

Impact of Regulatory Changes on New Product Design Choice and End-of-Life Inventory Buildup

Ananth V Iyer*, Svenja Sommer, Mohammad Saoud*****

***Krannert School of Management, Purdue University, West Lafayette,
IN 47907, Aiyer@purdue.edu**

***HEC Paris, sommers@hec.fr**

***Kuwait University**

Impact of Regulatory Changes on New Product Design Choice and End-of-Life Inventory Buildup

September 8, 2012

Abstract

New environmental regulations often impose stricter efficiency standards on manufacturers. The regulations impose specific cutoff dates by which manufacturers should cease production of the older version of the product and (potentially) offer new redesigned products with higher efficiencies. However, while the manufacture of the old product is banned, US regulations typically allow the continued sale of inventory in the supply chain after the cutoff date. We take the perspective of a single company, whose design choice only marginally impacts other firms in the industry (monopolistic competition). Our goal is twofold i.e., to identify the optimal design choice for the new product, and the associated inventory build up of the old product in advance of the cutoff date. We thus combine logistics considerations with new product design decisions to optimize profits. Further, we explore the effects of product, customer and industry characteristics on the manufacturers' choices.

Keywords: regulatory change, inventory build-up, product design

1 Introduction

US Regulation to improve environmental impact of products typically involves banning the manufacture of products with efficiency levels below a specified limit after a prescribed date. For example, when the US Department of Energy raised, by 30 %, the minimum energy efficiency rating for the air-conditioning equipment manufactured or imported into the US after the cutoff date of January 23, 2006, it permitted the continued sales of the old product. Thus, just before the cut-off date, both old and new designs might share production capacity. Following the cutoff date, the industry will manufacture only compliant designs. In such a case, given some customers' preferences for less efficient but also less costly equipment, there may be a build up of inventory of the old products until the cutoff date, which might coincide with the phase-in of manufacture of the new compliant designs, as it was the case in the air-conditioning industry (see below). The incentive to build old product inventory is in stark contrast the typical demand-driven product roll-overs, as we will further discuss in the literature review.

We take the perspective of a single company supply chain, whose design choice only marginally impacts other firms in the industry (monopolistic competition). Our goal is twofold - identify the optimal design choice for the new product and the associated inventory build up of the old product in advance of the cutoff date. We thus combine logistics considerations with new product design decisions to optimize profits. We also show how these decisions depend on different problem parameters. Our model and analysis may be useful to decision makers in the supply chain as well as regulators who need to anticipate the impact of their choices on industry.

The remainder of the paper is structured as follows: Section 1.1 introduces a specific example of the type of regulatory change we consider, namely that of the US air-conditioning industry, to motivate the problem and model assumptions. Section 2 discusses the related literature. Section 3 describes the model setup. Section 4 presents the optimal inventory choice, as well as the effects

of various problem parameters on the optimal design choice. We finally conclude with a discussion of managerial implications, both from the manufacturer and the regulators perspective.

1.1 The Case of the US Air-conditioning Industry

The problem studied in the paper is motivated by discussions with managers in the air-conditioning industry. On January 23th, 2006, the Department of Energy (DOE) banned the production of the most popular type of air-conditioners in the US market, the 10 SEER air-conditioners.¹ From that day onwards, the minimum efficiency air-conditioning equipment made or imported into the US would be 13 SEER, which represented a 30% increase in efficiency. However any finished goods inventories held across the US supply chain could be sold to customers until supplies ran out.

The increase in the minimum efficiency standard for air-conditioners from 10 to 13 SEER was initially announced on January 22, 2001. Due to industry opposition and counter lawsuits by consumer groups, much uncertainty remained about the new energy efficiency standard until January 13, 2004. On that date, the US Court of Appeals for the Second Circuit ruled that the increase in the minimum efficiency standard for air-conditioners would come into effect on January 23th, 2006, as planned. The ruling meant that from January 23th, 2006 onwards, the production or import of 10 SEER air-conditioners was banned, leaving a relatively short time (about 2 years) for manufacturers to adjust their products.

This regulation represented a major change for manufacturers, since 10 SEER air-conditioners were the most popular type of air conditioning in the US market, accounting for 75% of the residential air-conditioning sales (McGrory et al. 2000). However, the DOE's justification for this new requirement was the projected substantial energy savings (saving the equivalent of the annual energy consumption by 26 million households between 2006 and 2030) and a large environmental benefit (saving by 2030 "carbon dioxide gas emissions equivalent to removing 3 million cars from

¹SEER stands for Seasonal Energy Efficiency Ratio, one measure of cooling efficiency of air-conditioners.

our road ways each year”, source: Carrier).

The environmental regulatory change required manufacturers to redesign the product to be compliant with regulations. It also required them to cease production of the old product version by the specified date. However, it left the details of the redesign to the discretion of the manufacturers. To meet the new regulatory standards, different technological solutions were available, both with respect to product components and with respect to the product architecture, and different firms made different design choices, ranging from minor to complete redesigns. While more radical redesigns would result in higher customer benefits, they would also involve more radical changes in the manufacturing process. Choosing a technological solution for the 13 SEER production was only one of the decisions the industry was facing. Since the production but not the sale of 10 SEER air-conditioners was banned after the cutoff date, there was the consequent decision regarding how much 10 SEER inventory to build up for the future demand.

An industry survey by Emerson Climate Technologies showed that the majority of the wholesalers (52%) and even many contractors (41%) increased their 10 SEER inventory significantly in anticipation of the transition to 13 SEER. The importance of building inventory during the transition to higher efficiency equipment was also stressed by Watsco, the largest distributor of air-conditioning equipment in the US, for whom this build-up contributed a 20% increase in inventory levels (conversations with Watsco manager and 2005 annual report). Industry level data for the air-conditioning industry confirms a drastic increase of factory shipments towards the end of the 10 SEER production period (see Figure 1). The large December and January shipments suggest furthermore that many distributors waited with their inventory decisions.

This empirical evidence suggests that both the design choice and the inventory decision were important decisions, which were not answered uniformly by all firms. Under what circumstances should a firm make major design changes with significant early investments in old product inventory,

(Source: ApplianceMagazine.com)	Factory shipments	Change compared to same month last year	Change year to date
July 2005	689,098	25.0%	-0.5%
August 2005	593,745	43.7%	3.8%
September 2005	626,713	47.7%	7.7%
October 2005	524,778	84.7%	12.1%
November 2005	460,623	78.4%	15.4%
December 2005	426,620	54.2%	17.3%
January 2006	530,896	85.6%	85.6%
February 2006	395,620	13.5%	46.3%
March 2006	478,202	-19.2%	14.7%

Figure 1: Factory Shipment Data; Source: ApplianceMagazine.com and when is it better of to make only modest adjustments with the option to adjust the inventory levels after the demand scenario is better known? What factors should a decision maker consider? Our model attempts to provide a first answer to these questions.

