Quantifying the effects of natural hedging – An examination of US production for BMW and Porsche [♠]

Richard Friberg[♣], Stockholm School of Economics and CEPR

Cristian Huse[♦], Stockholm School of Economics

This version: July 2014

Abstract

We quantify the effects of "natural hedging", producing cars sold in the US locally, for the risk profile of the US operations of German carmakers BMW and Porsche. There are three steps in the simulation procedure we use. First, we estimate a random coefficients logit demand system for differentiated products using data from the US car market. Second, we generate counterfactual paths to macroeconomic risk factors using copulas, in a way that flexibly can be adapted to the risks faced in various industries. We then feed the counterfactual draws into the demand system, letting prices and quantities adjust, to generate profit distributions under different assumptions on production locations. Natural hedging reduces exchange rate exposure, decreasing profit variability substantially.

JEL Classification Codes: F23, L16, L62

Keywords: Exchange rate exposure, macroeconomic exposure, operational hedging, natural hedging, risk management.

[♠] We are grateful to the Swedish Research Council (VR) and Jan Wallanders and Tom Hedelius Stiftelse for financial support. We thank Elisa Alonso, Marcus Asplund, Johannes van Biesebroeck, Carlos Noton, Rickard Sandberg and seminar audiences at CEMFI, the CEPR/JIE workshop in Applied IO in Tel Aviv, EARIE in Istanbul, ESSEC, Foro de Finanzas in Elche, HECER, Lund, Stockholm School of Economics, Uppsala and Queen Mary for valuable comments.

[♣] Email: nerf@hhs.se. Correspondence address: Stockholm School of Economics, Dept of Economics, Box 6501, SE-113 83 Stockholm, Sweden.

[♦] Email: cristian.huse@sse.edu. Correspondence address: Stockholm School of Economics, Dept of Finance, Box 6501, SE-113 83 Stockholm, Sweden.

Introduction

The German car maker BMW produces a number of models in the US and states in the annual report for 2007 (p. 62) that "From a strategic point of view, i.e. in the medium and long term, the BMW Group endeavours to manage foreign exchange risks by 'natural hedging', in other words by increasing the volume of purchases denominated in foreign currency or increasing the volume of local production." Similarly, Volkswagen recently built a plant in Tennessee and states in its annual report 2009 (p 188) that "Foreign currency risk is reduced primarily through natural hedging, i.e. by flexibly adapting our production capacity at our locations around the world, establishing new production facilities in the most important currency regions and also procuring a large percentage of components locally". Several Asian carmakers also have significant production capacity in North America, and natural hedging is one stated reason for this.[1] Other carmakers follow different strategies. Porsche for instance produces exclusively in the euro area but has 30-40 percent of its sales in North America. How would the risk profile of Porsche change if it were to produce in the US?

In this paper we generate counterfactual profit distributions for the US operations of BMW and Porsche to examine the consequences on the risk profile of producing some models locally in the US. We use product level data for the top segments of the US auto market for 1995-2006 to estimate demand that serves as the main input in our counterfactuals. We follow Berry, Levinsohn and Pakes (1995) and model demand using a random coefficients logit model. We generate forward looking counterfactual values on exchange rates and on a measure of the business cycle (consumer confidence) based on data from 1973-2006. We use copulas to model the correlation of the draws between exchange rates and consumer confidence. While novel to Industrial Organization, copulas have seen rapid adoption in other fields such as asset pricing (see Patton (2009) for an overview). To generate profit distributions we use simulation methods and feed the counterfactual values of exchange rates and consumer confidence into the demand system, letting prices and quantities respond. Our results illustrate the rationale underlying natural hedging; increasing the volume of production in the consumer market reduces exchange rate exposure, which in turn results in less-dispersed profit distributions. In particular, firms become less exposed to losses due to movements in the exchange rate, which suggests that natural hedging is an attractive strategy for managers that place large weights to negative outcomes.

To introduce the issues, and highlight the challenges of gauging the benefits of natural hedging, let us contrast two highly stylized investment possibilities. In the first case a German firm produces in Germany and exports all sales to the US. Letting e denote the euro-dollar exchange rate, p

[1] In Toyota's annual report (2007, p 77) it is for instance written that "Localizing production enables Toyota to locally purchase many of the supplies and resources used in the production process, which allows for a better match of local currency revenues with local currency expenses."

denote the price in US dollars, c the constant marginal cost in euros, q sales in the US and F the fixed cost of production, the profit is thus equal to

$$\Pi = (ep - c)q - F. \tag{1}$$

Other things equal, a depreciation of the euro, a higher value of e, makes for higher profits from US sales, when expressed in euros. Conversely, an appreciation of the euro will lower profits.

If the firm instead engaged in natural hedging, and produced all its US sales locally in the US, the profit, when translated into euros, would instead be given by

$$\Pi_{\mathrm{u}} = e(p_u - c_u)q_u - F_u, \tag{2}$$

where the subscript u highlights that if production is located in the US prices, quantities and marginal costs may all differ from what would be optimal if production instead were in Germany. The key difference between equations (1) and (2) regards how the exchange rate enters the profit equation. When production is in Germany, as in (1), *an appreciation of the euro (lower e) is associated with lower revenue in euros but marginal costs are unchanged*. In contrast, under natural hedging, as in equation (2), marginal costs are also falling from the perspective of a German producer when the euro appreciates. For now keeping all variables constant, equation (1) leads to an exchange rate exposure of $\delta\Pi/\delta e = pq$, a change in the exchange rate is proportional to revenue in dollars.[2] In the case of (2) $\delta\Pi/\delta e = (p_u - c_u)q_u$, a change in the exchange rates is proportional to net revenue in dollars. The purpose of this paper is to quantify how the distribution of net present values for BMW and Porsche depend on whether US sales are produced locally in the US or not.

If the profit streams associated with the two different investments (produce at home or abroad) were certain it would be a simple matter to calculate present value of profits and then choose the location with the highest net present value (see Brennan (2003) for an overview of the literature on investment rules). In contrast, when there is risk, we need to create counterfactual profit distributions.[3] We want to account for that the exchange rate can take many different values in future periods and demand may be subject to business cycle shocks, with possible correlation to both the euro exchange rate and to cost shifters for competitors such as the exchange rate *vis-à-vis* the yen. How should such counterfactuals be generated? Apart from reasoned "guesstimates", textbooks in finance and international business suggest that one selects a probability distribution for each of a set of variables

[2] Clearly this simple example is only for intuition (even if Marston (2001) stresses that in some situations the envelope theorem implies that the effect of exchange rates on profits is this simple). In our analysis we will take account of that prices change as well as subjecting demand to other shocks.

[3] Note that this is true even if the decision maker is risk neutral as expected profit will be affected by the nature of shocks unless we are in the very special case where profits are linear in all shocks. If firms are risk averse or want to avoid low realizations of profits to finance ongoing investments (as in Froot, Stein and Scharfstein (1993)) the reason for evaluating the whole distribution is further strengthened.

that affect profits, such as price and market size, and then use these distributions to generate counterfactuals.[4] Hertz (1964) is an early proponent of this method. Despite its use in business and teaching, and the marketing of a large number of software applications, the academic literature on the method is slight.[5] The *ad hoc* nature of assumptions regarding the risk distributions of prices and quantities, and their relation, are the probable reason for the limited attention of academics. [6]

We propose to use what have now become standard tools in empirical Industrial Organization, coupled with counterfactual draws on macroeconomic risk factors, to aid simulations of project value. The idea to feed a large number of counterfactual cost and demand shocks into a system of demand for differentiated products to generate counterfactual profit distributions seems trivial. At the same time it allows us to ensure economically sound relations between variables that affect profits and thus address a weakness of the Hertz method. Despite this, it is an avenue that has hardly been pursued in the previous literature. Previous applications of demand models similar to the one we estimate typically consider only one, or a few counterfactual scenarios. Prominent examples include evaluations of mergers (Nevo (2000a)), measurements of the impact of trade policy (Berry, Levinsohn and Pakes (1999)) or quantification of the welfare effects of entry (Petrin (2002)). Somewhat closer in spirit to the present work is Berry and Jia (2010), who provide an *ex post* analysis of the sources of profit changes in the US airline industry between 1999 and 2006. They for instance find that just a few observed changes, in particular a greater price sensitivity on the part of consumers and a stronger preference for direct flights, can explain around 80 percent of the fall in profitability for the legacy carriers. The perhaps closest precursor is Friberg and Ganslandt (2007) who examine exchange rate exposure on the Swedish market for bottled water and generate counterfactual profits following the same logic as in the current paper. The present paper extends that work in several ways. They use a nested logit specification for demand whereas we model demand in a much less restrictive fashion. They use shocks that are bivariate normal and consider only one counterfactual period, whereas we generate counterfactual paths of shocks that easily extend to other settings. Most importantly, we use the methodology to examine different operating strategies. Finally, one can argue that natural hedging on automobile markets is a more interesting application of risk measurement than the Swedish market for bottled water.

Our work also bears a close relation to dynamic oligopoly games (see for instance Ericson and Pakes (1995), Bajari, Benkard, Levin (2007), or Ackerberg et al (2006) and Aguirregabaria and Nevo

[4] Alternatively these sources suggest that one can use a decision tree to analyze future values of the firm or consider a limited set of alternative scenarios. We are not offered any guidance on how to generate quantitative estimates for the different scenarios or branches however, which is the aim of the present project.

[5] Most software applications are based on the spreadsheet program excel – see for instance the commercial products @RISK or Crystal Ball.

[6] McAfee (2002, p 257) for instance notes that "It is almost invariably a mistake in this approach to assume the variables are independently distributed. In particular, macroeconomic variables like income, interest rates, growth rates, and so on have a known covariance structure. Accounting for such covariances is a major challenge for scenario analysis generally, but a larger challenge the more scenarios there are".

(2012) for surveys). These papers develop tools to estimate structural models of demand and use them to examine industries over time, while allowing for strategic choices to affect the payoffs of competitors. However, when considering many strategies of several competitors the state space grows rapidly and computational costs are an important restriction. At the risk of oversimplifying, the papers in this literature have concentrated on inferring parameters or behavior that is hard to observe directly, such as the sunk costs of entry. Such information is of clear importance to a policymaker trying to, for instance, gauge the probability of entry following some policy change (for steps in the latter direction see Benkard, Bodoh-Creed and Lazarev (2010)). The assumptions on the type of shocks faced by firms are typically quite stylized (such as i.i.d. firm specific shocks to the sell-off value of the firm) and neither the time series properties of shocks, nor using the models in a forward looking manner, have been the focus of this literature. In contrast, the present paper puts the *future distribution of shocks* center stage – that is, it focuses on how should you value the profits associated with an investment when exogenous risks such as exchange rates or the business cycle are important. For many applications we ultimately wish to have a framework that is suited for both dealing with uncertainty that arises because of the strategic interaction and for dealing with the risks that stems from the stochastic nature of exogenous demand and cost shocks – see Besanko et al (2010) for such a combination in a stylized framework.[7] For the time being we believe that is useful to complement work that focuses on the strategic interaction with work that focuses on how exogenous shocks feed through into profits – and how the impact of cost and demand shocks depends on strategic choices.

