

**The Impact of Component Sharing on Quality: An Empirical Study in the
Automotive Industry**

Kamalini Ramdas

Darden Graduate School of Business

University of Virginia

Charlottesville VA 22906

Email: ramdask@darden.virginia.edu

Taylor Randall

David Eccles School of Business

University of Utah

Salt Lake City Utah 84112

Email: taylor.randall@business.utah.edu

June 8th 2004

The Impact of Component Sharing on Quality: An Empirical Study in the Automotive Industry

Kamalini Ramdas

Darden Graduate School of Business

University of Virginia

Taylor Randall

David Eccles School of Business

University of Utah

Abstract

Component sharing – the use of a component on multiple products within a firm’s product line – is widely practiced as a means to offer high variety at low cost. While many researchers have examined tradeoffs involved in component sharing, little research has focused on the impact of component sharing on quality – defined in this paper as component reliability. The design literature suggests that a component designed uniquely for a product will result in higher quality due to the better fit of the component within the architecture of the product. The traditional learning curve literature suggests that higher volume of a component will increase component reliability. Sharing a component across multiple products increases volume and hence should increase reliability. However due to difficulties in transferring learning across products, we anticipate that this volume effect will be moderated by the number of products needed to achieve any given volume. Using data from the automotive industry, we find strong support for the hypothesis that higher component reliability is associated with a component that has been designed uniquely for a product. This finding suggests that the popular design strategy of component sharing can in some cases compromise product quality. We also find support for the hypothesis that higher component reliability is associated with higher cumulative component volumes, and we find that this effect is moderated by the number of models needed to achieve the cumulative volume. These effects present designers with a conundrum: both designing a unique component for each product application and sharing a component across multiple products can increase reliability. We evaluate the relative strength of each of these effects within the context of our data and develop insight on the situations in which each effect dominates the other.

1. Introduction

In this paper, we examine the impact of component sharing – the use of a component across multiple end products within a firm’s product line – on quality, focusing specifically on component reliability as opposed to other interpretations of quality¹.

Consider the following design scenario relating a firm’s component sharing strategy to the reliability of the components used in its products. When designing an assembled product, designers repeatedly evaluate whether to create a unique component specifically for the product or to reuse an existing component. Specific design allows the designer greater flexibility in tailoring the component specifications to the needs of the product, which should lead to higher quality. For any component, however, unanticipated defects may arise in its manufacture, assembly or use. Over time, the occurrence of such defects is reduced via improvements in manufacturing and assembly processes for the component, and engineering improvements to the component itself based on feedback from downstream firm functions and from end users. The benefit of using an existing component is that many reliability problems may already be identified and corrected via this improvement process. The downside to using an existing component is that design fit can be compromised, leading to decreased reliability. The tension described in this scenario leads us to the central questions of this paper. Does greater reliability result when a component is used in a product for which it was specifically designed? Is higher component reliability associated with increased experience with a particular component, due to the greater manufacturing, assembly and field performance experience gained by increasing component volume? Finally, if both the *specific design* effect and the *experience* effect exist, which might have the greatest relative impact and in what situations?

We believe that these are important questions both for practitioners that engage in component sharing and for researchers studying this topic for several reasons. First, if component sharing compromises product reliability, the detrimental quality effects could negate the widely touted benefits of component sharing which include reduction in product design, manufacturing and distribution costs, and increased responsiveness to

¹ We will use the terms quality and reliability interchangeably in this paper.

consumer demand (Ulrich 1995; Swaminathan and Tayur 1998; Ramdas and Sawhney 2001; Ramdas 2003, Rutenberg 1969; Fisher et al. 1999; Krishnan and Gupta 1998, 2001; Desai et al. 2001; Kim and Chhajed 2001; Ramdas et al. 2003; Thonemann and Brandeau 1999). Conversely, if component sharing enhances product reliability in addition to its more widely known benefits, this could lead companies to encourage more component sharing. Second, if both effects do exist future research may focus attention on “smart” component sharing strategies where designers may have the best of both worlds.

While a few researchers have prescriptively modeled how component sharing can be effectively used to satisfy market segments with differing quality needs (Desai et al. 2001; Heese and Swaminathan 2004), the reliability tradeoff we described above remains unaddressed. Using empirical data in the domain of automotive braking systems, we examine how component sharing impacts reliability. Examining the impact on reliability of tailoring a component’s design specifically to a product application is one major contribution of our work. Another key contribution is that we shed light on increased component volume attained via component sharing improves reliability. Finally, we are able to highlight design contexts where each of these effects dominates the other.

In section 2 we describe the industry context for our study. In section 3, we develop our hypotheses. In section 4, we discuss the data and variable definitions. In section 5 we present methodology and results. In section 6 we discuss our findings and their implications for practice. Section 7 contains conclusions and limitations.

2. Industry Context

The context for our study is the automotive industry. We focus specifically on one component of the automotive braking system – the brake rotor – and study brake rotor sharing strategies at Ford Motor Company. We chose to focus on braking systems because braking quality is of critical importance to the consumer, brakes-related quality issues are a critical determinant of warranty costs (Automotive News, June 1999), and there is considerable sharing of braking system components at Ford and other auto makers. Within braking systems, we chose to focus on brake rotors because based on our discussions with industry experts, reliability problems associated with brake rotors are

most easily identified as occurring due to rotor design decisions, as opposed to decisions about some other components in the braking system or the automobile.

An automotive braking system is a hydraulic system that converts human foot pressure applied at the pedal into a much higher braking pressure applied at the wheels, via the braking system components. The pressure applied at the wheels forces stationary brake components – brake pads which are attached to the calipers – to rub against rotating components – the brake rotor, thus converting the kinetic energy of a moving car into heat energy via friction. The pressure applied to the rotor by the brake pads causes wear on the brake rotor. Common rotor problems that arise from use include warping, scoring or even cracking of the rotor.

Automotive braking system design is initiated only after vehicle design has been broadly specified, via “system level parameters” such as vehicle weight, top speed, and stopping distance. Given these inputs, the components of the braking system must be designed so as to provide adequate torque to stop a car from top speed within the desired stopping distance. In addition, all braking components are designed for "maximum loading" conditions: for example, the brake pedal should not crack if the driver steps exceptionally hard on it in a panic stop. Further, several constraints arise due to the interaction between braking components: for example, the hydraulic ratio (ratio of areas of master cylinder and caliper pistons) must lie within pre-specified limits to eliminate excessive pedal “travel”, which could cause the brake pedal to hit the floor of the car. Braking system design parameters like rotor radius, desired pedal force, and area of the caliper pistons and master cylinder piston are manipulated to meet these different ends. In the context of our study, the challenge for the designer is to balance the constraints of each unique braking system with the potential benefits of component sharing.

3. Theory and Development of Hypotheses

Our hypotheses rely on prior research in product design theory and learning curve theory, as well as interviews with senior executives, managers and engineers in the auto industry. The first hypothesis addresses the question of whether uniquely designed components produce better quality outcomes. The second hypothesis addresses the question of whether and how quality improvement through experience effects occurs as the volume of a component increases.

Design theory suggests that designing unique components for each specific product application will result in higher product quality. Ulrich (1995) discussed the role of components within modular and integral product architectures. In a pure modular architecture, the interfaces among components are standardized and so multiple products can be configured by mixing and matching from a base set of components. In a pure integral product architecture, the complex interactions among components requires that components be specifically designed for each product (Ulrich and Ellison 1999). In practice, products are often a blend of modular and integral architectures. The elements of the architecture that are integral lead designers to uniquely design components for each end product rather than share components across end products. An implication of this theory is that sharing components inappropriately will result in poor fit of the component with the product and hence will hurt quality. Engineers at Daimler Chrysler provide anecdotal evidence for this proposition pointing out that due to architectural constraints it is impossible to adequately adapt components from rear-drive Mercedes Benz cars to front-drive Chrysler vehicles (Automotive News, Nov 2001).

Distinct from the fit argument, we argue that an engineer's accountability for design quality is higher when he or she is charged with designing a component for a specific model, than when the engineer reuses an existing component. In the latter case, a quality problem could potentially be attributed to those who originally designed the component. This reduced accountability somewhat reduces the engineer's incentive to ensure that the reused component will work well for the application at hand. These arguments lead us to believe that the reliability of a product is higher if the component is designed specifically for the product.