The problem faced by the HVAC industry is not an isolated incident. Similar energy efficiency standards have been introduced for various industries. For example, the US energy efficiency standard for refrigerators increased in July 2001. More recently, on July 22, 2009, the European commission introduced new eco-design regulations to make industrial motors, water circulators, televisions, refrigerators and freezers more energy efficient (EURActiv 23/07/09). Other examples include safety and fuel efficiency requirements in the automotive industry. All these regulatory changes have one thing in common: They require a product redesign, which is driven by compliance considerations rather than customer preferences and for which the product introduction timing is determine by the regulation rather than being the manufacturer’s choice (and hence driven by demand considerations).

2 Related Literature

Our paper draws on various streams of literature. First, our paper relates to sustainable operations management for which there is an emerging body of research (Corbett and Kleindorfer 2001a, 2001b,

2003, Journal of Cleaner Production 2008, Corbett et al. 2008). As the review of Kleindorfer et al. (2005) demonstrates, the operations literature has paid little attention to the effect of regulatory changes on the supply chain. A recent exception is Plambeck and Wang (2009), who explore how different fee structures of various e-waste regulations affect the frequency of product introductions.

Second, the regulatory change requires a product redesign, which provides a connection to the product rollover literature. Building on earlier empirical work (Saunders and Jobber 1994, Greenley and Bayus 1994), Billington et al. (1998) differentiate two strategies for product rollovers: “solo-product roll” where the old product is being phased out as soon as the new product is phased in, and “dual-product roll” where the sales of the old and new product overlap. Lim and Tang (2006) present analytical results for conditions under which each of these strategies would be optimal for an *open* loop supply chain, and Iyer et al. (2008) for *closed* loop supply chains. Bhaskaran et al. (2010) look at a similar joint production and product design problem. However, unlike in our problem, they look at settings where all customers clearly prefer the new product, and any left-over inventory incurs additional costs. Therefore, manufacturers have to decide the final build and the new product introduction timing, considering both its effect on the design quality of the new product, and the potential leftover inventory and lost sales costs of the old product. El Khoury and van Delft (2011) study the optimal rollover strategy when the availability date of the new product is stochastic. In our context the regulatory change determines the cut-off date for the production of the old product, while it permits continued sales of manufactured units of the old product as long as it is available in inventory. Compared to this stream of literature, the main difference is that the new product introduction is forced on the manufacturers despite customer preferences. Thus, while in demand-driven settings the old product inventory is worthless (or less valuable) once the new product has been introduced, in regulatory-driven product changes manufacturers have incentives to build inventory and continue to serve those customers preferring the old product.

Third, manufacturers have to make inventory decisions both before and after a demand update

has been obtained. This setting relates our paper to numerous papers in the supply chain literature, and we can point the reader only to a few examples. The inventory decision is very similar to the situation considered in the papers by Donohue (2000) and Gurnani and Tang (1999). Donohue (2000) studies the optimal supply contract (with returns) if the retailer has two instances to order a product from a manufacturer, a low cost one before and a higher cost one after an updated demand forecast has been obtained. Gurnani and Tang (1999) focus on optimal order quantities with forecast updates, if the second order costs are uncertain. In contrast, we will assume that the second order costs are higher, but the exact costs is in fact a choice variable that is affected by the new product design.

Fourth, our paper is related to the literature on technology choice during product upgrades. Product upgrades are done for various reasons; they can be consumer driven, competition driven, or due to regulatory changes. For research regarding the timing of new product introductions, see Banerjee and Sarvary (2009). See Kornish (2001) for a model related to market segmentation via price discrimination. Another stream within this literature looks at technology diffusion. While most of the literature focuses on technology adoption of end customers (including several recent papers in the operations management literature, see for example Ho et al. 2002, and Kumar and Swaminathan 2003), Robertson and Gatignon (1986) review the technology diffusion literature on an industry level. Their review suggests among others that the technology diffusion will be more rapid in case of higher demand uncertainty, changes in consumers' needs, and higher competitive intensity, with monopolistic conditions especially reducing innovativeness.

The goal of this paper is to merge these different streams of literature and model a context where regulation enforces a product design change. The decisions we explore are unique in that we link the optimal design choice of the new product to the logistics cost impact of accommodating the manufacture of the old product until the cutoff date.

3 Model Setup

In this model, we take the perspective of a single company supply chain, whose design choice only marginally impacts other firms in the industry (monopolistic competition), and hence does not impact other firms' choices. Furthermore, we consider an integrated supply chain consisting of a manufacturer who owns his distribution channel. This setup reflects the situation in the air-conditioning industry, where distributors are often owned by manufacturers or closely collaborate with just a few of them under exclusivity contracts. It also proxies for settings in which supply chain coordinating contracts are used.

We first describe the timing of decisions, the available information and associated costs (see Figure 2). We will differentiate four main phases: (1) the early production period between t_1 and t_2 , during which the firm builds the initial inventory Q_1 of the old product for the transition period at the standard manufacturing cost c_1 ; (2) the capacity overlap period between t_2 and $t_{effective}$, during which the firm may build additional inventory Q_2 of the old product for the transition period at higher costs c_2 ($c_1 \leq c_2$) but with the advantage that at this time additional market information (modeled by a signal s) is available (i.e., Q_2 is a function of both the manufacturing costs c_2 and the signal s); (3) the transition period between $t_{effective}$ and T , during which both the new product and the total old product inventory ($Q_1 + Q_2$) may be sold; and finally the steady state period after T (see below).