In the next section we present the data and describe the product ranges of BMW and Porsche in some detail. We also highlight some of the difficulties of relying on standard forecasting techniques in a differentiated products oligopoly. In Section 3 we present our estimation methods and specify how counterfactuals are generated. Section 4 we show the results from demand estimation and from the generation of counterfactual macroeconomic conditions. The counterfactual profits are then presented and analyzed in Section 5. We conclude in Section 6.

2 The Data and the Firms

We examine consequences of production location for BMW and Porsche. We have chosen to limit the analysis to the US operations of BMW and Porsche rather than examining the global risk profile of firms. For our demand estimation and counterfactuals we need not only data on BMW and Porsche but also on competing products. Thus, we use quantity sold, recommended dealer price and product characteristics for all cars sold in the luxury, sport, SUV (sports utility vehicles) and CUV (cross over utility vehicles) segments in the US. The main source of data is WARDS who supplied us with a

[7] Clearly, it is easy to consider scenarios in the framework that we use – we are referring here to the broader evolution of industry based on the extent of sunk costs and other industry characteristics.

panel of monthly sales by model line (BMW 3 series, Porsche 911 etc). We examine the period from August 1995 to July 2006. In our regression analysis we aggregate sales to 12-month periods. Rather than use calendar years we note that new models, and a new recommended dealer price, appear in late summer each year. Our time unit of analysis therefore runs from August to July the following year and we use the term model-year.

In Table 1 below we show some descriptive statistics for our set of cars. We examine the upper segments of the car market and the mean real price is roughly stable at 35,000 dollars. The lowest price is for a Pontiac G5 and the highest is for a Porsche Carrera GT. On average some 30,000 to 40,000 cars are sold per model in a given model-year. The largest selling name plate in the data is the Ford Explorer. The number of models in the data increases substantially over the period, mainly reflecting growth in the CUV and SUV segments.

[Table 1 about here]

We focus on three macroeconomic variables in the analysis – the real exchange rates between the dollar and the euro (usd/eur), between the dollar and the Japanese Yen (usd/jpy) and the measure of consumer confidence published by the Conference Board. Consumer confidence is frequently mentioned in the industry as an important covariate of demand for cars. This is confirmed by Ludvigson (2004) who also examines the relation between different measures of consumer confidence. The dollar appreciated against the euro and yen up until the middle of the period, after that it depreciated against the euro but remained rather stable against the yen. The consumer confidence measure of the business cycle shows substantial variability as well.

Finally, we collect production location of each model in our dataset for the period 1995-2006 from company webpages and specialized publications.

2.1 The US market for BMW and Porsche, a closer look

BMW

German-based BMW is one of the ten largest car manufacturers in the world. Compared to other auto manufacturers, the accounting figures point to BMW as a profitable firm with high margins: its return on assets is on average 5.3 percent and the profit margin is 15.6 percent (EBITDA operating margin before interest, taxes, depreciation and amortization).

The main products for BMW over this period are the luxury cars in the 3, 5 and 7 series. At the start of the period it also sells the roadster Z3. Although for the purposes of our analysis we will be focusing on the BMW brand, we should also mention until 2000 the BMW Group also controlled the

Land Rover and Range Rover lines as well as the Mini, all of which were produced in the UK. Production location for the BMW brand varied somewhat across the years. The first model produced in the US plant in Spartanburg, SC, starting from mid-1995 (model-year 1996) was the roadster Z3. In 1999 it was followed by the BMW X5, a middle luxury CUV, and the Z3's successor, the BMW Z4, in 2003. However, starting from 2008, when the second generation of the BMW Z4 was introduced, its production was moved to BMW's Regensburg plant (Germany). Thus, at the end of the sample period only BMW models X5 and Z4 were produced in the US, with all other products of the BMW brand being produced in the euro area. Over the period, on average, 23.7 percent of BMW deliveries of cars are in North America. We therefore expect a potentially important role for the usd/euro exchange rate on BMW profits. Indeed the annual report for 2005 (p. 56) notes that "Of all the currencies in which the BMW group does business, the US dollar represents the main single source of risk; fluctuations in the value of the US dollar have a major impact on reported revenues and earnings."

[Table 2 about here]

Porsche

During the time period that we examine however, accounting profitability and operating margins are high at Porsche: the return on assets is on average 19.7 percent and the operating margin is 24.7 percent. Porsche's main product over the period is the 911 - a name plate that was introduced in 1963 and still accounts for almost half of US revenue at the end. Initially, the 911 is the only model marketed by Porsche in the US. The small roadster Boxster is then introduced in late 1996. The Cayenne is introduced in 2003 (identified as a middle luxury CUV by WARDS) and the sports car Cayman in 2005. In 2004 Porsche adds the top-of-the-line sports car Carrera GT. After only having had assembly in Germany, Porsche starts production of its Boxster in Finland in 1997 (under an agreement with Finnish producer Valmet). Since 2005 also the Cayman model is produced in Finland which, like Germany, is part of the euro zone. The North American market accounted for an average of 35 percent of sales revenue for Porsche. With a substantial share of revenue from the North American market, but all costs in Europe, we expect that Porsche profits are exposed to the US dollar. Indeed, prior to our period of study Porsche's profits had a strong relation to the dollar. In the mid 1980s, at the peak of the strong dollar, more than 60 percent of Porsche's sales were to North America. Over the latter part of the 1980s, and early 1990s, the dollar weakened against the German mark and by the early 1990s Porsche was having grave financial difficulties. [8]

[Table 3 about here]

[8] Indeed, Porsche is enough of a schoolbook case on exchange rate exposure that it is featured as mini cases in two of the leading textbooks in international finance (Eiteman, Stonehill and Moffett (2007, p 322) and Eun and Resnick (2007, p 236). In the present paper we want to move beyond qualitative discussions in these works and examine the quantitative implications of different strategies.

2.2 Why use a structural model?

Key to our comparison of different investment scenarios is the future evolution of profit flows. A natural starting point would perhaps be to consider historical profit flows and use regressions based on historical profits to generate forecasts. One could for instance regress profits on consumer confidence, exchange rates and, using Monte Carlo methods to take draws on these variables, generate forward looking profit distributions.[9] We believe that important limitations in the application of such a methodology to a differentiated products such as automobiles. To highlight why, let us consider Porsche's revenue flows from US sales (in euros) in Figure 1.

[Figure 1 about here]

Eyeballing the figure it is easy to envision that there is a link between Porsche's revenue and the business cycle as measured by consumer confidence, especially taking into account the fact that the associated real prices are stale. One might also note that during 1996 to 2001 the euro weakened against the dollar and revenues from US sales, when converted into euros, show a trend-wise increase. Conversely, the strengthening of the euro in 2002 and 2003 is associated with lower revenue in euros but towards the end of the period the revenue seems robust to the stronger euro. A natural reason for the latter effect is that a new model, the Porsche Cayenne, was introduced in 2003 and proved successful. This exemplifies that *changes in the set of products sold* will affect profit flows, something that is illustrated in this case in Tables 2 and 3. One could use regressions at the product level, but for many of the products we would have very short time series data to estimate effects. We also need to deal with *endogenous price changes and changes in product characteristics*, both by the firm itself and by competitors.

The challenges in using product level data are therefore very similar to the challenges that one faces when evaluating prospective mergers. By using the hedonic approach to demand modeling we are able to use the implied consumer preferences to infer demand also for new products or products for which we only observe a short time series (see for instance Davis and Garcés (2010) for a discussion of the characteristics approach vs. the product level approach to demand modeling). Some observers are critical of structural modeling and argue for the empirical models that focus on identifying a causal effect (see for instance Angrist and Pischke (2010)). We agree that this is very attractive when the setting so allows, but just as in the case of mergers in differentiated products markets, we believe that idiosyncrasies and the ability to generating theoretically grounded counterfactuals favor structural

[9] A large number of articles examine the sensitivity of stock market prices to macroeconomic variables in this way (see for instance Dominguez and Tesar (2006)) and a smaller number of articles examine profit flows in this way (see for instance Oxelheim and Wihlborg (1995), see Andrén et al (2005) for an example where regressions on profit flows are combined with Monte Carlo techniques to gauge the sensitivity of profits to price risks.

models to perform counterfactuals (see for instance Nevo and Whinston (2010) or Einav and Levin (2010) for a discussion of mergers). Our reading of the evidence on merger simulation is that structural models of the type we use have indeed proved useful if one uses a demand specification that is sufficiently rich to generate truthful cross-price effects (see for instance Budzinsky and Ruhmer (2009), Weinberg (2011) or Björnerstedt and Verboven (2014)). The ability to perform counterfactuals and to let pricing adjust to different scenarios is an important motivation for us as well. Nevertheless, for the present paper, the main reason for relying on the characteristics approach to demand estimation is to provide good estimates of demand, despite short relevant time series data.

3 The Empirical Model and Generation of Counterfactuals

3.1 Estimating Demand and Backing out Marginal Costs

We follow BLP (1995) who estimate a random-coefficients (RC) logit model for automobiles in the US market. Define the conditional indirect utility of individual i when consuming product j in period t as:

$$u_{ijk} = \sum_{k=1}^{K} x_{jkt}\beta_{ik} + \xi_{jk} + \varepsilon_{ijt}, \qquad i = 1, \ldots I; j = 1, \ldots, J, t = 1996, \ldots, 2006$$

where x_{jkt} are observed product characteristics. As observable characteristics we use size (width x length), horsepower, a dummy for automatic transmission, price, as well as fixed effects for brand, country of production and for time. We also include a random coefficient on price, as explained below. We also interact consumer confidence with different dummy variables for different subsegments (16 in all) to capture that macroeconomic demand shocks can have differential impact on sales of different types of products. ξ_{jt} represent unobserved (by the econometrician) product characteristics, assumed observed by all market participants.

Following the literature, we decompose the individual coefficient on price according to

$$\beta_{ik} = \beta_k^* + \sigma_k \nu_{ki}$$

where β_k^* is common across individuals, ν_{ki} is an individual-specific random determinant of the taste for characteristic k, which we assume to be Normally distributed, and σ_k measures the impact of ν on characteristic k. Finally, ε_{ijt} is an individual and option-specific idiosyncratic component of preferences, assumed to be a mean zero Type I Extreme Value random variable independent from both the consumer attributes and the product characteristics. The specification of the demand system is completed with the introduction of an outside good with conditional indirect utility $u_{i0}=\xi_{0m}+\sigma_0+\nu_i+\varepsilon_{i0}$, since some consumers decide not to buy any car. Following standard practice in the literature we relate the potential market (M_t) with the number of existing households each year. [10]

It is assumed that consumers choose the product that yields the highest utility, and, integrating over consumers yields predicted market shares for each product *j* in period *t* (s_{jt}) as a function of parameters and product characteristics. We treat price as endogenous in our demand specification and use GMM to estimate parameters. To estimate our model, besides the exogenous characteristics, we use the BLP instruments (following BLP (1995)), a set of polynomial basis functions of exogenous variables exploiting the three-way panel structure of the data, consisting of the number of firms operating in the market, the number of other products of the same firm and the sum of characteristics of products produced by rival firms.