This basic notion is complicated by industrial practice. Auto manufacturers often choose to design a new rotor for use on multiple models at once, rather than on a single new model. We argue that doing this dilutes the specificity of the design to any particular model in the use set. Therefore for any vehicle, if the rotor used on it was designed for the specific model and model year of that vehicle, and furthermore if that was the *only* vehicle for which it was designed, we expect the fit and therefore the reliability to be higher than if that rotor was designed for use on several different models in that particular model year. In our empirical model, we expect that the positive impact on reliability associated with using a rotor that was designed for the model and model year of a vehicle will be moderated by the total number of models on which that rotor was used in its first year². These arguments lead to our first hypothesis.

Hypothesis 1: The greater the specificity of a component’s design to a product application, the greater the reliability of the component in that application. However, the larger the set of products a component was designed for, the poorer the specificity of the design to any one product in that set, and hence the poorer the component’s reliability.

In the automotive braking context, we expect that the reliability of a brake rotor in a vehicle will be higher if the rotor was designed for the specific model and model year of the vehicle, than if it had been originally designed for another model or even another model year of the same model. This effect will be negatively moderated by the number of models for which the brake rotor was designed.

The literature on learning curves provides empirical evidence that product performance – often defined in terms of cost reduction – increases in the cumulative production volume of a product.³ Fine (1986) and Ittner (1996) link learning curve theory to quality improvement. Anzoni and Simon (1979) describe the cognitive processes involved in learning curve theory as “learning-by-doing.” Adler and Clark

² We assume that if a rotor was used on multiple models in its first year, it was designed for use on all of these models.

³ See Hatch and Mowery (1998) for a comprehensive list of the many studies that document the learning curve phenomena in various industries. See as examples Wright 1936 – airplanes, Baloff 1971 - automobiles, Dick 1991 – semi-conductors, Lieberman 1984 – semi-conductors.

(1991) and Hatch and Mowery (1998) detail two important nuances of “learning-by-doing”. Adler and Clark report that cumulative product volume is highly correlated with engineering change orders and training of personnel both performed with the explicit goal of improving product performance. Thus product performance increases not simply due to increasing volumes but rather due to engaging consciously in improvement activities to address the problems that the increases in volume reveal. Hatch and Mowery conclude that knowledge associated with learning-by-doing is somewhat specific to the production environment. They report that losses in semi-conductor yield occur as processes are transferred from development to manufacturing. Similarly Adler (1990) shows that not all productivity gains associated with learning-by-doing are transferred to new manufacturing situations.

Now reconsider the product design decision presented in the introduction, in the context of brake rotors. When designing a car, a designer is faced with the decision of whether to design a new brake rotor for the braking system of the car (and potentially for other models that will share the same rotor) or to use an existing brake rotor. The potential quality benefit of using an existing brake rotor is increased rotor reliability due to cumulative production experience with the rotor. This experience may have already accrued if a brake rotor has reached sufficiently high volume or may be accelerated by adding the additional volume of the current design project. Industry executives we interviewed claimed that the economic incentive to share components increases when after-sales warranty costs are factored into the design decision. This is consistent with the analysis of Ramdas et al. (2003) who show via a prescriptive optimization model that if warranty costs are considered when deciding whether or not to design afresh, the likelihood of reusing an existing component increases.

Similar to productivity losses that occur as manufacturing processes are transferred to different situations, the transfer of a component to a new model may also result in some loss of quality. We believe that learning occurs through a process of identifying problems, solving them, and transferring the knowledge gained in this process so that the same problems do not recur. For example, when a new car model is launched at an assembly plant, the quality assurance department may notice an incidence of problems related to seat assembly. A quality circle team may find that a modification of the

assembly process fixes the problem. If this new procedure is made standard, and this information is disseminated in the plant, then any worker who works at the seat station will follow the new procedure, resulting in improved seat assembly quality.

Due to the stickiness of information and difficulty in transferring knowledge across separate assembly locations, we expect that there will be less transfer of knowledge when the vehicles that use a particular component are assembled at different locations. Also, idiosyncratic characteristics of the assembly location for the new model relative to an existing assembly plant where the component was used may reduce the relevance of previous learning and also introduce new causes for defects. Exhibits 1 and 2` show examples in our data from Ford of how specific brake rotors can reach high production volumes via different component sharing strategies – sharing over time on a single model, or sharing over multiple models, over time. In practice, different models that share a component are often made at different plants or on different assembly lines, and this trend will increase as companies move towards implementing global product platforms⁴. General Motors is using its Epsilon platform as the basis for models from the Oldsmobile, Chevrolet, Pontiac, Saab and Opel brands that will be made in plants located in different continents (Automotive News, Nov. 1999). We expect a brake rotor that reaches a high production volume with use on one model to have higher reliability than a brake rotor that reaches a high production volume with use on multiple models. These observations lead us to our second hypothesis.

Hypothesis 2: Higher component reliability will be associated with higher cumulative production volume, but the relationship will be moderated by the number of products used to attain any particular volume.

In the automotive braking context, brake rotor reliability will increase in the cumulative volume of the rotor, but this increase will be mitigated by the number of models used to attain that volume.

⁴ A platform is defined most broadly as a set of resources that is shared across products – ranging from components to production processes.

4. Data and Variables

We obtained data from several sources. The data on product quality focuses specifically on the reliability of automotive brake rotors. This data was obtained from the Department of Transportation of the US government, and records consumer complaints about front brake rotor quality problems. This data is used by the US government to guide product recall decisions. Our data on brakes component sharing was obtained from an automotive research company. From this data, we can determine all of the unique braking components in use on vehicles sold on the US market in the period of our study, and what vehicles used each unique brake. We calculated cumulative volumes for each brake rotor by summing the volumes of all the models over which the rotor had been used, up to and including the model year of the car. We obtained individual model volumes from Wards Automotive and Automotive News, and adjusted for cases in which multiple rotors had been used on a model. We obtained data on vehicle characteristics such as weight and horsepower from Wards Automotive and Automotive News, and data on brake characteristics via direct measurement, Internet sources, and Motor Vehicle Manufacturer's Association specification sheets. We obtained data on control variables such as precipitation and population density from public sources.

The unit of analysis in our study is an individual vehicle on the road, identified by its vehicle identification number (VIN). The VIN label, which is often printed on the underside of the dashboard, uniquely identifies every vehicle on the road. It can be used to track ownership, standard and optional factory installed equipment, and other individual-vehicle specific details. The original data set from the government on brake failures contained 893 observations from Ford Motor Company for individuals reporting a first time rotor failure. Of these, 295 observations were dropped due to incomplete data. The primary reason for the incomplete data was lack of precise rotor information to match the VIN number. The pattern of missing rotor information appears random, reducing the potential for sample bias. This leaves 601 observations on which to test our hypotheses. Approximately 56 different brake rotors were used on the vehicles in our data set.

Below, we describe the variables used in our hypothesis tests. Refer to Table 1 for descriptive statistics.

Dependent Variable:

Miles to Failure: We use miles to failure of the brake rotor as our measure of rotor reliability, the dependent variable in our study. Miles to failure is calculated as the miles driven in a particular vehicle from 0 miles to the miles at the time of the first reported failure of the brake rotor. The rotor failure may be due to problems such as scoring, warping or cracking. The mean miles to failure is 33,579 (median = 29,000) with the minimum failure reported at 2 miles and the maximum failure reported at 160,000 miles. Exhibit 3 presents a histogram of miles to failure. The data is clearly skewed.

Study Variables:

Cumulative Volume: Existing literature uses cumulative production volume as a proxy for learning effects (Macher 2003). Although alternative measures such as time and engineering resources have been explored (Hatch and Mowery 1998), cumulative volume is the standard proxy. We measure the cumulative production volume of a rotor used in a VIN as the total volume from its time of first use up to the model year of that VIN. The mean cumulative production volume for brake rotors is 880,178 rotors (median = 544,816) with a minimum production volume of 8,334 and a maximum of 7,568,420. Exhibit 4 shows the cumulative volumes over time for the 70 unique brake rotors used by Ford over the time period of our study. It is common convention in the learning curve literature to take the natural log of cumulative volume for estimation procedures (Hatch and Mowery 1998). We follow this convention.

Number of Models: The number of models represents the number of different vehicle models on which a particular brake was shared up to the time of manufacture of the vehicle in question. This variable serves as the moderating variable in our hypotheses. Rotors are shared on an average of 2.36 models (median = 2) with a minimum of 1 model (no sharing) and a maximum of 9 models. We defined each unique name plate – e.g. Ford Taurus vs. Mercury Sable – as a distinct model. Although these vehicles share a platform, there are some significant differences – for example in vehicle weight, which is an important factor in brakes design. This is the reason for our choice of definition.