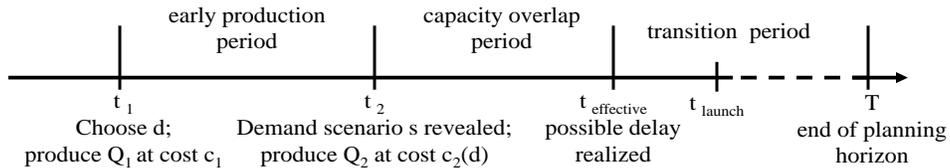


Figure 2: Timeline of Old Product Quantity Decisions

In addition to deciding the old product inventory to be build for consumption during the transition period ($Q_1 + Q_2$), the firm has to choose at t_1 the optimal design d for the transition

period. Similar to typical quality choice models, we capture this multi-dimensional construct by a single parameter d . This choice will influence the value of new products sold during the transition period (and hence the demand faced by both the old and the new product), but also affects the risk of finding design flaws only late during the ramp-up of the new product resulting in a new product introduction delay of $t_{launch} - t_{effective}$; whether such a delay is realized will be known by the date the regulation becomes effective $t_{effective}$ (see Section 3.2 for details). The design choice will also influence the cost $c_2(d)$ of producing the old product after the ramp-up of the new product started at time t_2 , that is during the capacity overlap period (see Section 3.1 for details). Note that the “capacity overlap” does not refer to the full scale production of the new product, which will begin only at t_{launch} , but rather is caused by the early ramp-up necessary to test the manufacturing process and makes adjustments to the process to achieve efficient full scale production of the new product once it is launched.

We focus on the optimal design choice for the new product during the transition period during which both products are sold. After T (i.e., during steady state) firms might redesign the new product further, since its design no longer impacts the old product’s manufacturing costs and demand; we refer to the optimal design choice that need not consider the old product as steady state design (d^{max}). Given the optimal design choice for the transition period d^* , the choice between a transition period design d^* and directly moving to a steady state design d^{max} obviously depends on the cost of redesign, let’s call them FC . If the profits earned during the transition period with an intermediate design ($E(\Pi_{Total}(d^*))$) are significantly larger than the profits that would be earned during the transition period with the stable design ($E(\Pi_{Total}(d^{max}))$), these additional profits may outweigh the redesign costs FC . That is, if $E(\Pi_{Total}(d^*)) - E(\Pi_{Total}(d^{max})) > FC$, firms should choose an intermediate design; otherwise, they should move directly to the steady state design.²

²If a proportion of the redesign cost is a function of the design choice d it could be directly incorporated into the profit function of the new product.

Since this comparison does not add any surprising insights, we focus the remaining discussion below on the scenario where this transition period design choice is relevant. The following subsections describe the setup, first the production side link between the two product versions, and then the market model for the demands realized during the transition period. We also link the modeling choices to industry information to provide a context and managerial insights.

3.1 The Design - Production Cost Link $c_2(d)$

When deciding the design of the new product, the firm does not only influence the quality of the new product, but it also influences the production costs of old product during the capacity overlap period. More radical redesigns would not only result in a better new product (see Section 3.2 for details), but they also require more radical changes in the manufacturing process and hence more setups, if the old and the new product share capacity close to the selling season. Firms typically ramp-up the production of the new product well before the new product is actually launched. In this initial “trial production period”, the firm will produce first units of the new product to discover design flaws or production difficulties that need to be resolved before the actual product launch and full volume production starts. If this is done while the old production continues, this will necessitate additional setups. Alternatively, the company could decide to discontinue production of the old product during the ramp-up of the new product, the effect of which is equivalent to setting a very large c_2 in our model.

In case of the air-conditioning industry, minor modifications in components of existing products could be handled by the existing manufacturing process, while “the whole assembly line could have to be changed to accommodate the more efficient models” (Source: www.qualityserviceinc.com). Thus, the design choice affects the manufacturer’s costs to continue producing the old product during the new product production ramp-up period. The more radical the design changes and the more integral the new products architecture, the less common equipment can be used, making

either an entire separate assembly line or more tool setups necessary.

The strength of this link between the new product design choice and late production costs (i.e., at t_2) for the old product depends on product architecture of the two products. In particular, we build on the following arguments from the literature: One, a more modular design reduces the cost of producing several product variants at the same time; two, a more modular product architecture reduces the degree to which an older product variant needs to be changed to achieve the desired functional change; and three, to obtain very high performance levels for “global performance measures” (e.g., size or mass) an integral architecture may be required (Ulrich 1995). The first argument is underlying models of product variety such as Lee and Tang (1997), who model product modularity as a strategy for delayed differentiation; the second is the driver in papers considering the speed of product changes; for example, Baldwin and Clark (2000) treat product modularity as an option to introduce a product upgrades every time a new technology for a product component comes up.

Building on these arguments, we claim that the link between the new product design d and the cost of late production c_2 might differ depending on the product architecture. A completely modular design of the old product would suggest that, for a larger range of d choices, the late production costs are not (or only marginally) influenced by the ramp-up of the new product and the late production cost does not differ much from the early production cost c_1 , i.e., $c_2 = c_1$ for a large range of d . Only for very high performance levels of the new product, the latter will require an integral architecture, and hence the late production costs would be very high. On the other hand, a completely integral product architecture of the old design suggests that even minor redesigns from the customer point of view require significant changes in the manufacturing process and hence the production costs increase drastically even for small d (see Figure 3). Either of these two settings, as well as intermediate ones, can be represented by a logistic function relationship between c_2 and d , e.g., by $c_2 = k_c + \frac{\alpha_c}{1 + \beta_c e^{-\gamma_c(d - \delta_c)}}$, where $k_c \geq c_1$. A large β_c captures a more modular design, and

a small β_c reflects a more integral design.

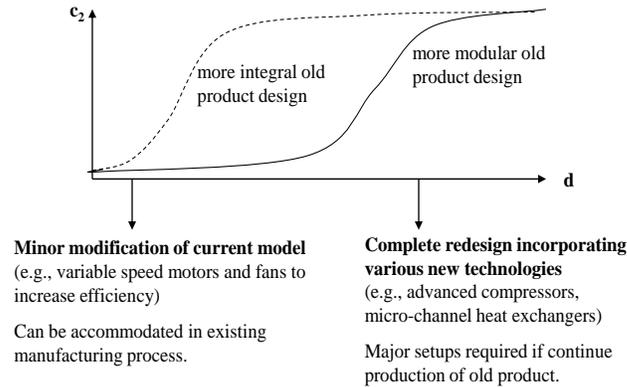


Figure 3: Link between Design Choice and Late Manufacturing Costs

3.2 Customer Utility

Unlike manufacturers choosing new product introductions in an attempt to better meet market needs, regulatory changes in fact stop the production of a product that might best meet the needs of a large, cost conscious customer segment. This was the case, in the air-conditioning industry where a large percentage of customers (75%) historically bought the product that just met regulations. Given such customer preferences, since production for the old product ceases when the regulation comes into effect (at time $t_{effective}$), but the demand during the transition period can still be satisfied from inventory, we need a demand model for both products during this transition period; (see Figure 2 for the time line of decisions and events). This section describes the factors influencing the demand during this period for the new and the old product respectively.