It is common to assume that competition in the US car market can be described as static Nash-Bertrand (see e.g. BLP (1995), Goldberg (1995), Petrin (2002)). We follow this assumption as well for the purpose of backing out marginal costs. Thus, we assume that multiproduct producer based in Germany sets prices of products *j* in year *t* so as to maximize the following profit function

$$\Pi = \sum_{j\epsilon F}\left(\frac{eur}{usd}_{jt}\, p_{jt} - mc_{jt}^{EU}\right) M_t\, s_{jt}, \tag{3}$$

where p is price in dollars, mc is a constant marginal cost expressed in euros, eur/usd_{jt} is the real exchange rate between the euro area and the US and M is the potential market. It may clarify to rewrite (1) in the following way to highlight that we may think of exchange rates as a marginal cost shock.

$$\Pi = \sum_{j\epsilon F}\left(p_{jt} - \frac{usd}{eur}_{jt}\, mc_{jt}^{EU}\right) M_t\, s_{jt}. \tag{4}$$

Using the first order conditions for prices from this maximization problem and rewriting in vector form implies that we can back out the marginal costs that are implied by the demand model in combination with multiproduct Nash-Bertrand.[11] Note that firms take account of cross-price effects to own products when pricing, changing the set of such products is the key mechanism in applications of this setup that are used for merger simulations.

Equations (3) and (4) described the profit flows for a set of products produced in Germany. For a product produced in another country the exchange rate is instead the one between that country and the US and for a US producer the exchange rate is equal to 1. Note that for a foreign producer that

[10] Following some sensitivity analysis, we found our results to be largely robust to the choice of M.
[11] Marginal costs are highly persistent over time. We have examined them by regressing the marginal cost of BMW and Porsche products on their lags and model fixed-effects, obtaining insignificant estimates for the autoregressive components and significant ones for the fixed-effects. This led us to use the marginal costs observed in the last year of the estimation sample in our simulations.

produced in the US, engaged in natural hedging, the maximization problem would still appear as in equation (4) with eur/usd_{jt} set to1, as prices and costs are in the same currency. The resulting profits would be translated into the home currency at the exchange rate eur/usd_{jt} but the profit maximization problem would be purely in dollars.

3.2 Counterfactual shocks

We need to take a stand on stochastic processes to generate counterfactual levels of exchange rates and consumer confidence. Consumer confidence affects demand directly whereas exchange rates only have an indirect effect via prices, as we explore further below. Note that this step is completely separate from the demand estimation. This can be useful if we want to include several business cycles to generate macroeconomic shocks but only have data on a shorter time period for the relevant product markets. We use bimonthly data for consumer confidence and the real exchange rates for the period January 1973 to July 2006 to estimate the statistical properties of these variables. Our reading of the evidence is that the forecasting ability of macroeconomic models of exchange rates is weak and we instead opt for a simpler, purely statistical, approach.[12] A number of studies have modeled exchange rate behavior over shorter horizons using autoregressive processes. A frequent finding is that a GARCH (1,1) model performs well (see for instance Hansen and Lunde (2005) or Rapach and Strauss (2008)). Patton (2006) also uses GARCH(1,1) processes to model the daily exchange rates of the US dollar against the yen and the euro.

We want the counterfactual shocks to capture the co-dependence between variables. For instance shocks to US monetary policy are likely to affect the exchange rate against both the euro and the yen. In recent years copulas have been used to model the interdependencies between asset prices (see for instance Jondeau and Rockinger (2006), Kole et al (2007) or Patton (2009) for a survey)[13]. Consider three random variables X_1, X_2, X_3. The joint cumulative density function (cdf) is given by $H(x_1, x_2, x_3)=Pr[X_1\leq x_1, X_2\leq x_2, X_3\leq x_3]$. For each X_e, $e=1,2,3$, the marginal cdf is given by $F_e(x_e)=Pr[X_e\leq x_e]$. A concern is that the standard multivariate distributions, such as the multivariate normal, would force all marginal distributions to follow the same processes. The attractiveness of the copula approach is that it allows modeling of the univariate processes separately from their dependence. The core result with regard to copulas is due to Sklar (1959) who showed that any joint distribution of random variables can be decomposed into two parts: The marginal univariate

[12] The igniting spark to the large literature on the forecasting ability of exchange rates of macro models was Meese and Rogoff's (1983) finding that a random walk beat all the proposed models. While some of the ensuing studies point to some predictive power of macro based models (for instance Mark (1995)), other studies point to very weak predictive power (Sarno and Valente (2009)).

[13] Copulas are also finding applications in marketing, see Danaher and Smith (2010).

distributions and a function, the copula function, that captures the dependency between the marginals.[14] Using C to denote the copula function we can thus write

$$H(x_1, x_2, x_3)=C(F_1(x_1), F_2(x_2), F_3(x_3)).$$

We use a multivariate t-copula to model the dependence between our three stochastic variables of interest. Define $F_e(x_e) \equiv u_e$. The t-copula is then defined by

$$C(u_1, u_2, u_3; \rho, \nu) = T_{\nu,\rho}\left(t_\nu^{-1}(u_1), t_\nu^{-1}(u_2), t_\nu^{-1}(u_3)\right)$$

where $T_{\upsilon,\rho}$ is the cdf of the multivariate Student's t distribution with correlation matrix ρ and degrees of freedom υ. The cdf of the univariate student's t distribution with υ degrees of freedom is denoted by t_υ. An attractive feature of the t copula is that it allows for a higher dependence between extreme events than for instance the Gaussian copula. As $\upsilon\rightarrow\infty$ the t copula converges to the Gaussian copula.

We use GARCH(1,l) models to estimate the exchange rate processes. Use y_{et} to denote the logarithmic returns (first-differences of logarithmic series) in the real usd/eur and real usd/jpy respectively between time t and $t-1$. We assume that the process followed by y_{et} is given by

$$y_{et} = c_e + \eta_{et}$$

$$\sigma_{et}^2 = \kappa_e + \alpha_{e1}\sigma_{et-1}^2 + \alpha_{e2}\eta_{et-1}^2.$$

So the exchange rates are assumed to follow an autoregressive process of order 1. Today's realization is equal to the last period's value plus a random shock. The error term η is assumed to follow a t distribution with mean zero and variance σ^2. We allow the shocks to have time varying volatility.

We model the process followed by consumer confidence in first differences, such that y_{ct} is the difference in consumer confidence between time t and and $t-1$.

$$y_{ct} = c_c + \eta_{ct}$$

The decreases in consumer confidence are greater than increases. To capture this asymmetry we model the shocks using an exponential GARCH model, EGARCH(1,1). Again let the error term η follow a t distribution with mean zero and define $z=\eta/\sigma$. Following Nelson (1991) we then assume that volatility can be modeled as

$$\ln(\sigma_{ct}^2) = \kappa_c + \alpha_{c1}\ln(\sigma_{ct-1}^2) + \alpha_{c2}z_{t-1} + \gamma(|z_{t-1}| - E|z_{t-1}|).$$

If γ is negative, the conditional volatility will be greater for negative shocks than for positive shocks. We fit a Student's t-copula to the residuals that we estimate by the GARCH and EGARCH processes.

[14] See for instance Nelsen (1999).

Based on the estimated GARCH processes we then generate 200 random shocks for each future period in the forecast horizon and let the correlation between shocks in each period follow the copula relation. We generate counterfactual values up to 48 months ahead from the end date July 2006.

3.3 Counterfactual profits

For each set of draws of exchange rates and consumer confidence in each period we generate counterfactual operating profits at the product level for all firms. In making forward simulations we clearly rely on a large number of assumptions. Some of the assumptions are motivated by computational concerns but many others are reflecting what we perceive to be appropriate to capture key features of the case at hand and it would be equally easy to apply other assumptions, as we discuss below.

To describe our counterfactual simulations it may be instructive to write down profits for firm *F* that is assumed to control products $j \in (F_{US} \cup F_G)$ in time *t+n* under the set of counterfactual draws *r*, where each *r* refers to a set of draws on the dollar-yen, dollar-euro and consumer confidence. Let usd/eur_{rjt+n} denote draw *r* on the dollar-euro exchange rate in *t+n.* For a German producer with some products produced in Germany and some in the US the counterfactual profits are then given by

$$\Pi_{\mathrm{rt+n}} = \underbrace{\sum_{j \epsilon F_{US}} \frac{eur}{usd}_{rt+n} \left(\tilde{p}_{jrt+n} - \widehat{mc_{jt}}\right) M_t \, s_{jrt+n}}_{US\ production} + \underbrace{\sum_{j \epsilon F_G} \frac{eur}{usd}_{rt+n} \left(\tilde{p}_{jrt+n} - \frac{usd}{eur}_{rt+n} \widehat{mc}_{jt}^{EU}\right) M_t \, s_{jrt+n.}}_{German\ production} \quad (5)$$

We set the starting date for our simulations to July 2006. For each set of draws *r* and each future time period *t+n* we calculate counterfactual profits for each model and aggregate to firm level profits expressed in the home currency. We use draws from 12, 24, 36 and 48 months ahead to calculate yearly profits for the future years. Marginal costs, quantities and prices are clearly key components in the counterfactuals and let us discuss them in turn.

3.3.1 Marginal costs

As noted in 3.1 we follow the usual procedure in the literature and assume static Nash-Bertrand prices setting by multi-product firms to back out marginal costs from the first-order condition of the firms. In our forward simulations we keep the marginal costs fixed at their July 2006 level in the currency of the country of production and are denoted by $\widehat{mc}$. For a model that was produced in Germany but that counterfactually produces in the US we assume that $\widehat{mc}_{j,1996} = \widehat{mc}_{jt}^{EU} \left(\frac{usd}{eur}\right)_{1996}$.

Germany and the US are countries at similar levels of development and with substantial car manufacturing and with limited differences in factor prices or technology. For instance, over 1992 to 2005 wages in manufacturing are on average only 6.8 percent higher in US than in Germany.[15] For these production locations the swings in exchange rates are likely to overwhelm level differences in production costs. As seen in Table 1 the usd/eur exchange rate fell from 1.4 in 1996-7 to 0.88 in 2001-2 and then rose again to 1.22 by 2005-6. Inflation is low in both countries during this time so such changes translate into cost differences of production.

There may clearly be level differences in marginal costs of production in Germany and US. The natural hedging argument is about variability rather than levels and to keep the counterfactual analysis transparent we have opted for equality of marginal costs as the benchmark.[16] A large theoretical literature examines how costs of producing in different locations depend on agglomeration economies such as the thickness of local labor markets and access to suppliers of intermediate inputs in addition to other demand and cost factors such as market access, taxes, investment subsidies, transport costs and tariffs (see for instance Fujita et al (1991) for an influential overview of the theoretical foundations of location decisions – Smith and Florida (1994) and Mayer et al (2010) are representative of a large empirical literature that points to the importance of agglomeration economies for the location of production). Costs of production will also depend on where in the US one chooses to locate – locating production in the Southern US has for instance been linked to a weaker role of unions there than in the traditional car manufacturing locations around Detroit.[17] Production in the US could also imply the import of a large number of intermediate inputs from the home country, a behavior found for instance by Blonigen (2001) in his study of foreign direct investment by Japanese manufacturers of auto parts. With detailed firm level information on costs in different locations it would be straightforward to use such information instead of the backed out marginal costs in the simulations.