Measures of product specificity of component's design: We measure whether a brake was designed specifically for a model two different ways. An indicator variable Rotor-Year-Model captures whether or not the brake rotor used in a VIN was designed specifically for the model and model year of that VIN. If the rotor used on a VIN was introduced prior to the model year corresponding to that VIN, then Rotor-Year-Model equals 0. If instead the rotor used on a VIN was introduced in the same model year as that VIN, Rotor-Year-Model equals 1 regardless of how many other models the rotor was used on in its first year of use. In our sample, 36% of the observations had rotors designed specifically for the model and year of the VIN in question. Another indicator, Rotor-Model, is set equal to 1 if the brake used in a VIN is designed specifically for the model corresponding to that VIN, regardless of year. This variable accounts for rotors shared on the same model across time. For example, if a rotor is used on a model in 1993 (the design year), 1994 and 1995, Rotor-Model equals 1 in each of those years. 85% of our observations had a rotor designed for the model, but not necessarily for the model year of the VIN in question.

Control Variables:

Vented: Rotors may be vented or solid. Vents alleviate heat generated as calipers come in contact with the rotor during braking. The reduced heat keeps rotors from warping easily. We use an indicator variable equal to 1 to represent a vented rotor and a zero otherwise. We expect vented rotors to have lower hazard rates than solid rotors, other things being equal. 87% of the observations in our sample have vented rotors.

Horsepower: Automobiles with higher horsepower can accelerate faster, placing higher requirements on the brakes. We expect automobiles with higher horsepower to have a higher failure rate. The mean horsepower was 150 HP (median = 140 HP).

Swept Area per Ton: The swept area is the area of contact between a brake caliper and brake rotor. All else being equal, the larger this swept area the greater the braking ability. Heavier cars typically have a larger rotor swept area. The rotor swept area per ton controls for the relation between the weight of the vehicle and the compensating swept area, and a common metric of braking potential. We expect a higher

swept area per ton to result in higher miles to failure. We report a mean swept area per ton of 120 square inches per ton (median = 120 square inches per ton).

Percent Weight Difference: Interviews with industry personnel suggest that when a brake rotor is reused across multiple models, to ensure reliability the rotor selected might be over specified. In this case, we would expect the rotor performance to increase when reused. While there are many dimensions of performance specification, our interviews indicated that it is easier to share a rotor across models if the new automobile weighs less than the weight of the car that the rotor was originally designed for (design weight). We use the positive percent weight difference, calculated below, as a proxy for the potential over-specification of shared brake rotors. For each VIN,

$$\text{Percent Weight difference} = \begin{cases} \frac{(\text{design weight} - \text{weight of VIN})}{\text{design weight}}, & \text{if design weight} > \text{weight of VIN} \\ 0, & \text{else} \end{cases}$$

We report an average positive percent weight difference of 6.12% (median = 4.78%). Note that 78% of our observations exhibit a positive difference, indicating that rotors are typically reused on models that weigh less than their design weight.

Finally, we use indicator variables to control for the specific assembly plant that produced a VIN. MacDuffie et al. (1996) reported significant differences in quality across automobile assembly plants. While we found significant coefficients for several of the assembly plant dummies, for brevity we omit these variables in reporting our results.

Population Density: We use the population density of the county of driver residence as a proxy for general driving conditions. We expect the hazard rate to increase with population density as city driving results in greater use of the brakes. We report a mean population density of 3,442 people per square mile (median = 2,636). Population density exhibits considerable skewness. We mitigate this problem by taking the natural log of population density.

Precipitation: Interviews with brake experts suggested that weather conditions associated with precipitation can reduce the longevity of brakes and brake rotors. We use the average precipitation of the county of driver residence to control for wet weather conditions. The mean precipitation per year is 38 inches (median = 41 inches).

Table 2 presents Pearson correlation coefficients among the study variables. We highlight several of the significant correlations. We note that our dependent variable, miles to failure, is negatively associated with horsepower and positively associated with precipitation, but is not significantly correlated with our variables of interest. Second, we notice that cumulative production volume is highly correlated with the number of models.

5. Methods and Results

Since we are dealing with survival data, a hazard rate model should be used rather than standard regression analysis (Helsen and Schmittlein 1993). To test our hypotheses, we employ a Cox proportional hazards model (Cox 1972). The Cox model is commonly used for hazard rate analysis, and accommodates skewed distributions of failure time. The hazard of failure can vary over the life of the item being studied – e.g. older men are more likely to get prostate cancer, and an older brake is more likely to fail. An advantage of the Cox model over other hazard models is that it requires no strict assumption about how the hazard of failure varies over time. In our study, we are interested in understanding what drives the differences in failure rates across different models at any mileage, rather than how the hazard rate varies with mileage. Thus the Cox model is the preferred method of estimation (Allison 1995). We specify the hazard function used to test our hypotheses as follows:

$$h_i(t / \mathbf{X}) = \lambda_0(t) \cdot e^{\alpha_1 \ln(\text{volume}_i) + \alpha_2 (\text{number of models}_i) + \alpha_3 \ln(\text{volume}_i)(\text{number of models}_i) + \alpha_4 (\text{specific design}_i) + \alpha_5 (\text{specific design}_i)(\text{number of models}_i) + \beta (\text{control variables}_i)}$$

where $h_i(t / \mathbf{X})$ represents the hazard of failure for the i^{th} observation (VIN in our study) at mileage t , given a set of covariates \mathbf{X} , $\lambda_0(t)$ represents a base line hazard that is a function of miles driven but does not vary by individual VIN, the α 's are the coefficients of the study variables, "control variables" represents a vector of control variables and β represents a vector of coefficients of the control variables. The variable "specific design" refers to either Rotor-Year-Model or Rotor-Model, discussed earlier. Since we estimate this model using Cox regression, we do not need to specify the form of $\lambda_0(t)$. The Cox model is called a proportional hazards model because the ratio of hazards for any two individuals is constant over time. In our application, it is reasonable to assume that all

brakes share a similar pattern of hazard $\lambda_0(t)$ with regard to miles driven, and that relative differences in reliability at any mileage are a function of differences in the independent and control variables included in our study.

Since we have data only on those individuals who experienced a rotor failure and who also chose to report this failure, the hazard distribution for our data could be different from that in the general population. However, under the reasonable assumption that the probability of reporting a failure is not a function of any of the variables of interest in our study, it can be shown that the coefficients of our study variables have the same interpretation as they would had we in fact analyzed data representing both reported and unreported failures (see Appendix 1).

Table 3 presents results of estimates obtained using the Cox proportional hazard model. Note that higher reliability is associated with a lower hazard (negative coefficient) and lower reliability is associated with a higher hazard (positive coefficient). The Chi-square statistics for all models in the table are significant at the 0.01 level. The r-square values calculated based on Allison (1995) range from 17.60% to 20.69%. The first column presents results of a model with control variables only. As expected, higher proportional hazard rates are associated with higher horsepower and higher population density. We see no significant association between vented rotors, swept area per ton, precipitation, percent weight difference and the hazard rate.

Columns II and III present models containing the main effect for volume, our two alternative fit measures (Rotor-Year-Model: designed specifically for model and year; Rotor-Model: designed for model) and interaction terms which test our two hypotheses. In Column II, we report a negative and significant coefficient for Rotor-Year-Model and a positive and significant coefficient for the interaction between Rotor-Year-Model and number of models. This finding is consistent with hypothesis 1, which states that the reliability of the brake will be higher on cars that are among the set of models the brake was designed for, but that this effect will decrease as the number of models the brake was designed for increases. We also report a negative and significant coefficient for the log of cumulative volume and a positive significant coefficient for the interaction between the log of cumulative volume and the number of models. This finding is consistent with

hypothesis 2 which states that reliability is increasing in cumulative volume, but that this effect mitigated by the number of models over which the volume is shared.