Impact of New Product Design Choice on Customer Utility

While many possible design choices may meet the regulatory requirements, they also affect product attributes valued by the customers and thus the utility derived by the end customer. For the air-conditioning industry to meet the new regulatory standards, different technological solutions were available. The redesign could range from minor modifications to components of a current model to a

complete redesign of the product's architecture and the incorporation of various new technologies.³ More radical redesigns and inclusion of newer technologies can increase the energy efficiency of the 13 SEER units, reduce the noise level, reduce its size, increase its durability (e.g., corrosion resistance), increase its serviceability (e.g., diagnostic system), decrease the amount or adjust the type of refrigerant required (to be phased out R-22 refrigerant versus 2010 compliant R-410A refrigerant), or improve the humidity control (SOURCE: company brochures of HVAC equipment manufacturers Heil, Carrier, Coleman, Goodman). Furthermore, more integral product designs have been shown to potentially provide higher product performances for integrative performance measures, such as weight or size (Ulrich 1995), which was also the case in the air-conditioning industry. In summary, regulatory compliance for 13 SEER products could be attained either through component changes, technological changes or product architecture changes. The choice of how to achieve compliance would impact product size and other customer sensitive attributes that would affect customer utility and thus potential market share.

While the design choice is clearly multi-dimensional, for clarity we aggregate these choices into a single parameter d , which reflects how much the design change improves the product from a customer point of view. (As mentioned earlier, a similar approach is being followed in typical quality choice models.) More specifically we assume that the new product's (average) utility depends in d in the following way: $U_n = k_n + \beta_n e^{\gamma_n(d - \delta_n)}$. Wassenaar et al. (2005) discuss in more detail the link between design choices and customer satisfaction for different types of product attributes. In terms of their classification (excitve, basic, must-be attributes), we assume the changes (beyond

³The component choices included coil features and dimensions, cabinet size, fin configuration, compressor type, accumulator, controls and motor type. Instead of replacing components, firms could also consider incorporating new technologies such as variable speed motor control, combined with advanced control technologies, which promised to provide a cost effective SEER increase; microchannel heat exchangers, which replaced the conventional round tube plate fin (RTPF) heat exchanger with an aluminum based microchannel coil, improved heat transfer and lowered fan power consumption, allowing for the production of smaller and lighter air-conditioners and reducing the volume of refrigerant required; and advanced modulating compressors, which could function as two capacity compressors and yield significant increases in the energy efficiency.(SEER is only one measure of energy efficiency; an alternative measure, EER focuses on the efficiency of the A/C when it runs at high temperatures. When design changes are made, both of these do not necessarily go in the same direction; Source: Chapter 4 Engineering Analysis by the US DOE: http://www1.eere.energy.gov/buildings/appliance_standards/residential/pdfs/chap4.engineer.pdf)

those required for regulatory purposes) are done to “excitive” attributes.

One factor that affects customer utility is price. For the purpose of the paper, we assume that the margin is fixed and given. On the one hand, any prices set for this transition period will affect prices for future products (i.e., it creates a reference price); on the other hand, the firm in a monopolistic setting is not free in making this choice, but has to consider “market prices”, i.e., what the competition charges. Since both aspects are outside our model, and since a price setting model would require detailed assumptions how these factors influence our focal firm’s decisions, we simply take the margin as given. This assumption is, however, not as restrictive as it sounds (see explanation in Section 3.3).

The Risk of Product Introduction Delays and its Impact on Customer Utility Our model will also explore the effect of possible product introduction delays. New product introductions might be delayed either because design problems are identified late during the production ramp-up or a manufacturer faces supply bottlenecks for radically new components, as it was for example the case for the Boeing Dreamliner. For simplicity, we model product introduction delays l as a zero-one variable (a delay takes place or not), with $prob_l(d)$ representing the probability of a delay (or lag) in the product introduction. We assume the uncertainty about this delay will be resolved only when the regulation comes into effect ($t_{effective}$), at which time either a fixed delay is realized ($l = 1$) or no delay is realized ($l = 0$). (Note that it makes no difference whether the delay is realized at $t_{effective}$ or slightly earlier, as long as the production quantity of the old product can no longer be adjusted, i.e., because production schedules are fixed or additional components could not be made available quick enough.)

Specifically we define the delay function as $prob_l(d) = \frac{\alpha_l}{1 + \beta_l e^{-\gamma_l(d - \delta_l)}}$, hence capture the fact that more extensive or radical design changes increase the risk of delays. Increasing the parameter α_l let’s us further explore the effect of increased delay risk (for all levels of d but with a more pronounced effect for larger d), in case the available time for redesign and ramp-up are shorter.

For example, in the particular case of the air-conditioning industry, given the number of proposed changes in the regulation (e.g., lowering minimum efficiency rating to 12 SEER), and lawsuits following these proposed changes, the industry had fairly little time (two years) to come up with a regulation compliant new product design and to secure the necessary supplies. The risk of new product delays was therefore significant in this context.

If the new product is indeed late to market, customers will experience a loss in utility, since they obtain the product only with a delay. Some customers might not be willing to wait and might consider switching to another product (the old product or the outside option), which will be captured automatically by the changed market shares resulting from the market share model discussed in the next section. Recall that we model customer utility of the new product as $U_n(d, l) = k_n + \beta_n(l)e^{\gamma_n(d-\delta_n)}$. We capture the utility loss due to a delay by assuming a lower $\beta_n(l)$ in case of a delay ($l=1$); note that the utility loss will be the larger, the larger d ; (larger utility that can be lost / discounted). The effect of possible delays depends on the customers' propensity to switch in case of unavailability of their preferred product. Customers should be more likely to switch, and therefore the firm's market share would be smaller, if customers are less patient. Hence the degree to which $\beta_n(l)$ changes in l can capture customer patience.

3.3 Market Share Model and Profit Functions

Our demand model considers two types of customers: A proportion s of customers, who are so cost conscious (or so budget constraint) that they only consider buying the old product, and if not available will switch to other options (like window units), and a proportion $(1 - s)$ of customers who consider both products and will buy the new one, if it offers enough benefits over the old product, i.e., their selling decision will depend on the new product design d . This proportion s is not known to the manufacturer and may be influenced by a number of external factors.⁴ In the model, we

⁴In the case of the air-conditioning industry, there was uncertainty about incentive programs by utility companies (which could make the new air-conditioner more affordable), and individual counties' willingness to give building permits for houses with below standard air-conditioners (which could remove the 10 SEER option completely).

capture the uncertainty resolution by assuming that at time t_2 , i.e., before Q_2 is decided, a signal s enables the old product demand distribution to be updated. We denote the pdf of possible demand signals by $h_S(s)$.