3.3.2 Quantities

The vector of counterfactual demand for products in each counterfactual draw and time s_{jrt+n}, depends on the vector of counterfactual prices for all products in that counterfactual and on the counterfactual realization of consumer confidence interacted with product segments as in the demand estimation. We keep the set of products fixed in the simulations going forward. We also assume that unobserved

[15] Source: OECD, labor compensation per employee in manufacturing, expressed in USD using PPP-adjusted exchange rates.
[16] Clearly, in the forward simulations the exchange rate will have a large effect on the German cost of production from the perspective of the market.
[17] The vote by employees at Volkswagen's Tennessee plant not to join the union United Auto Workers attracted much attention in the spring of 2014 for instance (see New York Times February 14, 2014, "Volkswagen vote is defeat for labor in South").

product characteristics, ξ, are kept fixed at their 2006 levels.[18] On a case by case basis it would clearly be straightforward to do counterfactual analysis dropping some products or introducing new, "synthetic", products or substantially changing the observable characteristics of available products or consider the consequences of a successful advertising campaign. Quite possibly BMW or Porsche, if applying the method, would like to pursue such scenario analysis. Generating the whole evolution of the product portfolios in the car industry in a forward looking way would however be a formidable task and we believe that keeping the set of products fixed is a natural starting point. It is further the case that product cycles are rather long in the car industry, a typical platform remains on the market for at least five years.

3.3.3 Prices

We use hedonic regressions to generate counterfactual prices. We first regress real prices on exchange rates interacted with country of produciton, product characteristics (HP, size, transmission) and product fixed effects for 1996-2005. We then use the coefficients from these hedonic regressions to generate counterfactual prices in each of the counterfactuals.

The way in which we generate counterfactual prices raises two questions: Why not use static Nash-Bertrand in counterfactuals and why not let counterfactual prices reflect consumer confidence? Starting with the first question note that, while the type of models that we build on are generally seen as giving a plausible representation of substitution patterns on the part of consumers, static Nash-Bertrand pricing may overestimate the degree of price adjustment in many cases. In her study of the US car market, using similar tools as we do, Goldberg (1995, p. 937) for instance notes that "After 1985, the model predicts a significant increase in German import prices as a consequence of a dramatic dollar depreciation which is not matched by the data. In fact, the prices of German imports remained fairly constant during 1986 and 1987." More generally, Goldberg and Hellerstein (2008) argue that demand systems that imply more plausible substitution patterns tend to generate excessive pass-through if coupled with static Bertrand-Nash pricing. Nakamura and Zeron (2010) and Goldberg and Hellerstein (2012) introduce dynamic price adjustment in a framework similar to ours. The dynamic analysis of pricing in these models is still at the frontier of research and we would need to consider the counterfactuals not just in one baseline scenario but for a large range of counterfactual macro shocks which would be computationally demanding. Given our hope in showing how tools from empirical Industrial Organization can be applied to practical investment problems we opted for the hedonic price regressions to yield counterfactual prices in a parsimonious way. Our motivation here is somewhat related to the estimation of policy functions that form the first stage in Bajari, Benkard and Levin's (2007) dynamic oligopoly model. We nevertheless found that using backed out

[18] One could also assume that consumers had preferences over where a product is produced. One the one hand a "buy American" preference might lead to a higher appreciation for a locally produced car. On the other hand it may well be that some of the mystique of top level sports cars such as Porsche is linked to the origin.

marginal costs was preferred to possible alternatives such as relying on accounting costs or making assumptions directly on the marginal costs. As we document below, the marginal costs that we recover are well in line with the previous literature.[19]

The second parsimonious adjustment that we make is to not include consumer confidence in the hedonic price regressions. In preliminary hedonic regressions we included the same interactions between segments and consumer confidence as in the demand regressions. These interactions were not significant however and using the point estimates to generate the counterfactual prices resulted in excessive variability of prices and profits. Limited and delayed responses of prices to demand shocks is the subject of a large literature in macroeconomics (see for instance Blinder (1998) for survey evidence or Nakamura and Steinson (2013) for an overview of the micro evidence on sticky prices). Menu costs of adjusting prices and implicit contracts between consumers and producers are just two of a host of possible explanations (see Okun (1981) for a seminal reference on implicit contracts or Rotemberg (2011) for a model of how ongoing customer relations can imply a lack of response to demand shocks). We use list prices which are likely to be even less affected by demand shocks than transaction prices (which may feature rebates). The evidence points however to that even transaction prices are very unresponsive to demand shocks - Copeland and Hall (2009) use transaction prices for the big three US carmakers and show that demand shocks have only a small impact on price and are absorbed almost entirely by sales and production decisions.

Discounting

Feeding the draws into the demand system, and letting prices adjust, yields a probability distribution of profits for each future time period. To analyze various strategies we calculate the discounted profit flows under different assumptions on production patterns. Many different alternatives are possible when determining the correct discount rate for a risky investment. Our goal here is not to add to the literature on the determination of discount rates but rather we note that the weighted average cost of capital (WACC) method is commonly used to discount cash flows in corporate finance (see for instance Damodaran (2010)). We use balance sheet information from the annual reports for 2005 (BMW) and 2005-2006 (Porsche) to deduce the the Cost of capital = cost of equity × (equity/(debt+equity) + cost of debt × (debt/(equity+debt)). The cost of equity is calculated using the CAPM relation where cost of equity = risk-free rate + beta *mature market equity risk premium. As risk free rate we use the 10 year German bund (interest rate of 4.05 in July 2006) and following Damodaran (2010) we use 4.05% as the mature market risk premium. Betas are 1.087 for BMW and

[19] In contrast to many other applications of the demand models, where all other macro variables and cost shifters are kept constant (such as in Nevo (2000) or Petrin (2002)), price responses to macro variables are a driver of results in our study. Static oligopoly may well be appropriate for predicting price effects of changes in the number of competitors or the set of competing products but not for price effects of macroeconomic fluctuations.

1.251 for Porsche (calculated on monthly data using DAX 1988:10 to 2006:6). The resulting discount rates are 5.66 for BMW and 5.93 for Porsche.

4. The estimated model

4.1 Demand Estimates

Table 4 reports estimates of two RC logit specifications for the US car market. Both use price, engine power (HP), size and whether non-manual transmission is included in the baseline model as observable product characteristics. We model price as a random coefficient with a mean effect of price on utility and individuals' coefficients on price follow a Normal distribution as outlined in Section 2. Both specifications also include time (model-year) and brand fixed-effects. We treat price as endogenous in our demand specification.

To estimate our model, besides the exogenous characteristics, we use the BLP instruments consisting of the number of firms operating in the market, the number of other products of the same firm and the sum of characteristics of products produced by rival firms. As documented in the literature (Berry 1994, BLP 1995), not accounting for the endogeneity of prices results in an attenuation bias, that is, the price coefficient is biased towards zero, and this is what our findings also suggest: the uninstrumented version of Specification I has a price coefficient of -0.002, well below the instrumented ones at -0.021. Besides the tenfold increase in the slope of the demand curve, at 27.52 (and significant at the one percent level), the F-statistic of the first-stage regression of price on the exogenous regressors is well above the rule-of-thumb value of 10 suggested by Staiger and Stock (1997). This suggests that instruments are not weak and that there is no evidence that the instrumented price coefficient is biased towards the uninstrumented one. Instruments are also not rejected when computing tests of overidentifying restrictions, as reported in Table 4.

[Table 4 about here]

The stance in which Specifications I and II differ is in the treatment of consumer confidence and market segment variables. Specification I uses consumer confidence and separate fixed-effects for market segments. In contrast, Specification II uses interactions of market segments and consumer confidence. Specification II thus allows asymmetric responses in market shares according to the market segment a model belongs to, according to which economic outlook consumers expect to prevail. Both specifications have significant coefficients for the mean and for the dispersion of price coefficients, whereas the remaining characteristics are usually not significant. In fact, most of the explanatory power for market shares tends to come from brand and market segment fixed-effects.

The (own) price elasticities (equivalently, markups) of the models in Specification II are in the range 3.7-7.3 with an average elasticity 6.0, thus broadly in line with previous studies of the car industry, notably Petrin's (2002) RC logit estimates using micro data (see, for instance, column 6 of his Table 9).[20] Interestingly, the estimates for Specification II suggest an intuitive "pecking order" effect of the interaction terms. For instance, demand for the "Upper Luxury" segment tends to be more sensitive to consumer confidence than that of the "Middle Luxury" segment, which in turn is more sensitive than that of the "Lower Luxury" segment. We interpret these results as evidence that, conditional on buying a car, consumers are more likely to purchase models from high-end segments the more confident they are about the economic outlook.

4.2 Counterfactual shocks

As described above the first step in generating the counterfactual draws is to fit univariate processes for exchange rates and consumer confidence. The estimation output for the marginal distributions is given in Table 5. The significant coefficient on lagged volatility in the usd/eur relation points to that volatility is indeed time-varying at this frequency. The process for consumer confidence reflects a pattern where the typical change is an upward drift but that negative shocks are associated with greater volatility (captured by the negative coefficient on the leverage term).

[Table 5 about here]

We then fit a t-copula to the residuals from the univariate relations. The degrees of freedom for the t-copula are estimated to be 21.65. The estimated correlation coefficients using the t-copula are -0.085 between usd/eur and consumer confidence, 0.063 between usd/jpy and consumer confidence and 0.522 between usd/eur and usd/jpy. Combining these estimates allows us to generate counterfactual shocks where the marginal distributions follow the GARCH processes and the co-dependence follows a t-copula in each future period. Adding the succession of these shocks to the starting values in July 2006 then gives us 200 counterfactual paths of the exchange rates and consumer confidence. As an example of our results, Figure 2 shows the distributions for counterfactual draws for these three variables 12 months ahead from July 2006. The histograms show the densities for the respective variable and the scatter plots show the relation for each bilateral comparison. The scatter plot in the lower left hand corner for instance plots counterfactual draws of usd/eur against counterfactual draws of usd/jpy.

[Figure 2 about here]

As seen, the draws reflect substantial dispersion for all three variables. The skewness of consumer confidence is visible. The starting value in July 2006 is 134 and we see predictions for 12 months ahead centered at this level (median across the draws is 146, mean 139) but a long tail of weaker

[20] Goldberg (1995) finds elasticities in the range 1.1-6.2 across specifications and market segments, using data from 1983-1987.

realizations. As seen in the scatter plots in the middle row, the relation between consumer confidence and the exchange rates is weak. The positive relation between the two exchange rates on the other hand is clearly visible in the scatter plots in the upper right and lower left corner. These then are the counterfactual levels of macro variables that are fed into the demand system when we consider the 12 month horizon ahead. Note that by the additive nature of the shocks we can view our results as simulating 200 possible paths of the underlying variables. As we expand the forecast horizon some of the paths for consumer confidence are predicted to be too low, or even negative. In these cases we replace the value with a hypothesized lower threshold of 10. The lowest level in the time period covered by our data is 15.8 (December 1982).

4.3 Price setting

We use the coefficients from hedonic regressions to generate counterfactual prices. We regress real prices on real rates interacted with production location plus product characteristics (HP, size, transmission) and product fixed effects. The results of three specifications are reported in Table 6. Specification 1 includes interactions of product location and exchange rates, in addition to model fixed-effects. Specification 2 also includes the characteristics used in the demand model whereas our preferred specification – Specification 3 – also uses the estimated unobserved product characteristics from the demand system as an additional covariate, its rationale being that despite being unobserved by the econometrician, it is observed by market participants. Overall, the results are similar and the good fit of the model relies on the presence of model fixed-effects.

