Microeconomics* **5**(4), 188–218.

Joskow, A. S., Werden, G. J. and Johnson, R. L. (1994), 'Entry, exit, and performance in airline markets', *International Journal of Industrial Organization* **12**(4), 457–471.

Kim, E. H. and Singal, V. (1993), 'Mergers and market power: Evidence from the airline industry', *The American Economic Review* pp. 549–569.

Klemperer, P. (1995), 'Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade', *The review of economic studies* **62**(4), 515–539.

MATLAB (2016), *version 9.0.0 (R2010a)*, The MathWorks Inc., Natick, Massachusetts.

McDonald, J. B. (2005), Antitrust for airlines, *in* 'Statement to the Regional Airline Association Presidents Council Meeting. Retrieved October', Vol. 19, p. 2007.

Miller, N. H., Remer, M., Ryan, C. and Sheu, G. (2015), 'Pass-through and the prediction of merger price effects'.

Miller, N., Remer, M., Ryan, C., Sheu, G. et al. (2016), Upward pricing pressure as a predictor of merger price effects, Technical report, Department of Justice, Antitrust Division.

Moresi, S. (2009), 'Cournot competition and the upp test', *HMG Review Project-Comment. Project* (P092900).

Mueller, D. C. (1996), 'Lessons from the united states's antitrust history', *International Journal of Industrial Organization* **14**(4), 415–445.

Neven, D. J., RÖLLER, L.-H. and Zhang, Z. (2006), 'Endogenous costs and price-cost margins: An application to the european airline industry', *The Journal of Industrial Economics* **54**(3), 351–368.

Nevo, A. (2001), 'Measuring market power in the ready-to-eat cereal industry', *Econometrica* **69**(2), 307–342.

O'Brien, D. P. and Salop, S. C. (2000), 'Competitive effects of partial ownership: Financial interest and corporate control', *Antitrust Law Journal* **67**(3), 559–614.

Peters, C. T. (2003), 'Evaluating the performance of merger simulation: Evidence from the us airline industry', *DOJ Antitrust Division Economic Analysis Group Discussion Paper* (03-1).

R Core Team (2015), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: *http://www.R-project.org/*

Simons, J. J. and Coate, M. B. (2010), 'Upward pressure on price analysis: Issues and implications for merger policy', *European Competition Journal* **6**(2), 377–396.

Taragin, C. and Sandfort, M. (2015), *antitrust: Tools for Antitrust Practitioners*. R package version 0.95.1.
URL: *http://CRAN.R-project.org/package=antitrust*

United States Department of Justice (2010), 'Horizontal merger guidelines'.

United States v. Philadelphia Nat. Bank (1963).

Weinberg, M. (2008), 'The price effects of horizontal mergers', *Journal of competition Law and Economics* **4**(2), 433–447.

Weinberg, M. et al. (2007), *Price Effects of Horizontal Mergers: A Survey*, Center for economic Policy Studies, Princeton University.

Werden, G. J. (1996), 'A robust test for consumer welfare enhancing mergers among sellers of differentiated products', *The Journal of Industrial Economics* **44**(4), 409–413.
URL: *http://www.jstor.org/stable/2950522*

Werden, G. J. (2009), 'Economic reasoning in merger cases and how courts should evaluate it', *European Competition Journal* **5**(3), 701–720.

Werden, G. J. and Froeb, L. (2002), 'Calibrated economic models add focus, accuracy, and persuasiveness to merger analysis', *THE PROS AND CONS OF MERGER CONTROL : 10th ANNIVERSARY OF THE SWEDISH COMPETITION AUTHORITY, Swedish Competition Authority, ed., Swedish Competition Authority, Stockholm* .

Werden, G. J. and Froeb, L. M. (1994), 'The effects of mergers in differentiated products industries: Logit demand and merger policy', *Journal of Law, Economics, & Organization* pp. 407–426.

Werden, G. J. and Froeb, L. M. (2011), 'Choosing among tools for assessing unilateral merger effects', *European Competition Journal* **7**(2), 155–178.

Werden, G. J., Joskow, A. S. and Johnson, R. L. (1991), 'The effects of mergers on price and output: Two case studies from the airline industry', *Managerial and Decision Economics* **12**(5), 341–352.

Wickham, H. (2009), *ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag New York.
URL: *http://ggplot2.org*

Wickham, H. and Francois, R. (2015), *dplyr: A Grammar of Data Manipulation.* R package version 0.4.3.
URL: *http://CRAN.R-project.org/package=dplyr*

Wickham and Hadley (2007), 'Reshaping data with the reshape package', *Journal of Statistical Software* **21**(12).
URL: *http://www.jstatsoft.org/v21/i12/paper*

Appendices

Appendix I

Most of our analysis was done in the R programming language (R Core Team, 2015). Our data cleaning and data construction was done with the help of the R package "Dplyr" (2015), while the graphs were created with the use of the "Ggplot2" and "Reshape" packages ((Wickham, 2009) & (Wickham and Hadley, 2007)). The final package we used in the R environment was arguably the most crucial; the "Antitrust" package puts powerful merger analysis tools at the fingertips of any interested student or practitioner (Taragin and Sandfort, 2015).