To determine the market share of each of the two products, we follow the idea of customer choice models (for an overview of customer choice models see McFadden 1986 or Corstjens and Gautschi 1983). In these models, the utility of a product is defined as a function of various product attributes plus a random disturbance, which reflects, for example, unmeasured preference variations. Thus, the manufacturer does not face a winner-takes-all-market, but the customer choice depends on the realization of these random disturbances, and the product with a higher average utility does not capture the full market, but it is only more likely to be chosen by a customer. As discussed above, the proportion $(1 - s)$ of customers consider both product types in their purchasing decisions. We already defined U_n above. In addition, we denote the (average) utility of the old product for the customers considering both products with U_o . In addition to the two products offered, these customers also have a number of outside options. In the air-conditioning industry these options include for example buying a window unit air-conditioning, buying an even higher SEER standard product, or in case of replacement demand, repairing the existing air-conditioner. For simplicity, we refer to these as one outside option with (average) utility U_x .

We follow the logic of the most widely applied customer choice model, the multinomial logit model (see for example Cohen et al. 1996 and Cachon et al. 2008). Using this logic, the market shares of the old and new product among customers considering both options are $\frac{U_o}{U_o+U_n(d,l)+U_x}$ and $\frac{U_n}{U_o+U_n(d,l)+U_x}$ respectively. Thus, once the proportion s and the delay are realized, the overall probabilities of customers buying the old and new product are $prob_o(d, l, s) = s + (1 - s) * \frac{U_o}{U_o+U_n(d,l)+U_x}$

Hence, the advantage of waiting to choose the production quantity for the old product closer to the selling season is that more demand information regarding potential demand scenarios s is available. Indeed, closer to the end of the year several utility companies announced their incentive programs for customers buying higher SEER air-conditioners, and some counties like Phoenix and Las Vegas stated that they will only allow air-conditioners with 13 SEER or higher ratings to be installed in new buildings.

and $prob_n(d, l, s) = (1 - s) * \frac{U_n}{U_o + U_n(d, l) + U_x}$ respectively.

Demand Distributions and Profit Functions

We can now define the conditional demand functions for the two products. Since the production of the new product continues throughout the selling season, the manufacturer is simply replenishing any inventory that his supply chain sells to the end customers and hence can adjust to the uncertainty in demand. Hence, while the demand signal impacts the actual new product profits, we only need to know the expected demand to determine the expected profits of the new product. Let M denote the total number of customers with air-conditioning needs. Note that a fixed M does not imply a fixed market-size in the traditional sense; given the market share model, some customers do not buy either of our products but rather turn to the outside options, which includes the option not to fulfill this need. Since with increasing d more customers switch from the old product but also from the outside option to the new product, the total demand met by the manufacturer and hence its market will increase. We can now express the expected number of customers buying the new product:

$$\mu_{D_N} = prob_l(d) * \int_0^1 M * prob_n(d, l = 1, s) h_S(s) ds + (1 - prob_l(d)) * \int_0^1 M * prob_n(d, l = 0, s) h_S(s) ds$$

Let m represent the margin the manufacturer derives from selling the new product. The expected profits the manufacturer derives from the sales of the new product $E[\Pi_N]$ are thus given by:

$$E[\Pi_N] = m * \mu_{D_N} = m * (prob_l(d) * \int_0^1 M * prob_n(d, l = 1, s) h_S(s) ds + (1 - prob_l(d)) * \int_0^1 M * prob_n(d, l = 0, s) h_S(s) ds)$$

As discussed above, we assume the margin m is fixed for the purpose of the transition period, since price setting should be influenced by a number of factors outside this model (e.g., long time pricing, market prices). Note that this assumption is, however, not that restrictive. Both the

performance attributes and the price are aspects that in reality affect the customer choice and hence demand. Increasing the margin and hence the price will have a negative impact on demand and vice versa; so whether d increases profits by increasing demand or whether a higher d allows the firm to charge a higher margin has the same effect on the overall profits.

Obviously, d might also influence the margin via the cost. There might be levels of design change for which the effect on cost might outweigh the benefits (in demand or price increases), and hence the expected profits should decrease. That is, there is a maximal design choice d^{max} that can feasibly be considered, since already new product profits alone will decrease beyond this point. Our profit maximization model is hence only valid for “feasible” design changes, i.e., changes for which the expected profits increase. Hence, $d^* \leq d^{max}$. Note that d^{max} , which is also the optimal steady state design referred to earlier.

Similarly, for the old product, the expected demand is given by:

$$\mu_{D_O} = prob_l(d) * \int_0^1 M * prob_o(d, l = 1, s) h_S(s) ds + (1 - prob_l(d)) * \int_0^1 M * prob_o(d, l = 0, s) h_S(s) ds$$

However, since the demand for the old product has to be met from inventory, the manufacturer faces effectively a newsvendor type of inventory decision for this transition period. Hence, for the old product, we need to know the full demand distribution. Let $f_D(\xi)$ denote the demand distribution for the old product before the signal s is observed, and $f_{D|s}(\xi)$ denote the updated demand distribution. For the rest of this paper, we will assume that the cumulative demand conditional on scenario s $F_{D|s}(\xi)$ is nonincreasing for every value ξ . Finally, let sv and p represent the salvage value and price of the old product. We can now express the expected profits the manufacturer derives from the sales of the old product $E[\Pi_O]$ as:

$$\begin{aligned}
E[\Pi_O(Q_1, c_2)] &= \int_0^1 \left\{ p \int_0^{\max(Q_1, Q^*(s))} \xi f_{D|s}(\xi) d\xi \right. \\
&+ sv \int_0^{\max(Q_1, Q^*(s))} (\max(Q_1, Q^*(s)) - \xi) f_{D|s}(\xi) d\xi \\
&+ p(1 - F_{D|s}(\max(Q_1, Q^*(s)))) \max(Q_1, Q^*(s)) \\
&- c_2(\max(Q_1, Q^*(s)) - Q_1) - c_1 Q_1 \left. \right\} h_S(s) ds
\end{aligned}$$

where $Q^*(s)$ is the optimal total quantity for the demand conditional on realization s , and therefore we can rewrite $Q_2^*(s) = (Q^*(s) - Q_1)^+$. Note that (wlog) we are excluding goodwill cost, since we assume customers are aware of the old product being phased out.