To gauge if the results are reasonable we report the implied elasticities of all specifications. The exchange rate pass-through starts at 0.0812 and 0.113 for the Euro and Yen exchange rates in Specification 1, respectively, barely changes when product characteristics are included as per Specification 2 and gravitate around 0.10 for both exchange rates in Specification 3.[21] All pass-through terms are significant at the 1 percent level. Comparing to other estimates, our estimates are somewhat on the low side. A number of studies examine pass-through in import prices (see Goldberg and Knetter (1997) for an early survey) and find pass-through elasticities that are frequently equal to about one half. Note however that pass-through at the border is typically substantially higher than measured pass-through at the retail level. We can also compare to another non-structural estimate for the US auto market, Hellerstein and Villas-Boas (2010). The 24 models in their study exhibit an average pass-through of exchange rates into transaction prices of around 38 percent, but with large standard deviations.

[21] The practical implication of such similarity across specifications is the robustness across simulation results. In particular, the inclusion of estimated unobserved product characteristics in the hedonic model does not materially affect the results.

[Table 6 about here]

5. Simulation Results

We now turn to a presentation of the simulation results, feeding the counterfactual shocks into demand and costs and letting all prices respond. We compare different production scenarios as to what models are produced locally in the US – first in terms of per period profits and then in terms of their present discounted values (PDVs). It deserves to be emphasized that we examine only profits from the US market, i.e. we focus on the US operations of carmakers BMW and Porsche.

5.1 Per period distributions of profits

5.1.1 Profit distributions at different horizons

We start by considering simulated profits up to 4 years ahead for both BMW and Porsche. In generating these counterfactuals we use data up to July 2006 only, so the counterfactual profits for 2007 is one year out and, for 2010, 4 years out. A useful way of presenting simulated cash flows is to examine their frequency distribution, see Figure 3. In Panel 3a (3b) we graph kernel density estimates of simulated cash flows for BMW (Porsche) for their current production strategies at different horizons. For BMW, most of the production occurs in the EU, whereas for Porsche, the whole production is in Europe. As intuitively expected, the increased dispersion of the risk factors further in the future translates into more dispersed profit distributions for longer horizons.

[Figure 3 about here]

5.1.2 Profit distributions for alternative production strategies: The role of natural hedging

We now fix the time dimension (to a 36-month ahead horizon) and focus on alternative production strategies for both BMW and Porsche. Thus, Figure 4 displays the profit distributions of current production, producing entirely in the EU and producing entirely in the US for both car makers, together with the underlying eur/usd exchange rate. Intuitively, the more is produced in the consumer market (US), the less dispersed the profit distribution becomes as a result of natural hedging. The results for BMW, displayed in Panel 4a, show that even producing two models in the US (models X5 and Z4 in this case) already reduces the downside of profits for BMW in a non-trivial way.

[Figure 4 about here]

The more production is shifted to the US, the lower the exchange rate exposure of profits, very much in the spirit of the discussion surrounding equations (1) and (2). That is, the role of producing some models in the US is to decrease the cash flow sensitivity to exchange rates by reducing the weight of both tails of the profit distribution; this results in a lower standard deviation of the simulations as local production increases -- a testament to the fact that producing in the US can be seen as a natural hedge.[22] As a result, the profit distributions become less dispersed. Importantly, the BMW cash flows attain negative values in over 5 percent of the simulations for both the "Current" and the "All in EU" scenarios, but not in the "All in US" scenario; thus, natural hedging has the attractive property of avoiding the realization of negative cash flows.

The results for Porsche, displayed in Panel 4b are in line with those of BMW in that increasing production in the US reduces cash flow sensitivity to the usd/euro exchange rate. As for BMW, over 5 percent of the simulations result in negative profits in the current scenario; this suggests that, for a decision maker that attaches a larger weight to outcomes in the lower tail of the distribution, natural hedging appears an attractive strategy.

Panel 4c shows the underlying eur/usd exchange rate. By comparing it to the "All in US" scenarios for both BMW and Porsche, one can see how those scenarios essentially inherit the fluctuations – see in particular the pronounced upper tail – of the underlying exchange rate.

5.2 Discounted profits under different production locations

The previous section illustrated one use for the simulation tools that we develop, namely to generate probability distributions for cash flows that we can use to examine risk at different horizons and under different scenarios. In the present section we use the counterfactual values to compare the scenarios over the lifetime of a strategy. As explained in Section 3 we calculate discount rates for BMW and Porsche using the WACC method. We use the discount rates to calculate the PDV of profits under each of the 200 streams of profits and illustrate these distributions in Figure 5. In addition to the extreme production strategies of producing only in the EU or only in the US, we consider the current production strategies of both BMW and Porsche and a hypothetical strategy according to which carmakers can flexibly switch production between the US and the EU according their attractiveness. That is, the flexible EU-US production scenario we consider consists of producing entirely in the US or in the EU depending on which strategy yields higher profits. Both panels in Figure 5 illustrate the effect of natural hedging in that profit distributions become less dispersed.

[22] The profit distributions resulting from the intermediate strategies between "Current" and "All in US" strategies (for BMW) and "All in EU" and "All in US" (for Porsche) are omitted for the sake of clarity but available from the authors upon request.

[Figure 5 about here]

In the case of BMW, reported in Panel 5a, simply producing models X5 and Z4 in the US as opposed to producing entirely in the EU increases mean profits by over EUR 3bn (= EUR 40.1bn – EUR 36.5bn) with a decrease of another EUR 3.3bn in the standard deviation. The counterfactual scenario of producing entirely in the US would result in mean profits and standard deviation of EUR 75.0bn and EUR 6.3bn, respectively, whereas flexible production would yield EUR 77.7bn and 12.3bn, respectively. That is, both strategies dominate scenarios with production in the EU in the mean-variance sense.

The results for Porsche are qualitatively similar to those of BMW; natural hedging or flexible production will shrink the profit distribution as compared to producing in the EU. In particular, shifting the whole production from the EU to the US would result in an increase of EUR 4.0bn (=EUR 9.0bn – EUR 5.0bn) in mean profits and a reduction of EUR 3.5bn (= EUR 4.0bn – EUR 0.5bn) in the standard deviation of the profit distribution.

The above findings reflect the earlier results for one-period profits, in that producing more in the consumer market reduces the dispersion of the profit distribution. Although a rigorous analysis would require further knowledge of the fixed costs of setting up a plant, the magnitudes involved suggest that the gains accrued by pursuing natural hedging are substantial. To gain perspective, the cost of establishing Volkswagen's new plant in Chattanooga was USD 1bn (equivalent to about EUR 0.7bn at the prevailing exchange rate in January 2010[23], whereas in 2005, BMW opened a new plant in Leipzig, Germany, with a total investment of EUR 1.3bn prior to its opening (Annual report 2005, p 19).

6. Concluding comments

This paper proposes a structural model to quantify the exposure of firms to risk factors affecting their profits. In our illustrative application, we show that, under our assumptions, a decision to produce in the US is easily motivated for BMW but not for Porsche. The key insight of the paper is that by feeding draws from the distribution of risk factors through a demand system, rather than having them directly affect sales or market size, many of the weaknesses of simulation methods to evaluate risky investments are muted.

We have made a number of simplifying assumptions, most of which for convenience. We only considered the US market for instance and assumed a simple cost structure. Time and resource constraints hindered us from assembling similar quality data for BMW's and Porsche's other markets.

[23] See New York Times, "Students See a Creek and Imagine a Bridge for VW", Jan 26 2010.

Conveniently, the method can be implemented by using data that are typically available for purchase, such as sales, prices and characteristics of products. Using more detailed information – typically available to firms, but not researchers – is bound to increase the accuracy of any such exercise. For instance if a firm were to perform calculations such as these for itself, it would want to make use of its knowledge of the cost structure.

The method that we propose may also be useful input to firms' decisions on financial hedges. Reasons for hedging may be to smooth tax payments, avoid bankruptcy or to ensure sufficient cash flow to finance investments also in tough times (see Stulz (2002) for an overview of the arguments and Tufano (1996) or Adam and Fernando (2006) for empirical examinations of the motivations for hedging and its effects on firm value). In the current paper we have disregarded the why's, the when's and the how's of financial hedges. These are important issues but before taking a view on how to use financial hedges one needs to understand the relation between profits and risk factors. We focus on this first step in the decision process. In a second step one could use the counterfactual profits that we generate to evaluate different strategies for financial hedging. Brealey and Kaplanis (1996) do such comparisons for a simple stylized example and this may be one use of the counterfactual flows like the one that we present.

References

Ackerberg, Daniel, C. Lanier Benkard, Steve Berry and Ariel Pakes (2007), Econometric Tools for Analyzing Market Outcomes, in: James J. Heckman & Edward E. Leamer (ed.), Handbook of Econometrics. Amsterdam: Elsevier.

Aguirregabiria, Victor and Aviv Nevo (2012), Recent developments in empirical IO: Dynamic demand and dynamic games, in Advances in Economics and Econometrics: Theory and Applications. Tenth World Congress of the Econometric Society, forthcoming.

Bajari, Patrick, C. Lanier Benkard and Jonathan Levin (2007), Estimating dynamic models of imperfect competition, Econometrica 75, 1331-1370.

Benkard, C. Lanier, Aaron Bodoh-Creed and John Lazarev (2010), Simulating the Dynamic Effects of Horizontal Mergers: U.S. Airlines, manuscript, Stanford University.

Berry, Steve (1994), Estimating discrete choice models of product differentiation, RAND Journal of Economics 25, 242-262.

Berry, Steve, James Levinsohn and Ariel Pakes (1995), Automobile prices in market equilibrium, Econometrica 63, 841-890.

Berry, Steve, James Levinsohn and Ariel Pakes (1999), Voluntary export restraints on automobiles: Evaluating a trade policy, American Economic Review 89, 400-430.

Berry, Steve and Panle Jia (2010), Tracing the woes: An empirical analysis of the airline industry, American Economic Journal: Microeconomics 2, 1-43.

Besanko, David, Ulrich Doraszelski, Lauren Xiaoyuan Lu and Mark Satterthwaite (2010), On the role of demand and strategic uncertainty in capacity investment and disinvestment dynamics, International Journal of Industrial Organization 28, 383-389.

Bodnar, Gordon M., John Graham, Campbell R. Harvey and Richard C. Marston (2011), Managing risk management, manuscript, Duke University.

Brealey, Richard A. and Stewart C. Myers (2003), Principles of Corporate Finance, 7th edition, New York: McGraw-Hill.

Busse, Meghan, Jorge Silva-Risso and Florian Zettelmeyer (2006), $ 1,000 cash back: The pass-through of auto manufacturer promotions, American Economic Review 96, 1253-1270.

Chen, Frank Youhua and Candace Arai Yano (2010), Improving supply chain performance and managing risk under weather-related uncertainty, Management Science 56, 1380-1397.

Chintagunta, Pradeep K. and Harikesh S. Nair (2011), Discrete-choice models of consumer demand in marketing, Marketing Science 30, 977–996.

Chod, Jiri, Nils Rudi and Jan A. Van Mieghem (2010), Operational flexibility and financial hedging: Complements or substitutes?, Management Science 56, 1030-1045.