Schoonhoven (1981) suggests analyzing the derivative of the estimated equation for more careful examination of interaction effects when the main effect and interaction effect have opposite signs. For ease of exposition, we ignore the subscript i that refers to specific VIN numbers, in our regression model. Our regression model is:

$$h(t/X) = \lambda_0(t).e^{\alpha_1 \ln(\text{volume}) + \alpha_2 (\text{number of models}) + \alpha_3 \ln(\text{volume})(\text{number of models}) + \alpha_4 (\text{specific design}) + \alpha_5 (\text{specific design})(\text{number of models}) + \beta (\text{control variables})}$$

Taking the derivative of the hazard rate h with respect to volume, and setting this to zero, we have:

$$\frac{\partial h(t/X)}{\partial \text{Volume}} = h(t/X) \left(\frac{\alpha_1}{\text{Volume}} + \frac{\alpha_3 (\text{number of models})}{\text{Volume}} \right) = 0$$

Solving for the number of models at which $\frac{\partial h(t/X)}{\partial \text{Volume}} = 0$,

$$\text{number of models} = -\frac{\alpha_1}{\alpha_3}$$

For our data set, with $\alpha_1 = -0.355$ and $\alpha_3 = 0.139$, the proportional hazard rate decreases with volume when the number of models over which the volume is spread is less than 2.55, and starts to increase with volume when the number of models is greater than 2.55. In our data, roughly 25% of all observations have brake rotors that are shared across more than 2.55 models. By similar logic, we analyze the interaction term between Rotor-Year-Model and number of models. Here, the benefits of designing for a specific set of models are eliminated when the number of models equals 2.67. In our data, roughly 20% of all observations have brake rotors that were designed for more than 2.67 models.

Column III reports results of tests using Rotor-Model as the measure of design specificity. In this column, we report no significant effects for Rotor-Model and for the interaction between Rotor-Model and number of models. Note that while the same rotor may be reused over multiple years, the automobile as a whole can undergo significant design changes, potentially damaging or enhancing the fit of the rotor with the braking system and the automobile – e.g. our data reveals that the weight of a particular model can vary considerably over time. Under these conditions, the Rotor-Year-Model variable provides a more suitable measure for design specificity. We also note that the main

effect of volume and the interaction between the natural log of volume and number of models fall below conventional significance cutoffs.

We checked outlier diagnostics and found that the results in table 3 are robust to the influence of single data points. Further, we noted no problems with multi-collinearity, except for the interaction terms, for which the variance inflation factor is quite high.⁵

We re-estimated our model using different subsets of the data to test the robustness of our results. These model runs are reported in Table 4. Note that we only report results for models using the Rotor-Year-Model variable. Regressions using the Rotor-Model variable are not reported in Table 4, but differ from those in Table 3 – Column III in one significant way. The results consistently support hypothesis 2 that volume can reduce the rate of hazard, but does so at a decreasing rate.

Our first concern deals with a potential endogenous association between quality and the number of models on which a brake rotor is used. Note we have no direct hypothesis about the number of models and quality. However, if designers are given feedback on the quality of a brake, a designer may choose to reuse a brake because of its observed quality as opposed to the number of models a new brake is designed for affecting quality by altering design specificity. Our industry interviews yielded mixed opinions as to the existence and strength of this feedback loop. Also, reliability data only becomes available a few years after a brake's introduction. Nevertheless, to alleviate this potential concern, we eliminate instances where quality feedback may have led to a reuse decision by eliminating all observations where a brake was not designed for the specific model or models.⁶ Column I of Table 4 reports the results of this analysis. We observe that the coefficient on the number of models and the interaction term between number of models and $\ln(\text{Volume})$ become insignificant. Hypothesis 2 holds for the main effect of volume, but not the interaction.⁷ A second concern raised by discussions with industry

⁵ We observe that in models that include the main effects of cumulative volume and number of models, but exclude the interaction term, the signs of the main effects do not change from what is reported in columns II and III.

⁶ By placing this constraint on the data, we not only reduce the opportunity for endogenous effects, but also reduce the variance in the number of models. Importantly, this reduction in the number of models may occur in instances where the detrimental effects of sharing may be most likely.

⁷ We also re-estimate the models shown in columns II to VII using the observations for which there is no reuse. We report the following differences in the results from column I. For early failure, the interaction term for $\ln(\text{volume})$ and Number of models is positive and significant. For Passenger Cars, the interaction term for $\ln(\text{volume})$ and Number of Models is positive and significant, while the interaction term for

experts deals with a potential “lemon” effect, meaning that a certain number of parts fail at the outset of an automobile’s introduction. These failures may not have any association with the factors in this study. To examine the robustness of our results to “lemon” related failures, we eliminated observations with failures that had occurred at less than 500, 1,000 and 2,000 miles. The results of the 2,000 mile cutoff are qualitatively similar to the 500 and 1,000 mile results and are shown in Column II of table 4. Third, a limitation of the data is that cumulative production volumes are for North American auto sales only. Thus volumes for brakes shared globally will be understated creating a potential bias in the association between volume and the hazard rate. We were unable to acquire precise volume estimates for brakes shared globally. However, through interviews with company representatives and automotive industry experts we were able to identify several models where this effect was likely to be most severe, and we reran the analysis eliminating these models. These results are shown in Column III of Table 4. Fourth, our data contain both passenger cars and trucks. It is possible that the magnitude of these factors differs greatly across the types of vehicles. Column IV of table 4 shows results for passenger cars and Column V shows results for trucks. Fifth, we measure cumulative production volume at the end of each production year. An alternative is to measure cumulative production volume at mid year levels. Column VI shows results using this alternative measure of production volume. Finally, we observe that for some brakes we have very few observations. For such brakes we may lack sufficient data to make proper statistical inference. Column VII shows results for the subset of brakes for which we have more than 4 observations. We note in Table 4 that the results of regressions estimated using subsets of our data consistently show support for hypotheses 1 and 2.

6. Discussion

We find support for the two questions we address in our research. Rotor reliability is higher when the model using a rotor is among the set of models that the rotor was designed for, than if this is not the case. However, this specific-design effect is dissipated as the number of models that the rotor was designed for increases. Further,

Rotor-Year-Model is insignificant. For both trucks and for rotors with 4 or more observations, the only significant result is Rotor-Year-Model. Reduction in significance may be due to lack sufficient data for proper statistical inference.

higher production volumes on brakes are associated with higher reliability. However, the positive impact of cumulative component volume on the reliability of a brake rotor is moderated by the number of models on which the rotor is used. The inherent tension in these findings merits further analysis and interpretation. To facilitate this process, we compare the ratios of relative hazards between different design scenarios. The use of ratio comparison with a Cox proportional hazards model isolates the effects of variables of interest.

Scenario 1 - Reuse of an existing rotor on a continuing model: Consider a scenario where an automobile model (e.g. the Ford Escort) is being redesigned. To maximize reliability, should the rotor be carried over from the previous year's Escort model or should a new rotor be designed for the new model version? Using the estimates in Table 3 – column II and a general hazard function, we calculate the ratio of hazard between a uniquely designed brake rotor and a brake rotor that was used in the previous model year. Let CV denote the cumulative volume of the existing brake rotor, up to and including the volume from the previous year's model. We assume in this comparison that if a new rotor is designed, the designer is interested in the hazard rate of this rotor on the first car of the new model version. Since the hazard rate for a newly designed rotor decreases as the volume of the new model increases, this assumption gives us the worst case reliability for the new rotor. After simplification, the resulting ratio is as follows.⁸

$$\frac{\text{hazard of newly designed rotor}}{\text{hazard of reusing a rotor with a given cumulative volume}} = \frac{e^{-0.842}}{e^{-0.216 * \ln(CV)}}$$

Intuitively, we trade off the gain from product-specific design with the reliability gains from cumulative production volume. The case for new design will dominate the case for reusing the old brake rotor when the above ratio is less than 1, i.e. when the reduction in hazard achieved via using a newly designed rotor exceeds the reduction in hazard achieved via cumulative production volume of the existing rotor. This occurs when

⁸ Note that the coefficient for both Rotor-Year-Model and Ln(CV) are reduced by the value of the coefficients of their respective interaction terms. This is because the number of models in this scenario is equal to 1.

$-0.842 < 0.216 * \ln(CV)$, i.e. when $CV < 49$. This cumulative volume is dramatically less than the minimum cumulative volume for any brake rotor in our data set. In this scenario, we conclude that the reliability effect associated with unique design is dominated by the reliability effect of cumulative volume. Of course, as the volume of the new model version increases, reliability of the newly designed rotor would increase, and at some point the reliability of the new design will exceed that of the existing design. We investigate this effect further below. Note that the new design may be chosen even if its reliability is initially lower than that of the existing design, so long as its reliability exceeds the target reliability specification that the company has adopted.