The Profit Maximization Problem

The manufacturer has two decisions to make: the design d and the quantity Q of the old product to build for the transition period. At the outset (time t_1) he chooses d to maximize the overall expected profits $E[\Pi_O] + E[\Pi_N]$, and the optimal first order quantity Q_1 to maximize $E[\Pi_O | d]$. At time t_2 , after the demand scenario is revealed, he chooses the optimal second order for the old product to maximize $E[\Pi_O | d, Q_1, s]$ (for the timeline of events and decisions see also Figure 2 above). In the next section, we will solve this problem backward, hence start with the optimal quantity decisions, and then study the optimal design choice.

4 Results

In this section, we will first present analytical results for the end-of-life inventory decision for the old product. We then turn to the optimal design decision and turn to numerics to understand the influence of various factors on the design and the inventory build-up decisions.

4.1 End-of-Life Inventory Build-up for the Old Product

Using backward induction, we first look at the end-of-life inventory decision for the old product.

In the profit function above, we already presented the optimal second order quantity given the demand signal s : $Q_2^*(s) = (Q^*(s) - Q_1)^+$, where $Q^*(s)$ represents the optimal total quantity for the demand conditional on realization s , i.e., the quantity that solves the standard newsvendor problem given s . We will now show properties of the optimal first order quantity Q_1^* .

Proposition 1 *Given a cost c_2 (determined by the design decision d), $E[\Pi_O]$ is concave in Q_1 . Hence, for every c_2 , there exists a unique corresponding optimal quantity $Q_1^*(c_2)$.*

Proof. All proofs are provided in the Appendix.

We now have the existence and uniqueness of Q_1^* . For the standard single period inventory decision with demand uncertainty or the newsvendor model, it is well known that the optimal quantity is the level whose cumulative probability is the critical fractile. But what about our case? We know that no matter what the choice of Q_1 is, the conditional service level will always be $\geq \frac{p-c_2}{p-sv}$ because, at the second production opportunity, the optimal service level is $\frac{p-c_2}{p-sv}$. The next proposition provides an operational explanation of Q_1^* . First, we define the conditional service level as the service level over a conditional demand distribution of a realized signal s (i.e. $D | s$), and given a quantity $\max(Q_1, Q^*(s))$; then we have the expected service level expression as $\int_0^1 F_{D|s}[\max[Q_1, Q^*(s)]] h_S[s] ds$. The unconditional service level is the service level when we cannot observe the signal and have to make a single quantity decision Q_1 (i.e. the unconditional service level is $F_D(Q_1)$).

Proposition 2 *The expected conditional service level (across all possible signals s) for the optimal purchases $Q_1^*(c_2)$ is equal to the unconditional service level $\frac{p-c_1}{p-sv}$.*

Proposition 2 suggests a graphical approach to obtain $Q_1^*(c_2)$. Figure 4 shows the service level as a function of the realized scenario s for a fixed Q_1 . The service level is nonincreasing because of the assumption that the cumulative demand distribution $F_{D|s}(\xi)$ is nonincreasing in s for every ξ . In addition, the conditional service level is constant for signals greater than or equal to $s_{c_2}(Q_1^*(c_2))$,

the signal for which $Q^*(s) \geq Q_1$. That is, for these signals, the firm will make a second purchase ($Q_2(c_2, s) \geq 0$) and hence always achieve a service level of $\frac{p-c_2}{p-sv}$. Proposition 2 states that the optimal $Q_1^*(c_2)$ should be such that the expectation of the graph in Figure 4 is equal to $\frac{p-c_1}{p-sv}$. Thus, if for any chosen Q_1 , the expectation is lower than $\frac{p-c_1}{p-sv}$, then increasing Q_1 will cause the graph to shift up from the thick solid line toward thin solid line, i.e., the change is nondecreasing for every signal realization, thus the expectation will be nondecreasing in Q_1 . A binary search over Q_1 can generate the optimal decision $Q_1^*(c_2)$.

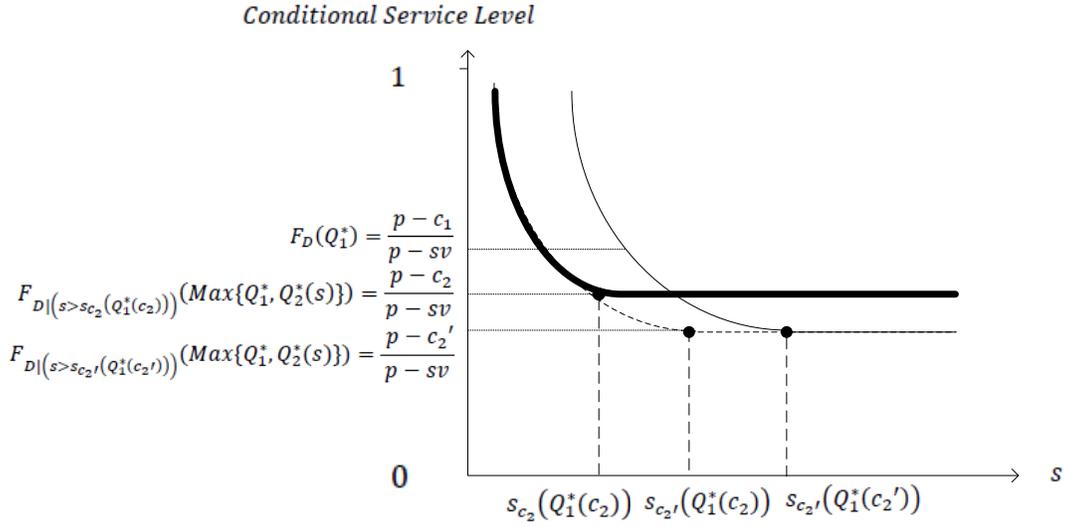


Figure 4: Conditional Service Level versus the Signal

Figure 4 also shows the impact of changing c_2 . Note that at this stage the design d has been decided. The following discussion therefore focuses on the effects of a change in the function $c_2(d)$ for some fixed level of d (rather than for the optimal d^* , which may change). Suppose for the same level of fixed d , c_2 increased to c_2' (that is the link between design choice and cost is stronger). If we maintain the initial inventory as $Q_1^*(c_2)$, then the dotted line shows that the overall function is less than or equal to the thick solid line, thus the expectation decreases to a value below $\frac{p-c_1}{p-sv}$. Note that it is then optimal to increase the value of Q_1 , the thin solid line shows that with $Q_1^*(c_2')(\geq Q_1^*(c_2))$, the expectation for the unconditional service level can be made to match $\frac{p-c_1}{p-sv}$. This result is

formally stated in Corollary 1.