Cochrane, John H. (2011), Presidential address: Discount rates, Journal of Finance 65, 1047-1108.

Copeland, Adam and George Hall (2011), The response of prices, sales and output to temporary changes in demand, Journal of Applied Econometrics 26, 232-269.

Damodaran, Aswath (2010), Applied Corporate Finance, 3d edition, John Wiley & sons.

Danaher, Peter J. and Michael S. Smith (2011), Modeling multivariate distributions using Copulas: Applications in marketing, Marketing Science 30, 4-21.

Davis, Peter and Eliana Garcés (2010), Quantitative Techniques for Competition and Antitrust Analysis, Princeton: Princeton University Press.

Dominguez, Kathryn M.E. and Linda L. Tesar (2006), Exchange rate exposure, Journal of International Economics 68, 188-218.

Eiteman, David K., Arthur I. Stonehill and Michael H. Moffett (2007), Multinational Business Finance, 11th edition. New York: Pearson.

Ericson Richard and Ariel Pakes, (1995), Markov-perfect industry dynamics: A framework for empirical work, Review of Economic studies 62, 53-82.

Eun, Cheul S. and Bruce G. Resnick (2007), International Financial Management, Second Edition, New York: McGraw-Hill.

Friberg, Richard and Mattias Ganslandt (2007), Exchange rates and cash flows in differentiated product industries: A simulation approach, Journal of Finance 62, 2475-2502.

Goldberg, Pinelopi K. (1995), Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry, Econometrica 63, 891-951.

Goldberg, Pinelopi K. and Michael M. Knetter (1997), Goods prices and exchange rates: What have we learned?, Journal of Economic Literature 35, 1243-1272.

Goldberg, Pinelopi K. and Rebecca Hellerstein (2012), A Structural approach to identifying the sources of local-currency price stability, forthcoming, Review of Economic Studies.

Hall, George J. (2000), Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants, Journal of Monetary Economics 45, 681-716.

Hankins, Kristine W. (2011), How do financial firms manage risk? Unraveling the interaction of financial and operational hedging, Management Science 57, 2197-2212.

Hansen, Peter R. and Asger Lunde, (2005), A forecast comparison of volatility models: Does anything beat a GARCH(1,1)?, Journal of Applied Econometrics 7, 873-879.

Hellerstein, Rebecca and Sofia Berto Villas-Boas (2010), Outsourcing and pass-through, Journal of International Economics 81, 170-183.

Hertz, David B. (1964), Risk analysis in capital investment, Harvard Business Review, 95-106.

Huchzermeier, Arnd and Morris A. Cohen (1996), Operational flexibility under exchange rate risk, Operations Research 44, 100-113.

Jin, Yanbo and Philippe Jorion (2006), Firm value and hedging: Evidence from U.S. oil and gas producers, Journal of Finance 55, 107-152.

Jondeau, Eric and Michael Rockinger (2006), The copula-GARCH model of conditional dependencies: An international stock market application, Journal of International Money and Finance 25, 827-853.

Jorion, Philippe (1990), The exchange rate exposure of US multinationals, Journal of Business 3, 331-345.

Jorion, Philippe (2006), Value at Risk: The New Benchmark for Managing Financial Risk, third edition, New York: McGraw-Hill.

Kole, Erik, Kees Koedijk and Marno Verbeek (2007), Selecting copulas for risk management, Journal of Banking and Finance 31, 2405-2423.

Ludvigson, Sydney C. (2004), Consumer confidence and consumer spending, Journal of Economic Perspectives 18, 29-50.

Mansfield, Edwin, W. Bruce Allen, Neil A. Doherty and Keith Weigelt (2009), Managerial Economics: Theory, Applications and Cases, Seventh edition, W.W. Norton: New York.

McAfee, R. Preston (2002), Competitive Solutions: The Strategist's Toolkit, Princeton: Princeton University Press.

Mello, Antonio S., John E. Parsons and Alexander J. Triantis (1995), An integrated model of multinational flexibility and financial hedging, Journal of International Economics 39, 27-51.

Moel, Alberto and Peter Tufano (2002), When are real options exercised? An empirical study of mine closings, Review of Financial Studies, 35-64.

Moffett, Michael H. and Barbara S. Petitt, (2004) Porsche exposed, mini-case, Thunderbird School of Global Management.

Nakamura, Emi and Dawit Zerom (2010), Accounting for incomplete pass-through, Review of Economic Studies 77, 1192-1230.

Nelson, Daniel B. (1991), Conditional heteroskedasticity in asset returns: A new approach, Econometrica 59, 347-370.

Nevo, Aviv (2000), Mergers with differentiated products: the case of the ready-to-eat cereal industry, RAND Journal of Economics 31, 395–421.

Okun, Arthur M (1981), Prices and Quantities: A Macroeconomic Analysis, Washington D.C.: Brookings Institution.

Patton, Andrew (2009), Copula-Based Models for Financial Time Series, in Torben G. Andersen, Richard A. Davis, Jens-Peter Kreiss and Thomas Mikosch (eds.) Handbook of Financial Time Series, New York: Springer Verlag.

Petrin, Amil (2002), Quantifying the benefits of new products: The case of the Minivan, Journal of Political Economy 11, 705-729

Porter, M.E. (1980), Competitive Strategy, Free Press, New York, 1980

Slade, Margaret (2001), Managing projects flexibly: An application of real-option theory to mining investments, Journal of Environmental Economics and Management 41, 193-233.

Staiger, Douglas and James H. Stock (1997), Instrumental variables regression with weak instruments, Econometrica 65, 557-586.

Stein, Jeremy, Stephen E. Usher, Daniel La Gattuta and Jeff Youngen (2001), A comparables approach to measuring cashflow-at-risk for non-financial firms, Journal of Applied Corporate Finance 13, 100-109.

Stulz, René M. (2002), Risk Management and Derivatives, Mason OH: Thomson Southwestern.

Train, Kenneth E. and Clifford Winston (2007), Vehicle choice behavior and the declining market share of U.S. automakers, International Economic Review 48, 1469-1496.

Tversky, Amos and Daniel Kahneman (1971), Belief in the law of small numbers, Psychological Bulletin 76, 105-110.

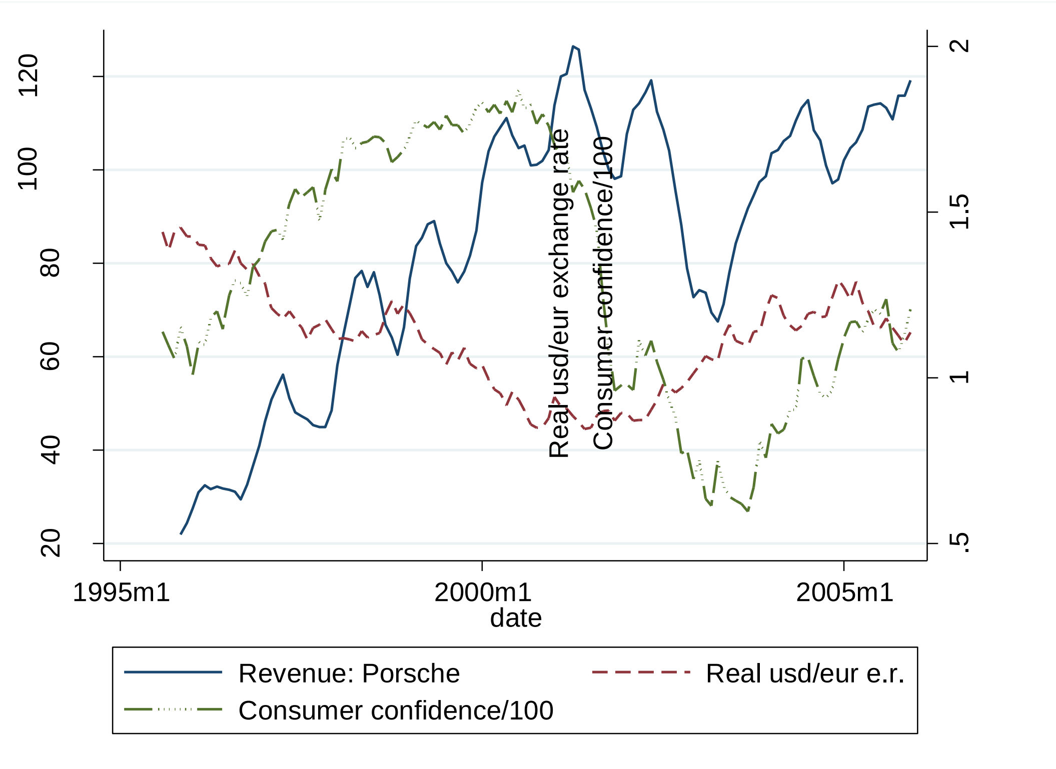

Figure 1. Porsche's monthly revenue (in euro) from US sales, consumer confidence and the US dollar – euro real exchange rate 1995-2006.