Scenario 2 – Reuse an existing rotor on a new model: Consider a scenario where a designer may create a new rotor for a new model or choose to reuse an existing rotor. Let V denote the expected volume on the new brake, CV the expected cumulative volume of the reused brake and N the number of models on which the rotor is already in use. After simplification the resulting ratio is as follows:

$$\frac{\text{hazard of newly designed rotor}}{\text{hazard of reusing a rotor with a given cumulative volume}} = \frac{e^{-1.344 + 0.502(1) - 0.355 * \ln(V) - 2.046(1) + 0.139 * \ln(V)}}{e^{-0.355 * \ln(V + CV) - 2.046 * (N+1) + 0.139 * (N+1) * \ln(V + CV)}}$$

As with the previous example, the reliability of the new design will dominate reuse when the ratio is less than 1. Note that the ratio depends on V , CV and N . Exhibit 5 illustrates how the ratio changes with V and CV . For purposes of discussion the calculations in exhibit 5 assume that N is equal to 2, which is the median number of models sharing a brake rotor in our data set. The x-axis shows representative values for anticipated volume V of a new brake rotor that is used on a single model. The lines represent the minimum, first quartile, median, third quartile and maximum values of cumulative volume for rotors shared across two models in our data set. The exhibit illustrates that as the anticipated volume of the new brake rotor increases, the effect of specific design tends to dominate the learning effects of cumulative volume. This effect occurs because the learning effects obtained from the new rotor, combined with the benefit for specific design, more than compensate for the learning effects lost when choosing not to reuse the brake rotor.

Scenario 3 – Design for multiple models: Consider a scenario where a designer needs to design brake rotors for a series of models in some model year. The designer can capture the benefits of design specificity by creating a new brake rotor for each model or can capture the potential learning effects of production volume by using a universal design that serves all models. In this scenario, we assume that the anticipated volume of each individual model equals V. The resulting ratio is as follows:

$$\frac{\text{hazard of newly designed rotor}}{\text{hazard of universal design}} = \frac{e^{-1.344+0.502(1)-0.355*\ln(V)-2.046(1)+0.139*\ln(V)}}{e^{-1.344+0.502(N)-0.355*\ln(N*V)-2.046*(N)+0.139*(N)*\ln(N*V)}}$$

Exhibit 6 shows values of this ratio for different values of V and N. Note that the newly designed rotor dominates the universal design as the anticipated volume of each brake increases. The newly designed rotor also dominates as the number of models on which the brake rotor needs to be shared increases. Universal design appears to be most useful when sharing across two or three models and with low anticipated volumes.

The scenario analysis above illustrates different decision making contexts for component sharing, and identifies the conditions where the unique design effects dominate and the conditions where the learning effects dominate, within our data. The scenarios were chosen because they are common decisions for a designer, not because they represent a comprehensive list of design scenarios. Other scenarios are certainly possible, and our approach provides a way to evaluate these scenarios.

7. Limitations and Conclusions

The data are limited to one automotive company and are the product of the design processes at that company. Generalization of the results will depend, in part, on the similarities among design processes across companies. Further, our analysis focuses on a single component, the brake rotor. Despite this, we believe our methods and results may be generalized in the following ways. First, brake rotors are mechanically similar to many other components in an automobile and components in other products. Second, the methods and measures used to estimate the impact of component sharing on failure rate of brake rotors are quite general and are applicable to components in other assembled products.

This study provides the first empirical evidence relating component quality – defined in our case as component reliability, to component sharing strategies. We find support for the hypothesis that higher reliability is associated with a component designed specifically for a model in a given year, but that this effect is reduced when the component is designed for simultaneous use on other models. This finding suggests that the popular design strategy of developing multiple products off a common platform with shared components can in some cases compromise product quality. We also observe that improved reliability is associated with higher cumulative production volumes, but these benefits decrease as the number of models required to attain this volume increases. Comparison of these effects suggests that the dominance of one effect over another will depend on the sharing situation being considered. These findings suggest that designers and researchers should look beyond the inventory and cost savings associated with component sharing strategies and also include the quality ramifications of component sharing while formulating component sharing strategies.

Acknowledgments

We are grateful to executives, managers and engineers at Ford and General Motors, several automotive experts and race car enthusiasts, and the Department of Transportation for providing us with information, and to Noah Springer, Ryan Sundquist, Prasanna Velamuri and Baocheng Yang for research assistance. We are grateful to Mary Margaret Frank, Richard Holubkov, Nils Rudi, Christian Terwiesch and seminar participants at UCLA, Wharton, the University of Chicago, INFORMS, and the Wharton Technology Miniconference for their valuable comments and in particular to Steven Stern for his help with sampling related issues. This research was supported in part by a grant from the Batten Institute for Innovation and Entrepreneurship at the University of Virginia.

References

ADLER, P. (1990), “Shared Learning,” *Management Science*, Vol. 36. No. 8. pp. 938-958.

ADLER, P. AND K. CLARK, (1991), "Behind the Learning Curve: A Sketch of the Learning Process," *Management Science*, Vol. 37, No. 3, pp. 267-282.

ALLISON, P. (1995) *Survival Analysis Using the SAS System*. SAS Institute.

ANZONI, Y AND H.A. SIMON, (1979) "The Theory of Learning by Doing," *Psychological Review*, 86, pp. 124-140.

BALDWIN, C. Y. AND K. B. CLARK, (1998), "Managing Modularity", *Harvard Business Review*, 75, 5, 84-93.

BALOFF, N. (1971) "The Extension of the Learning Curve," *Operations Research Quarterly*, 22 pp. 329-340.

DESAI, P., S. KEKRE, S. RADHAKRISHNAN AND K. SRINIVASAN (2001), "Product Differentiation and Commonality in Design: Balancing Revenue and Cost Drivers", *Management Science*, 47, 1, 37-51.

DICK, A.R. (1991), "Learning by Doing and Dumping in the Semiconductor Industry," *Journal of Law and Economics*, 34, pp. 662-669.

COX, D.R. (1972) "Regression Models and Life Tables" *Journal of the Royal Statistical Society*, B34, 187-220.

FISHER, M. L., K. RAMDAS AND K. T. ULRICH (1999), "Component Sharing in the Management of Product Variety: A Study of Automotive Braking Systems", *Management Science*, Vol. 45, No. 5, 297-315.

FINE, C. (1986), "Quality Improvement and Learning in Productive Systems," *Management Science*, Vol. 32, No. 10, pp. 1301-1315.

GUPTA S. AND V. KRISHNAN (1999), "A Product Family-Based Approach for Integrated Component and Supplier Selection", *Prod. and Oper. Mgmt.*, 8, 2, 163-182.

HATCH, N.W. and D. C. MOWERY, "Process Innovation and Learning by Doing in Semiconductor Manufacturing", *Management Science*, Vol. 44, No. 11, pp. 1461-77, 1998.

HEESE, H. AND J.M. SWAMINATHAN (2004), "Product Line Design with Component Commonality and Design Investments," Working Paper, University of North Carolina at Chapel Hill.

HELSEN, K. and SCHMITTLEIN (1993) "Analyzing duration times in Marketing - Evidence for the effectiveness of Hazard Rate models," *Marketing Science* 12(4) 395-414.

ITTNER, C. (1996), "Exploratory Evidence on the Behavior of Quality Costs," *Operations Research*, Vol. 44, No. 1 pp. 114-130.

KIM, K. and D. CHHAJED (2001), "An Experimental Investigation of Valuation Change Due to Commonality in Vertical Product Line Extension", *Journal of Product Innovation Management*, 18, 4, 219-230.

KRISHNAN, V. AND S. GUPTA (2001), "Appropriateness and Impact of Platform-Based Product Development", *Management Science*, 47, 1, 52-68.

KRISHNAN, V., R. SINGH AND D. TIRUPATI (1999), "A Model-Based Approach for Planning and Developing a Family of Technology-Based Products", *Manufacturing and Services Operations Management*, 1, 2, 132-156.

KRISHNAN, V. AND K. T. ULRICH (2001), "Product Development Decisions: A Review of the Literature," *Management Science*, 47, 1, 1-21.

MACHER, J. T., "Vertical Disintegration and Process Innovation in Semiconductor Manufacturing: Foundries vs. Integrated Producers", working paper, Georgetown University.