Corollary 1 $Q_1^*(c_2)$ is a non-decreasing function in c_2 .

Given the impact of c_2 on $Q_1(c_2)$, we expect the expected profit to be non-increasing in c_2 , this is stated formally in Lemma 1.

Lemma 1 $\Pi_O(c_2)$ is a non-increasing function in c_2 .

Given the relationship between c_2 and d , the last two results capture the *indirect* effect of changes in d on the old product's profitability. Note that d affects at the same time the demand for the old product via the market share model, an effect that is opposite to the one described here. We will explore the overall effect in our numerical section.

4.2 New Product Design Choice

While we could show how to solve the inventory build-up problem for the old product, no closed form solutions exist. We can therefore not solve the optimal design choice problem analytically. We turn to numerics to understand how different factors affect the profit functions of the two products and hence the optimal design choice. To do so, we need to make specific assumptions about the signals, as well as the demand distribution.

Distributional Assumptions

To develop intuition, we will focus on the special case with two possible demand scenarios - a low proportion of customers that only consider purchasing the old product ($s = s_L$) or a high proportion of such customers ($s = s_H$) with probabilities $prob(s = s_L)$ and $1 - prob(s = s_L)$ respectively. Furthermore, we define the conditional demand distribution for the old product (perceived at time t_2) as the weighted average of two binomial distributions:

$$f_{D|s}(\xi) = prob_l(d) * B(M, prob_o(d, l = 1, s)) + (1 - prob_l(d)) * B(M, prob_o(d, l = 0, s)).$$

Base Case

For our comparative statics, we need to construct a non-trivial base case, i.e., a case where the profitability of the two products is comparable enough that none of them dominates the decision making. (If the new product is a lot more profitable, $d^* = d^{max}$, while in the reverse case $d^* = d^{min}$.) The graphs in the following sections are based on the following parameter specifications: $M = 500$, $prob(s = s_L) = 0.5$, $s_L = 0.1$, $s_H = 0.9$, $U_o = U_x = k_n = 0.005$, $\beta_n(l = 0) = 1$, $\beta_n(l = 1) = 0.1$, $\gamma_n = 0.5$, $\delta_n = 8$, $m = 0.6$, $p = 1$, $sv = 0.1$, $c_1 = 0.4$, $\alpha_c = 0.4$, $\beta_c = 0.5$, $\gamma_c = 0.2$, $\delta_c = 8$, $\alpha_l = 0.3$, $\beta_l = 1$, $\gamma_l = 0.3$, $\delta_l = 8$.

Considering only the thin solid lines in Figure 5 one can note first of all that the continued sale of the old product in stock will typically affect the design choice negatively. Only if the new product is a lot more profitable, the highest profitable redesign d^{max} is chosen. This case should, however, be rare in our setting, in which the new product is only introduced for regulatory reasons and not customer preferences. We will now compare various scenarios to explore the impact of problem parameters on the optimal design choice.

Effect of New Product Profitability

In the base case, the profit margins for the old and the new product are set identically. We explore first the effect of an increase in the profit margin of the new product relative to that of the old product. Figure 5 shows the profit impact of an increase of the new products margin m from 0.6 to 0.75 (dashed line) and 0.9 (bold line) respectively. From the Figure (and corresponding numerical solutions summarized in the table at the end of this section) we can deduce that, not surprisingly, the increase in profit margin for the new product will increase the optimal design choice. Observation 1 summarizes the effects of an increase in the new product's margin on the technology and quantity choice:

Observation 1 *If the profit margin m of the new product increases, the optimal technology choice d^* increases.*

Effect of Market Uncertainty

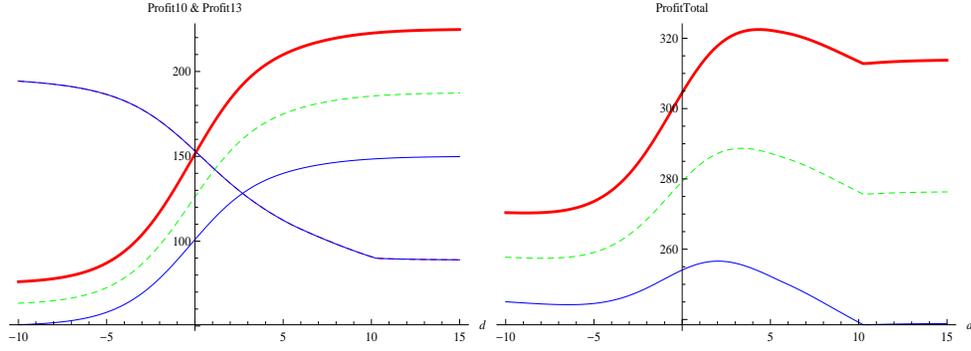


Figure 5: Old product profit (decreasing curve on left hand side), new product profits (increasing curves on left hand side) and total profits (right hand side) when $m = 0.6$ (solid line), $m = 0.75$ (dashed line) and $m = 0.9$ (bold line) respectively.

Next, we examine the impact of market demand uncertainty on design choice during the transition period. We change market uncertainty by varying gap between the low and high demand signals (s_L and s_H), while keeping the average fixed. More specifically, we decrease the uncertainty by changing (s_L, s_H) from the base case $(0.1, 0.9)$ to $(0.2, 0.8)$ to $(0.3, 0.7)$. Figure 6 show the effect of changing levels of uncertainty on the conditional demands of the old product.

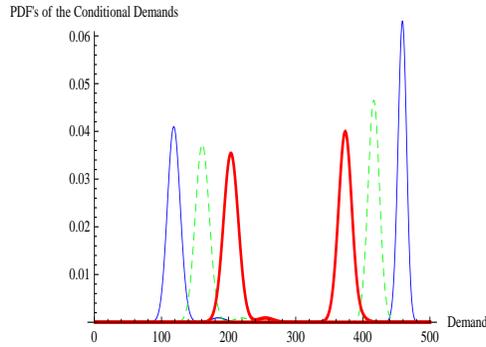


Figure 6: Conditional demands of the old product for original high uncertainty case (solid line), medium uncertainty case (dashed line) and low uncertainty case (bold line) respectively.