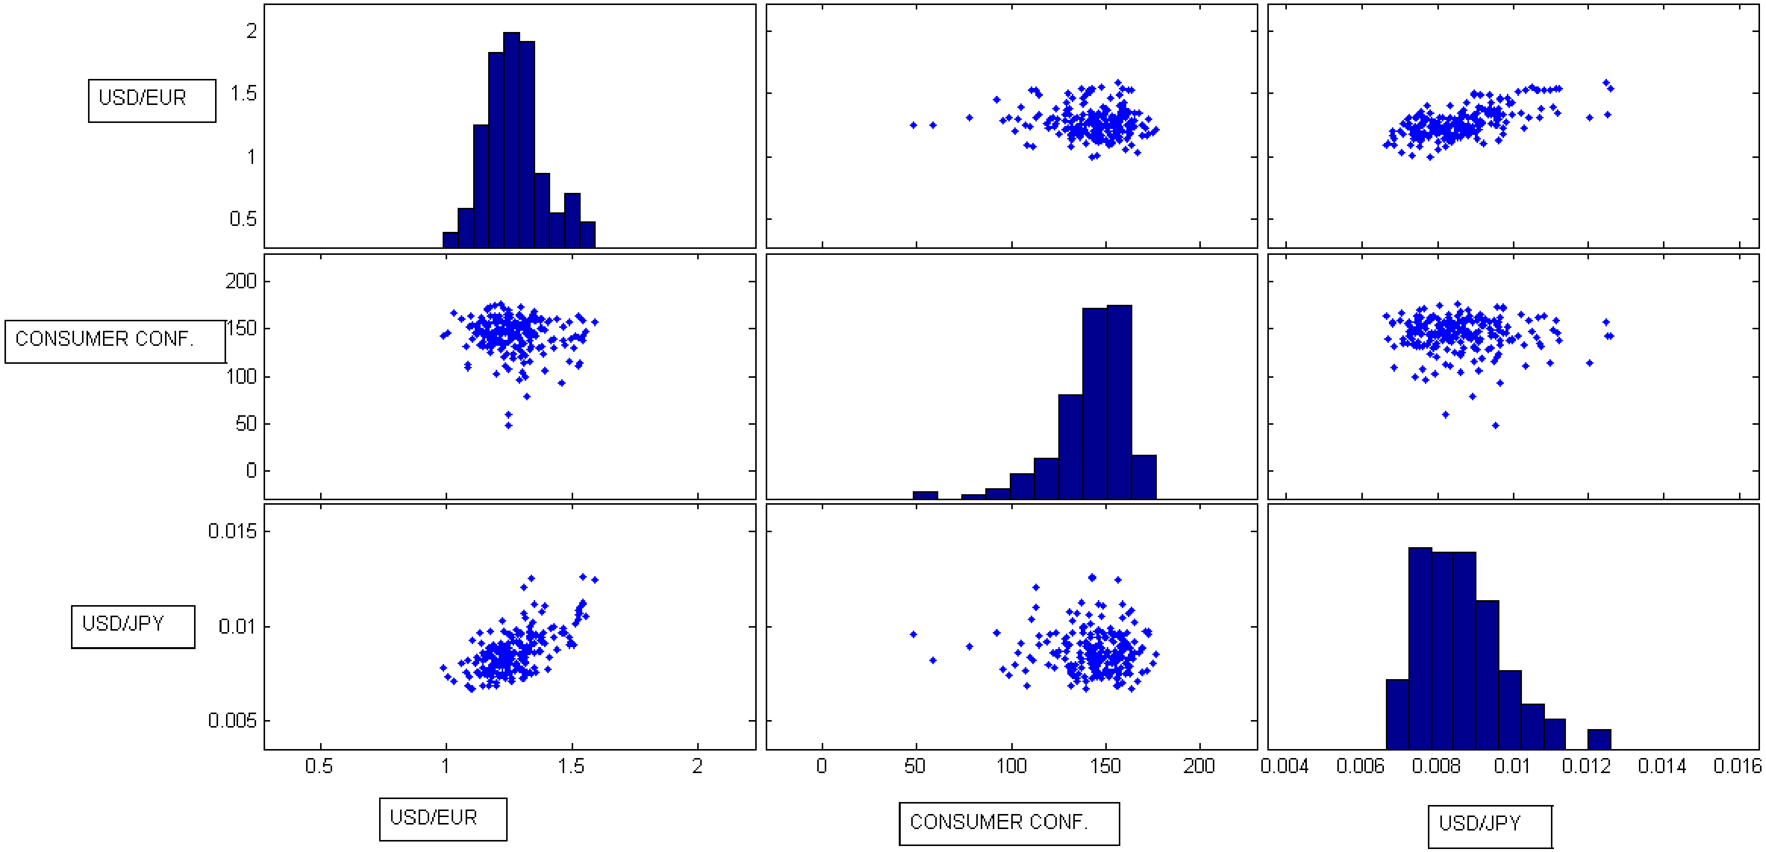

Figure 2. Counterfactual values of exchange rates and consumer confidence at the 12 month forecast horizon using July 2006 as start date.

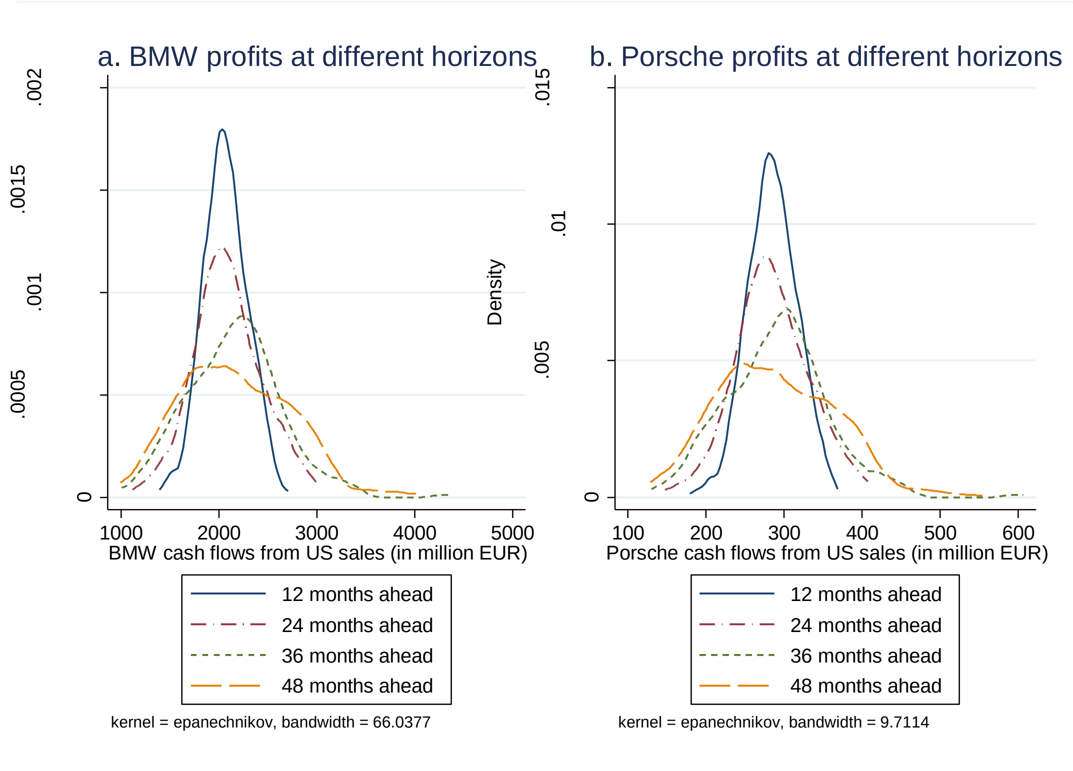

Figure 3. Profit distributions for the current strategies of BMW and Porsche at different horizons. Panel a displays profits stemming from BMW producing only the model X5 in the US. Panel b displays profits stemming from Porsche producing all models produced in the EU.

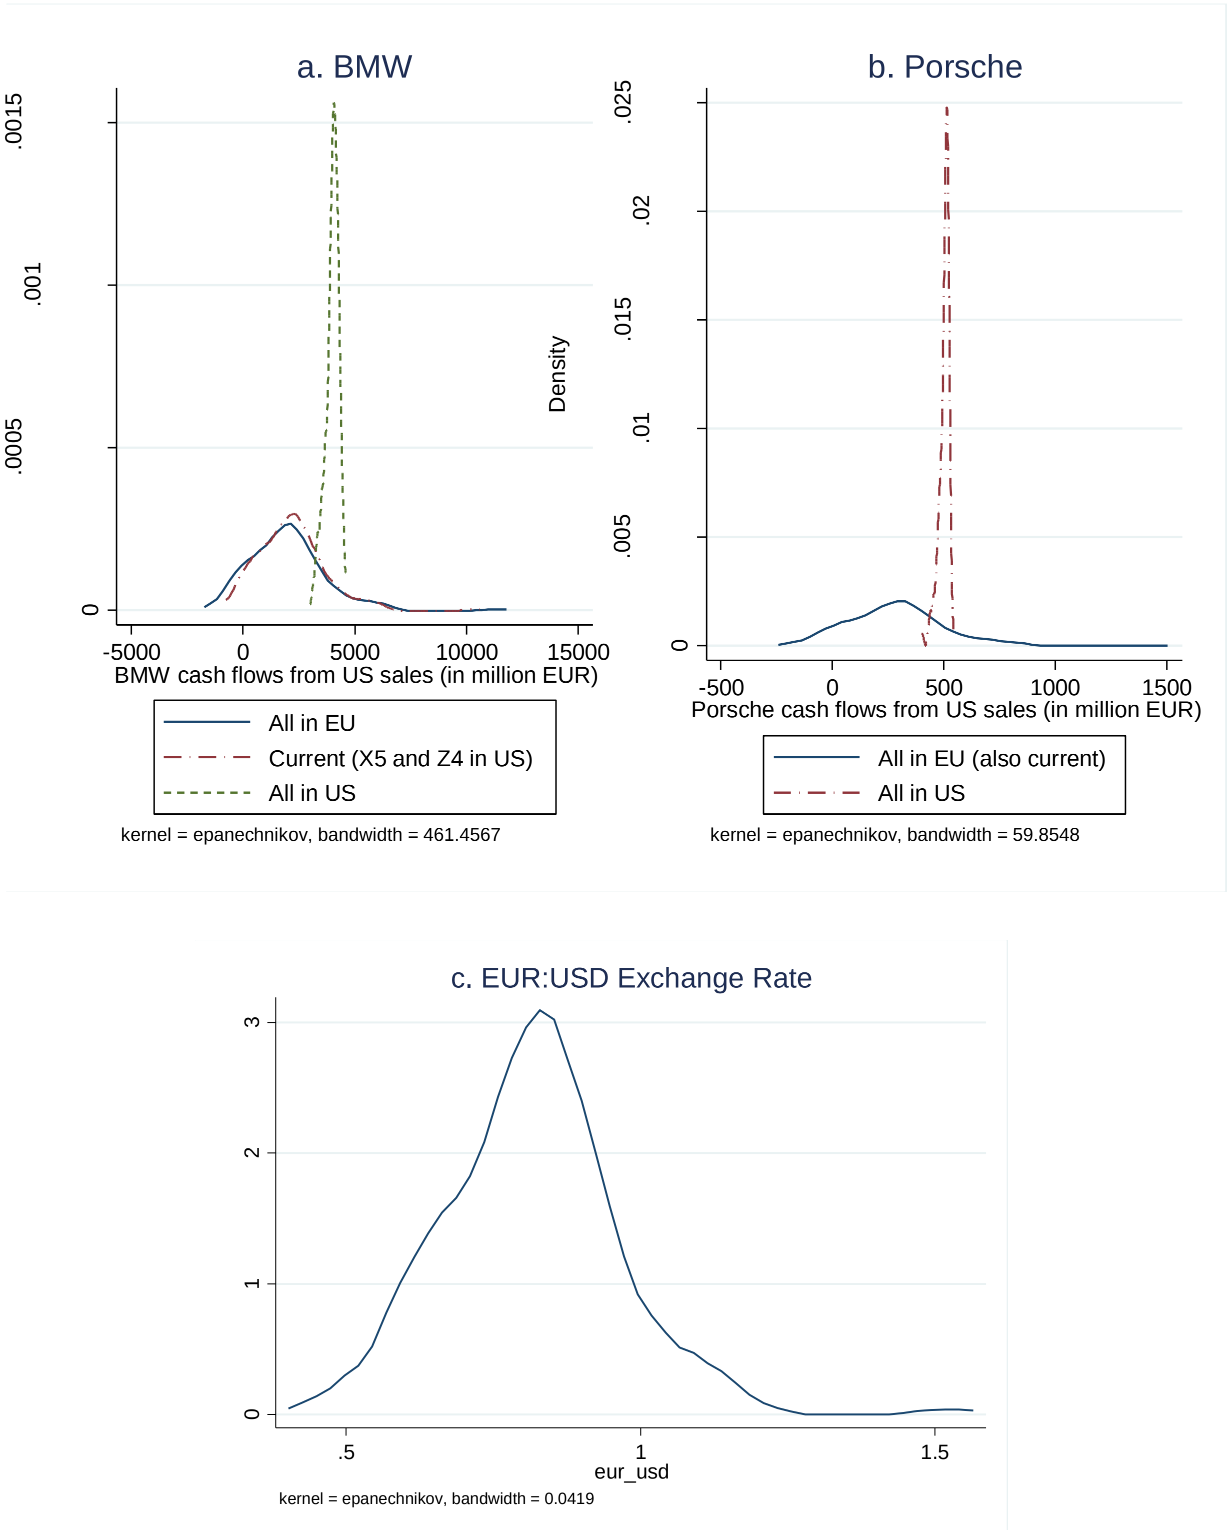

Figure 4. Counterfactual profit distributions for BMW and Porsche under alternative strategies at the 36-month horizon, and the underlying EUR/USD exchange rate. Panel a displays profit distributions for BMW. Panel b displays profit distributions for Porsche. Panel c displays the EUR/USD exchange rate. Panels a and b illustrate the effect of natural hedging in that profit distributions become less dispersed under the "All in US" scenarios. On the other hand, the "All in EU" strategies inherit the EUR/USD exchange rate fluctuations, in particular its long upper tail.