MACDUFFIE, J.P., SETHURAMAN, K. AND M. FISHER (1996), "Product Variety and Manufacturing Performance: Evidence from the International Automotive Assembly Plant Study," *Management Science*, Vol. 42, No. 3. pp 350-369.

LIEBERMAN, M. (1984), "The Learning Curve and Pricing in the Chemical Processing Industries," *RAND Journal of Economics*, 15, pp. 213-228.

RAMDAS, K., "Managing Product Variety: An Integrative Review and Research Directions," *Production & Operations Management*, Vol. No. , pp. , 2003.

RAMDAS, K., M. L. FISHER, AND K. T. ULRICH (2001), "Managing Variety for Assembled Products: Modeling Component Systems Sharing", *Manufacturing & Service Operations Management*.

RAMDAS, K. AND M. SAWHNEY, (2001), "A Cross-Functional Approach to Designing Multiple Line Extensions for Assembled Products", *Mgmt. Sci.*, 47, 1, 22-36.

ROBERTSON, D. AND K.T. ULRICH (1998), "Planning for Product Platforms", *Sloan Management Review*, 39, Summer, 19-31.

RUTENBERG, D. P. (1969), "Design Commonality to Reduce Multi-Item Inventory: Optimal Depth of a Product Line," *Management Science*, 491-509.

SCHOONHOVEN, C.B. (1981), "Problems with Contingency Theory: Testing Assumptions Hidden within the Language of Contingency "Theory,"" *Administrative Science Quarterly*, Vol. 26, No.3, pp. 349-377.

SWAMINATHAN, J. AND S. TAYUR (1998), "Managing Broader Product Lines Through Delayed Differentiation Using Vanilla Boxes", *Mgmt. Sci.*, 44, 12, S161-S172.

SWAMINATHAN, J. AND S. TAYUR (1999) "Managing Design of Assembly Sequences for Product Lines that Delay Product Differentiation," To appear in *IIE Transactions*.

THONEMANN, U.W. AND M.L. BRANDEAU (1999), "Optimal Commonality in Component Design," Forthcoming in *Operations Research*.

ULRICH, K. T. (1995), "The Role of Product Architecture in the Manufacturing Firm", *Research Policy*, 24, 419-440.

ULRICH, K. T. AND D. J. ELLISON (1999), "Holistic Customer Requirements and the Design – Select Decision", *Management Science*, 45, 5, 641-658.

ULRICH, K. T., T. RANDALL, M. L. FISHER AND D. REIBSTEIN, (1998), "Managing Product Variety", in *Product Variety Management: Research Advances*, T.H. Ho and C. S. Tang (eds.), Kluwer Academic Pubs.

WARD'S AUTOMOTIVE YEAR BOOK, 1965 – 2003.

WRIGHT, T. (1936), "Factors Affecting the Cost of Airplanes," *Journal of Aeronautical Science* 3 (4), pp. 122-128.

YANO, C.A. AND G. DOBSON, (1998), "Profit Optimizing Product Line Design, Selection and Pricing with Manufacturing Cost Consideration", in *Product Variety Mgmt: Research Advances*, T.H. Ho and C. S. Tang eds., Kluwer Academic Pubs.

Exhibit 1

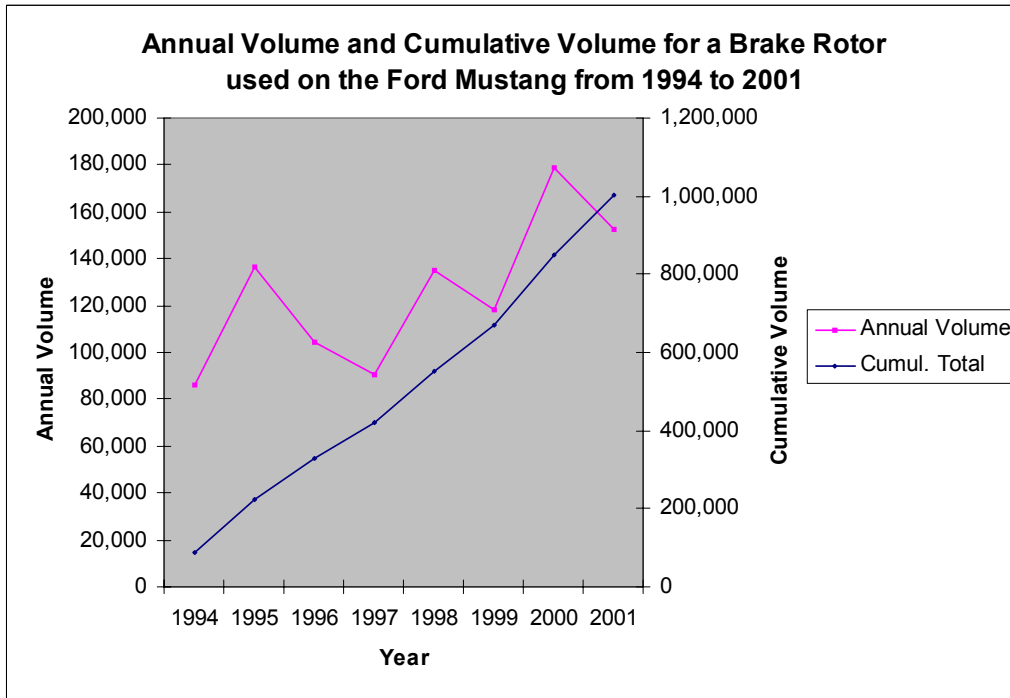


Exhibit 2

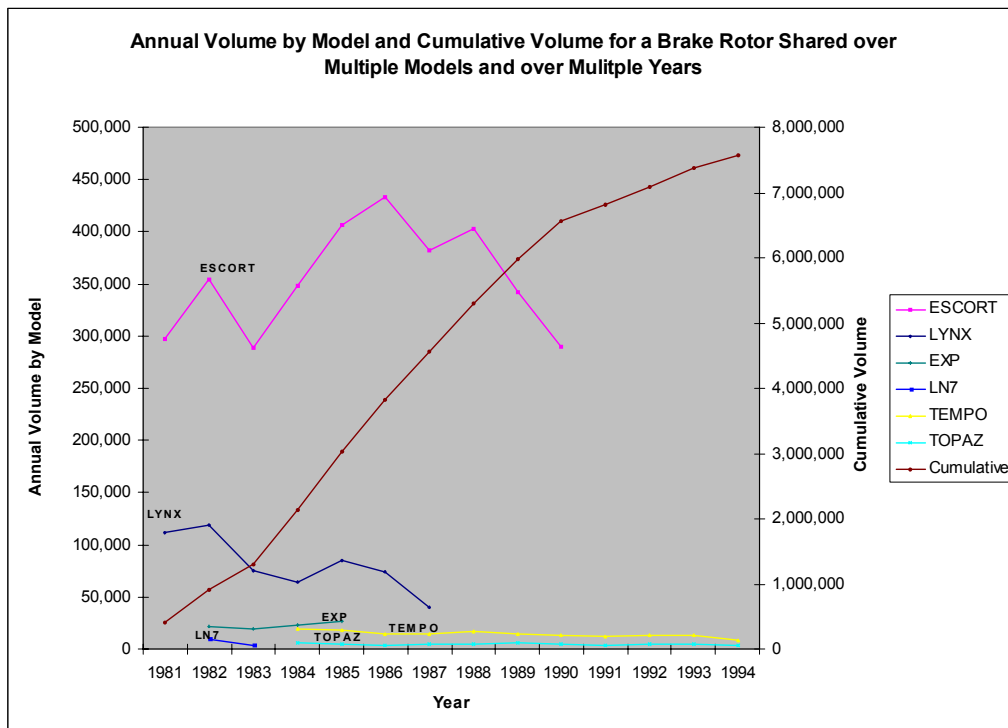


Exhibit 3

Histogram of Miles to Failure

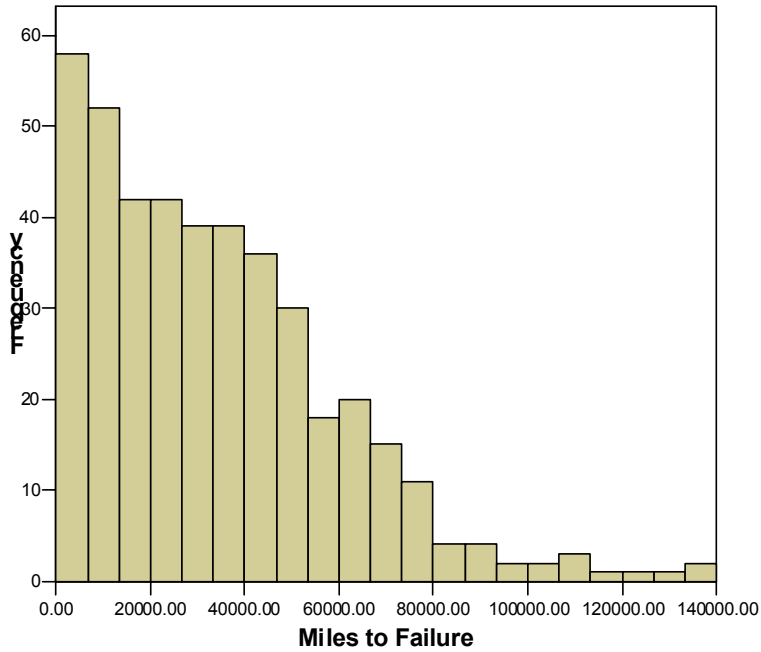


Exhibit 4

Cumulative volume for each unique front brake rotor used by Ford between 1983 and 2000