Our random-coefficients logit demand estimation was directly adapted from Steven Berry and Panle Jia's estimation procedure that was published with their article "Tracing the Woes: An Empirical Analysis of the Airline Industry" (2010). Their data and their estimation procedure is publically available on the American Economic Associations website[13]. This procedure was run exclusively in Matlab (2016).

[13] `https://www.aeaweb.org/articles?id=10.1257/mic.2.3.1`

Appendix II

Metropolitan statistical areas which fell under our analysis, along with their 2005 population and the relevant airports within their area.

Metropolitan Statisical Area	Population	IATA Airport Code(s)
Albuquerque NM	797517	ABQ
Albany-Schenectady-Troy NY	847421	ALB
Anchorage AK	351586	ANC
Atlanta-Sandy Springs-Marietta GA	4972219	ATL
Austin-Round Rock TX	1454706	AUS
Hartford-West Hartford-East Hartford CT	1185700	BDL
Birmingham-Hoover AL	1088218	BHM
Nashville-Davidson–Murfreesboro TN	1421124	BNA
Boise City-Nampa ID	545141	BOI
Boston-Cambridge-Quincy MA-NH	4448884	BOS, MHT, PVD
Buffalo-Niagara Falls NY	1144796	BUF
Los Angeles-Long Beach-Santa Ana CA	12933839	BUR, LAX, LGB, ONT, SNA
Washington-Arlington-Alexandria DC-VA-MD-WV	5251629	BWI, DCA, IAD
Cleveland-Elyria-Mentor OH	2125138	CLE
Charlotte-Gastonia-Concord NC-SC	1521474	CLT
Columbus OH	1706913	CMH
Cincinnati-Middletown OH-KY-IN	2090968	CVG
Dallas-Fort Worth-Arlington TX	5823043	DAL, DFW
Denver-Aurora CO1	2361778	DEN
Detroit-Warren-Livonia MI	4479254	DTW
El Paso TX	721183	ELP
New York-Northern New Jersey-Long Island NY-NJ-PA	18813723	EWR, JFK, LGA
Spokane WA	440434	GEG
Honolulu HI	904645	HNL
Houston-Sugar Land-Baytown TX	5352569	HOU, IAH
Indianapolis-Carmel IN	1640029	IND
Jacksonville FL	1247828	JAX
Las Vegas-Paradise NV	1709364	LAS
Salt Lake City UT	1046685	LC
Kansas City MO-KS	1944690	MCI
Orlando-Kissimmee FL	1931479	MCO
Chicago-Naperville-Joliet IL-IN-WI	9446565	MDW, ORD
Memphis TN-MS-AR	1256631	MEM
Miami-Fort Lauderdale-Miami Beach FL	5424697	MIA, FLL
Milwaukee-Waukesha-West Allis WI	1509388	MKE
Minneapolis-St. Paul-Bloomington MN-WI	3141050	MSP
New Orleans-Metairie-Kenner LA	1313787	MSY
Kahului-Wailuku HI	139687	OGG
Oklahoma City OK	1154991	OKC
Omaha-Council Bluffs NE-IA	812830	OMA
Virginia Beach-Norfolk-Newport News VA-NC	1641543	ORF
West Palm Beach-Boca Raton-Boynton Beach FL	1264956	PBI
Portland-Vancouver-Beaverton OR-WA	2096571	PDX
Philadelphia-Camden-Wilmington PA-NJ-DE-MD	5806092	PHL
Phoenix-Mesa-Scottsdale AZ	3878525	PHX
Pittsburgh PA	2381671	PIT
Raleigh-Cary NC	951809	RDU
Reno-Sparks NV	393820	RNO
Cape Coral-Fort Myers FL	544196	RSW
San Diego-Carlsbad-San Marcos CA	2936609	SAN
San Antonio TX	1888047	SAT
Louisville-Jefferson County KY-IN	1210182	SDF
Seattle-Tacoma-Bellevue WA	3207892	SEA
San Francisco-Oakland-Fremont CA	4158012	SFO, OAK, SJC
Sacramento–Arden-Arcade–Roseville CA	2041701	SMF
St. Louis MO-IL	2782411	STL
Tampa-St. Petersburg-Clearwater FL	2646540	TPA
Tulsa OK	885778	TUL
Tucson AZ	925000	TUS
San Juan PR	2579997	SJU