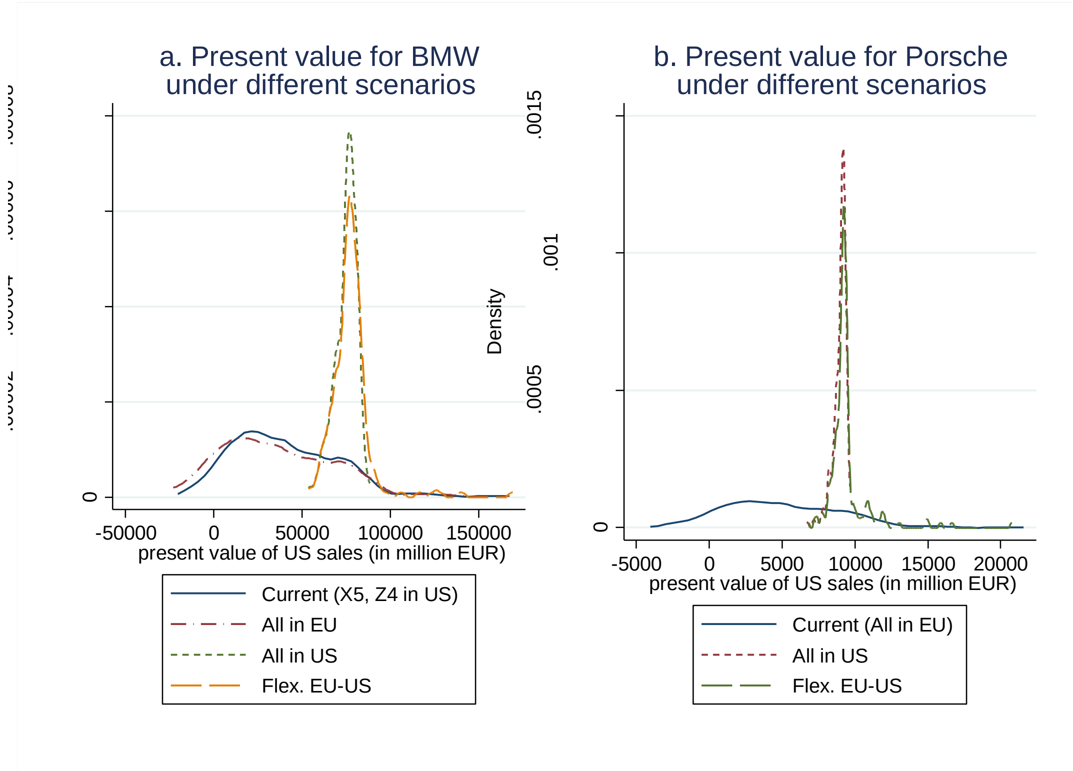

Figure 5. Distribution of present discounted values (PDVs) of alternative production strategies for BMW and Porsche. Panel a displays PDV distributions for BMW. Panel b displays PDV distributions for Porsche.

Table 1. Descriptive statistics, top segments of the US car market 1995-2006.

Model year	Price per model				Number of cars sold per model		# models	usd/euro	usd/jpy*100	Consumer Confidence
	Mean	SD	Min	Max	Mean	SD				
1995-6	38,042	18,481	12,708	87,966	31,530	57,076	82	1.3970	1.0522	112.88
1996-7	37,404	17,539	14,269	90,907	32,093	55,076	94	1.2591	.9289	140.65
1997-8	37,171	17,072	14,261	89,372	34,397	56,395	102	1.1342	.8355	164.11
1998-9	35,900	16,202	13,966	87,526	38,519	62,108	102	1.1461	.8701	173.16
1999-0	35,634	16,256	13,479	85,887	40,981	61,990	112	1.0056	.9495	180.08
2000-1	35,609	17,303	13,074	124,516	39,472	54,538	122	.8797	.8466	168.76
2001-2	34,669	17,604	12,924	124,900	43,623	58,617	121	.8985	.7599	108.28
2002-3	35,100	17,904	14,150	123,547	41,265	56,877	136	1.0422	.7754	73.05
2003-4	37,679	34,670	13,739	399,968	40,192	54,703	148	1.1687	.8213	82.72
2004-5	36,959	34,024	12,330	390,308	39,730	51,516	155	1.2204	.8185	109.13
2005-6	35,936	31,964	12,563	374,789	34,189	40,638	168	1.1487	.7331	125.42

Descriptive statistics is over models per 12 month period running from August to July. Prices in real 2000 dollars.

Table 2. Price, quantity and revenue share, BMW Brand, US market, 1995-1996 and 2005-2006.

	Price		Quantity		Share of revenue	
	1995-6	2005-6	1995-6	2005-6	1995-6	2005-6
3 series	23,562	26,730	49,868	118,377	0.22	0.31
5 series	41,273	36,544	28,439	56,266	0.22	0.20
6 series		61,713		9,741		0.06
7 series	65,472	61,332	18,478	19,270	0.23	0.12
8 series	81,495		552		0.01	
Z3	31,727		20,827		0.12	
Z4		29,606		10,215		0.03
Z8		111,840		9		0
X3		31,721		29,257		0.09
X5		36,544		34,326		0.12

Prices are in real 2000 dollars.

Table 3. Price, quantity and revenue share, Porsche Brand, US market, 1995-96 and 2005-06.

	Price		Quantity		Share of revenue	
	1995-6	2005-6	1995-6	2005-6	1995-6	2005-06
911	68,222	60,994	6,828	11,995	0.71	0.43
Boxster	43,718	38,743	4,500	5,770	0.29	0.13
Cayenne		36,392		12,501		0.27
Cayman		50,503		4,372		0.13
Carrera GT		374,788		208		0.05

Prices are in real 2000 dollars.

Table 4. Demand estimates, US car market 1995-2006. Random-coefficients logit model.

Variables	I	II	Examples	
Price	**-0.021**	**-0.030**		
	[-2.804]	[-4.475]		
HP	0.002	0.010		
	[0.247]	[1.394]		
Size	0.067	0.049		
	[1.360]	[0.790]		
Transmission	0.000	0.000		
	[.367]	[-0.754]		
Sigma price	**0.008**	**0.009**		
	[4.844]	[5.345]		
CC	0.016	-		
	[1.145]			
CC x Upper Luxury	-	**0.029**	Audi A8	BMW 7 Series
		[2.428]		
CC x Middle Luxury	-	**0.020**	Audi A6	BMW 5 Series
		[2.383]		
CC x Lower Luxury	-	**0.011**	Audi A4	BMW 3 Series
		[1.737]		
CC x Luxury Sport	-	**0.020**	Mercedes SLK	Porsche 911
		[1.843]		
CC x Luxury Specialty	-	0.013	Lexus SC430	Mercedes CLK
		[1.457]		
CC x Small Specialty	-	0.001	Mini Cooper	VW Beetle
		[0.079]		
CC x Large Luxury CUV	-	**0.016**	Acura MDX	Cadillac Escalade
		[2.166]		
CC x Middle Luxury CUV	-	**0.015**	Lexus RX330	Porsche Cayenne
		[2.108]		
CC x Large CUV	-	**0.020**	Chrysler Pacifica	Honda Pilot
		[2.274]		
CC x Middle CUV	-	0.012	Ford Escape	Hyundai Santa Fe
		[1.524]		
CC x Small CUV	-	0.005	Mitsub. Outlander	Toyota RAV4
		[0.623]		
CC x Large Luxury SUV	-	**0.030**	Cadillac Escalade	Range Rover
		[2.319]		
CC x Middle Luxury SUV	-	**0.016**	Land Rover Discovery	Lexus GX470
		[1.854]		
CC x Large SUV	-	**0.017**	Chevrolet Tahoe	Chevy Suburban
		[1.953]		
CC x Middle SUV	-	**0.015**	Land Rover Freelander	Nissan Xterra

		[1.930]		
CC x Small SUV	-	-0.004 [-0.531]	Chevrolet Tracker	Jeep Wrangler
<u>Elasticities</u>				
Min	-5.0	-7.3		
Mean	-3.9	-6.0		
Max	-2.5	-3.7		

CC denotes consumer confidence. Coefficients in bold denote significance at 5% level. T-stats in brackets. All specifications include time and brand fixed effects. Specification I also includes segment fixed effects. When testing for overidentifying restrictions the tests statistics are 1.615 and 1.125 for Specifications I and II, respectively. The associated p-values are 0.656 and 0.771. The degrees of freedom in both cases is three.

Table 5. Univariate processes for exchange rates and consumer confidence, Jan 1973-July 2006 bimonthly data.

	usd/eur	Consumer confidence	usd/jpy
Estimation	GARCH(1,1)	E-GARCH(1,1)	GARCH(1,1)
C	0.0003	**1.5170**	0.0004
	[0.09]	[2.70]	[0.11]
κ	0.0004	**2.0786**	0.0011
	[0.80]	[2.84]	[0.50]
α_1 "GARCH"	**0.7392**	**0.5033**	0.5263
	[2.40]	[2.94]	[0.62]
α_2 "ARCH"	0.0960	**0.4531**	0.0549
	[1.09]	[2.54]	[0.68]
γ "Leverage"		**-0.3759**	
		[-3.26]	
Degrees of freedom	200	18.17	16.37
Log-likelihood	324.1647	-712.1135	309.1884

Regressions run on bimonthly data 1973:1 to 1996:6. T-stats in brackets. Coefficients in bold are significant at the 5% level.

Table 6. Hedonic regression elasticity estimates, US car market 1995-2006.

Dependent variable: Real Prices (USD) Variables	[1]	[2]	[3]
prod. EUR x USD/EUR	0.0812*** [2.64]	0.087*** [2.99]	0.098*** [3.39]
prod. JPY x USD/JPY	0.113*** [3.4]	0.123*** [3.88]	0.096*** [3.23]
HP		0.063*** [4.15]	0.154*** [6.20]
Size		-0.0562** [-2.32]	0.005 [0.20]
Transmission		0.000722 [1.25]	0.001* [1.86]
csihat			0.00213*** [6.52]
Model Fixed-effects	Yes	Yes	Yes
N	1112	1112	1112
R-squared	0.9910	0.9913	0.9938

Note: The table reports elasticities and associated t-statistics for the regression of real prices on the interaction of product location and (real) exchange rates, controlling for product characteristics and product fixed-effects. Specification 3 also controls for product unobserved characteristics which are observed by market participants, but not the econometrician. * p<0.10, ** p<0.05, ***p<0.01