(70 unique front brake rotors in this period)

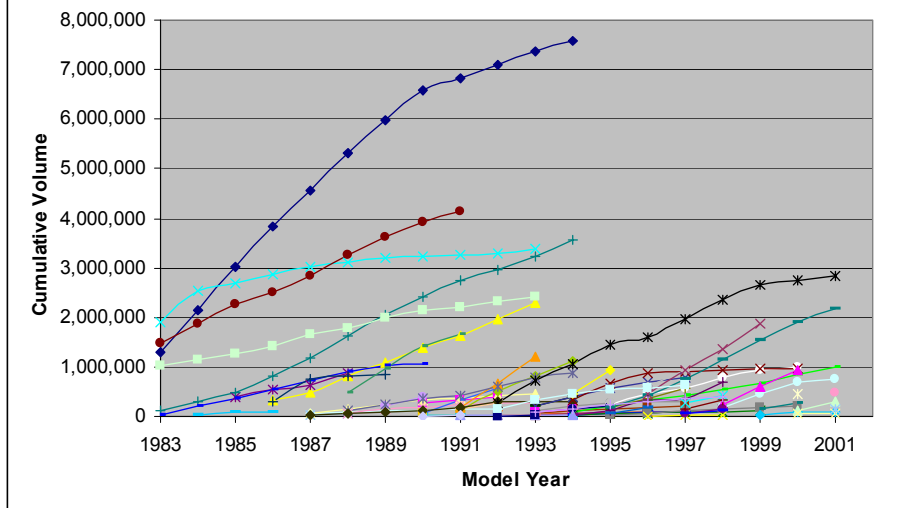


Exhibit 5

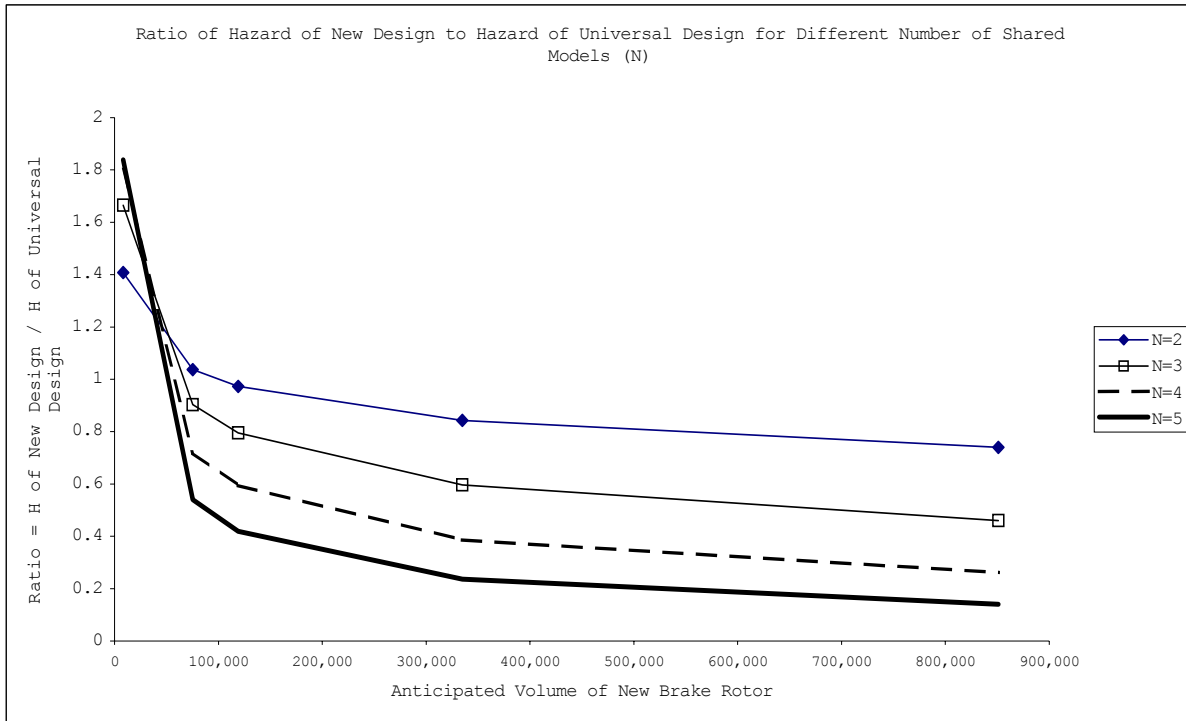


Exhibit 6

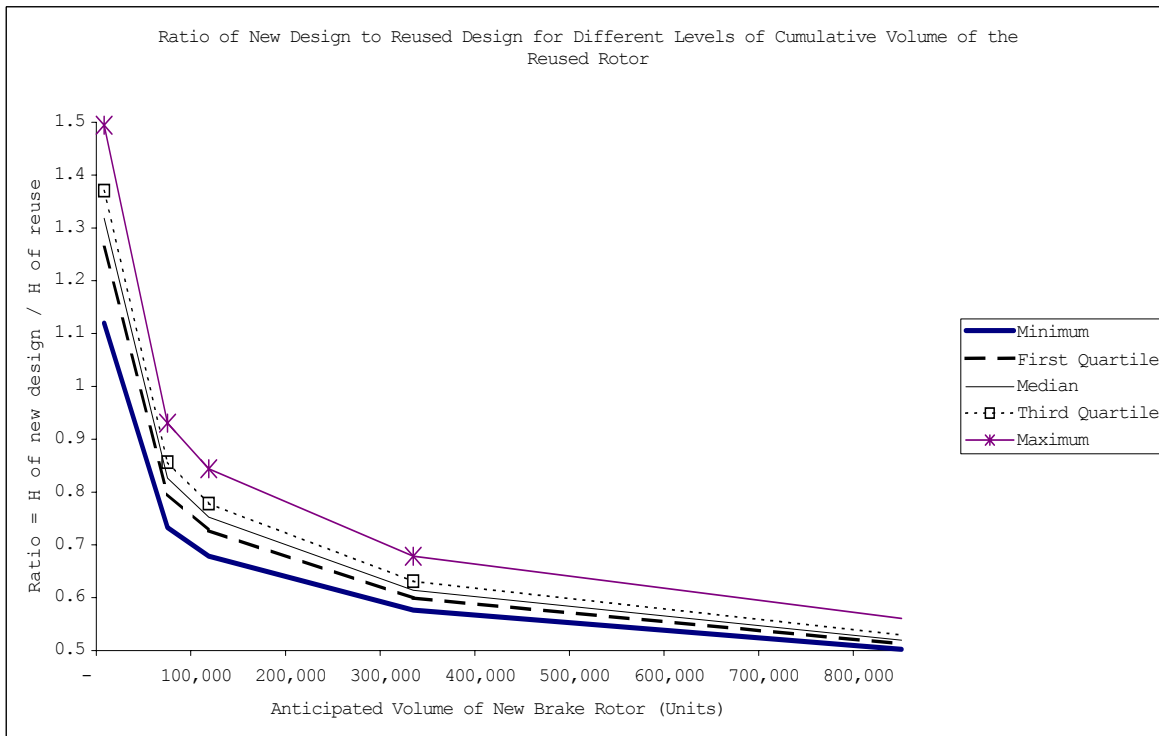


Table 1: Descriptive Statistics					
N=601					
	Mean	Median	Std. Dev.	Min	Max
Continuous Variables					
Miles to Failure	33,579	29,000	26,227	2	160,000
Cumulative Volume	880,178	554,816	1,180,627	8334	7,568,420
Number of Models	2.36	2.00	1.17	1.00	9.00
Horsepower	150	140	31	84	305
Swept Area per Ton	120	120	12.69	83.07	198.79
Precipitation	38.25	41	14.09	7	105
Population Density	3,442	2,636	3,691	18	34,916
% Positive Weight Difference	6.12%	4.78%	6.50%	0.00%	30.00%
Dichotomous Variables (% of sample)					
Rotor-Year-Model	36%				
Rotor-Model	85%				
Vented	87%				

Table 2: Pearson Correlation Coefficients (n=601)

	Miles to Failure	Cumulative Volume	Number of Models	Horsepower	Swept Area per Ton	Precipitation	Population Density	% Positive Weight Difference
Miles to Failure	1							
Cumulative Volume	-0.053	1						
Number of Models	-0.010	0.701***	1					
Horsepower	-0.094***	-0.334***	-0.324***	1				
Swept Area per Ton	-0.023	0.052	0.032	-0.257***	1			
Precipitation	0.102***	0.042	0.090**	-0.029	0.039	1		
Population Density	-0.001	0.007	0.021	-0.046	0.031	-0.047	1	
% Positive Weight Difference	-0.003	0.219***	0.240***	-0.320***	0.374***	0.011	0.075	1

***, **, * coefficient significant at the p<0.01, 0.05 and 0.10 levels respectively

Table 3: Cox Proportional Hazard Regression Estimating the Hazard of Brake Rotor Failure

(P-values of Wald Chi-Square Statistics in Brackets)

	I	II	III
Horsepower	0.005*** [0.003]	0.005*** [0.003]	0.005*** [0.002]
Swept Area per Ton	0.005 [0.284]	0.007 [0.145]	0.005 [0.288]
Precipitation	-0.005 [0.173]	-0.005 [0.165]	-0.005 [0.161]
Ln(Population Density)	0.094** [0.030]	0.100** [0.022]	0.088** [0.042]
Vented	-0.272 [0.135]	-0.259 [0.203]	-0.265 [0.192]
% Positive Weight Difference	-0.007 [0.994]	-0.181 [0.869]	0.162 [0.880]
Ln(Volume)		-0.355** [0.001]	-0.097 [0.669]
Number of Models		-2.046** [0.000]	-0.769 [0.127]
Ln(Volume) x Number of Models		0.139*** [0.000]	0.053 [0.139]
Rotor-Year-Model		-1.344*** [0.000]	-
Rotor-Model		-	-0.578 [0.821]
Rotor-Year-Model x Number of Models		0.502*** [0.001]	
Rotor-Model x Number of Models			0.040 [0.829]
Chi-Square Statistic	116.35***	139.11***	118.87***
R-squared	17.60%	20.69%	17.97%
N	601	600	600

***, **, * coefficient significant at the p<0.01, 0.05 and 0.10 levels respectively.

Table 4: Cox Proportional Hazard Models – Tests of Robustness

	[P values of Wald Chi-Square Statistics in Brackets]						
	I	II	III	IV	V	VI	VII
	Non- Endogeneity	Early Failures	North American Volume	Passenger Cars	Trucks	Alternative Volume Measure	Rotors with 4 or more observations
Horsepower	0.005*** [0.010]	0.006*** [0.001]	0.005*** [0.009]	0.004*** [0.039]	0.002 [0.656]	0.006*** [0.002]	0.004 [0.119]
Swept Area per Ton	0.000 [0.959]	0.005 [0.372]	0.009* [0.079]	0.008 [0.250]	-0.004 [0.667]	0.007 [0.152]	-0.001 [0.830]
Precipitation	-0.004 [0.337]	-0.006* [0.099]	-0.005 [0.171]	-0.003 [0.452]	-0.010 [0.174]	-0.005 [0.170]	-0.005 [0.144]
Ln(Population Density)	0.080* [0.094]	0.128*** [0.005]	0.099** [0.031]	0.121** [0.025]	0.083 [0.318]	0.099** [0.023]	0.096** [0.046]
Vented	-0.003 [0.990]	-0.268 [0.204]	-0.262 [0.192]	-1.472* [0.054]	0.111 [0.710]	-0.251 [0.213]	-0.033 [0.888]
% Positive Weight Diff.	1.961* [0.094]	0.163 [0.888]	-0.467 [0.685]	-1.098 [0.481]	2.688 [0.174]	-0.113 [0.917]	1.512 [0.217]
Ln(Volume)	-0.259** [0.019]	-0.439*** [0.000]	-0.376*** [0.000]	-0.488*** [0.000]	-1.217*** [0.002]	-0.334*** [0.001]	-0.123 [0.384]
Number of Models	-0.994 [0.106]	-2.527*** [0.000]	-2.262*** [0.000]	-2.565*** [0.000]	-9.301*** [0.002]	-2.027*** [0.001]	-1.325* [0.063]
Ln(Volume) x # Models	0.061 [0.147]	0.173*** [0.000]	0.156*** [0.000]	0.179*** [0.000]	0.686*** [0.003]	0.138*** [0.000]	0.082* [0.096]
Rotor-Year-Model	-1.031*** [0.002]	-1.514*** [0.000]	-1.400*** [0.000]	-1.306*** [0.004]	-3.030*** [0.000]	-1.528*** [0.000]	-0.901** [0.030]
Rotor-Year-Model x # Models	0.299* [0.062]	0.564*** [0.001]	0.541*** [0.001]	0.478** [0.020]	1.649*** [0.001]	0.590*** [0.000]	0.372** [0.047]
Chi-Square Statistic	149.34***	151.51***	135.93***	125.52***	41.01***	141.17***	116.02***
R-squared	24.88%	23.63%	21.22%	23.97%	25.08%	20.93%	18.96%
N	522	562	570	458	142	601	552

***, **, * coefficient significant at the $p < 0.01$, 0.05 and 0.10 levels respectively.

Appendix 1

Let X denote the vector of covariates – study variables and control variables – that the hazard of failure at mileage t is assumed to vary with. At any mileage t , we can state the hazard of failure given X generally as $h(t/X) = \lambda_0(t)\exp\{-\beta X\}$, where β is a vector of coefficients.

Let $p(\text{report}/t, X_1)$ denote the probability that the owner of a vehicle would file a report of a rotor failure, if a failure were to occur at mileage t and given a vector of covariates X_1 , where X_1 may partially overlap with X . The probability of reporting if a failure were to occur is likely to vary with the mileage at the time of failure, and also perhaps with some of the covariates in X . For example, if city drivers are more likely to complain than country drivers then the population density of the area where the vehicle was driven may impact the probability of reporting. Let $r(t/X, X_1)$ denote the hazard function for those individuals who both experienced a rotor failure *and* chose to report it. The hazard function that we can estimate from our data is in fact $r(t/X, X_1)$. By conditioning on the occurrence of a failure at time t , $r(t/X, X_1) = h(t/X)p(\text{report}/t, X_1)$.

Taking the derivative of the logarithm of $r(t/X, X_1)$ with respect to X , we have:

$$\frac{\partial \log r(t/X, X_1)}{\partial X} = \frac{\partial \log h(t/X)}{\partial X} + \frac{\partial \log p(\text{report}/t, X_1)}{\partial X_1} \frac{\partial X_1}{\partial X}$$

Since an individual car driver likely does not know the cumulative volume of the brake rotor used in his or her car, the number of models on which that rotor was used, or whether that rotor was designed specifically for the car, it is reasonable to assume that none of the variables of interest in our study are in the vector X_1 . Therefore the second term in the right-hand-side of the above equation does not impact our study variables. In interpreting the Cox regression coefficients for $r(t/X, X_1)$, the components of β for our study variables have the same interpretation as they would in a Cox regression that estimates the unobserved hazard $h(t/X)$.

Note that thinking in terms of the probability of reporting if a failure *were to occur* at time t allows us to model scenarios where a failure has not actually occurred by time t . It also allows us to model scenarios where the vehicle was taken off the road prior to time t due to some non-brake-related issue, precluding the possibility of a brake failure at time t . In this situation, if the reason the vehicle was taken off the road has nothing to do with why the brakes might fail then the components of β for our study variables would continue to remain unbiased.