The left-hand side of Figure 7 shows the effect of an increase in uncertainty on the old product's profit curve. Higher uncertainty makes the old product's profits (as a function of d) more sensitive to a change in d , dropping quicker and to a lower level of profitability. Since the expected profits of the new product are only affected by the overall mean, the change in market uncertainty does

not affect the expected new product profits. The right-hand side of Figure 7 reveals the effect on total profits and hence on the choice of d . The higher the uncertainty, the lower the optimal d . The intuition is as follows: If the uncertainty is very small, the signal reveals very little information, and the firm might as well decide the total production quantity up front and avoid producing later at higher costs. Hence, choosing a high d does not affect the old product's production costs. The higher the uncertainty, the higher the firm's incentives to delay the production and hence the higher the effect of choosing a high d on the manufacturing costs of the old product. Observation 2 summarizes the effects of an increase in the uncertainty on the technology choice:

Observation 2 *If the market uncertainty increases, the optimal technology choice d^* decrease.*

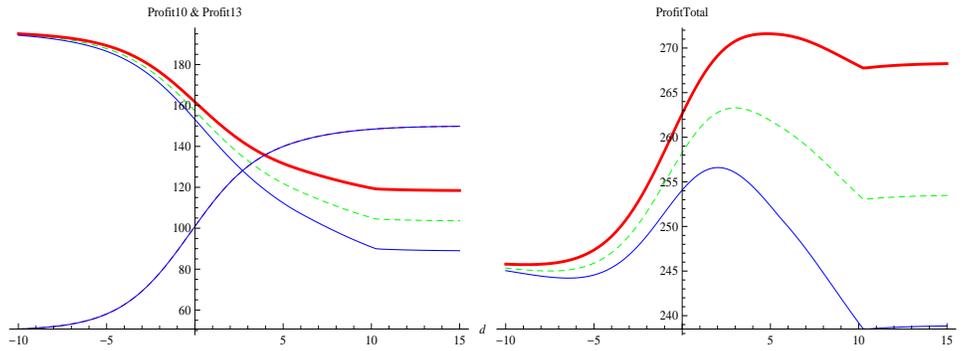


Figure 7: Effect of high (solid line), medium (dashed line) and low uncertainty (bold line) on profits.

Effect of Product Delay

Next we will explore the effect of product delays on the design choice. Here we need to differentiate between two effects: (1) the likelihood of a delay, and (2) the customers' propensity to switch products due to a delay (i.e., utility effect of a delay). The likelihood of a delay increases for any level of design choice, if the time available for the product redesign is shorter or the requirements are more stringent. Figure 8 shows the effect of an increase in the delay risk by changing the parameter α_l of the delay function from the base case of 0.3 (solid line) to 0.8 (dashed line) to 1 (bold line). As long as some customers risk switching away from the new product due to a delay (i.e., customer utility is affected by the delay), the manufacturer's profits are reduced even if the

profit margins for the two products are identical. This is due to the fact that some customers will not switch the old product but to an outside option instead. A similar figure could be generated for an increase in the customers' propensity to switch, i.e., if the parameter $\beta_n(l = 1)$ decreases. In both cases, manufacturers will typically lower their exposure to this risk by lowering the optimal design choice d . Our extensive numerical studies revealed one exception: If the new product is significantly more profitable than the old product, manufacturers might actually increase d in case of increased delay risks. The intuition is the following: If the new product is a lot more profitable, manufacturers want to ensure a large market share for the new product; hence a low d in response to a delay risk is not an option. Since, however, the risk of a delay is now already high for a medium high d , the marginal increase of this risk for even higher d can be ignored, and hence manufacturers increase d even further. Observation 3 summarizes this effect:

Observation 3 *If the delay risk increases (α_l increases or $\beta_n(l = 1)$ decreases), the optimal technology choice d^* decreases, unless the original d^* was already very large, (e.g., very profitable new product).*

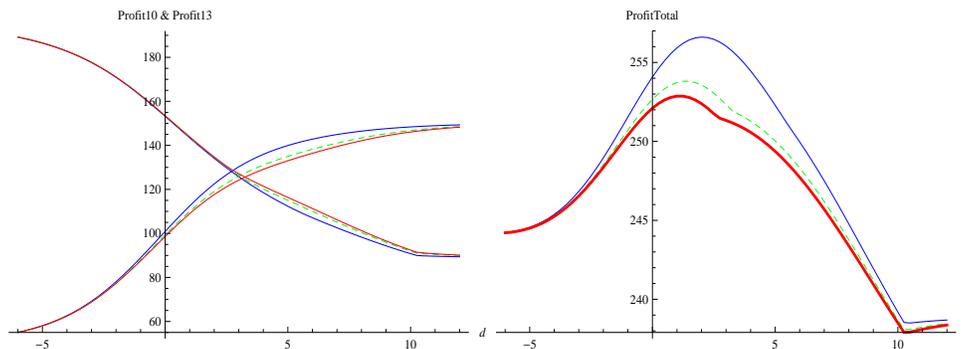


Figure 8: Impact of low (solid line), medium (dashed line) and high (bold line) probability of a delay on profits.

Effect of the Value of the Outside Option

Next we turned to the value of the outside option. One can interpret this as our products being less competitive. The outside option captures all other alternatives our customers consider in addition to the new and old product offered by us. This could include the option not to buy any

product, but can also capture the utility of all competitor products offered. Our model above does not consider competitor actions. Hence, the only situation we can consider is that of monopolistic competition, i.e., there are many competitors in the market and any competitor's individual choices has no impact on the choices made by the remaining players. That means, our focal firm will take the (expected) utility of all other options to the customer into consideration, but does not consider any strategic effect of his design choice on the other firms. In this sense, an increase in the value of the outside option can also be interpreted as the focal firm facing a more competitive market, i.e., competitors with higher design choices.

An increase in the utility of the outside option will affect the market shares of our products, especially when our design choice is low. For the new product, this seems obvious: If we choose a poor redesign, more customers will consider the outside option including competitor products. For the old product, the intuition is as follows: First, there is a fixed effect of more customers choosing the outside option over our product. As d increases, more customers are attracted by the new product. However, now more customers switch from the outside option to the new product rather than from the old product (which already has a low market share). Hence, the effect of increasing d on the market share of the old product is less pronounced, and disappears as that market share approaches the proportion of customers s that do not consider the new product. The overall effect is that firms choose more radical redesigns in more competitive settings. Note in fact in this case the change in d^* is particularly pronounced; if customers have better outside options (e.g., better competitor products to choose from) our firm will indeed be required to offer a better design itself. The following observation summarizes again the effect of an increase in the utility of the outside option (e.g., in competitiveness of the market) on the design choice of the focal firm:

Observation 4 *If the utility of the outside option U_x increases, the optimal technology choice d^* increases.*

Effect of Product Architecture

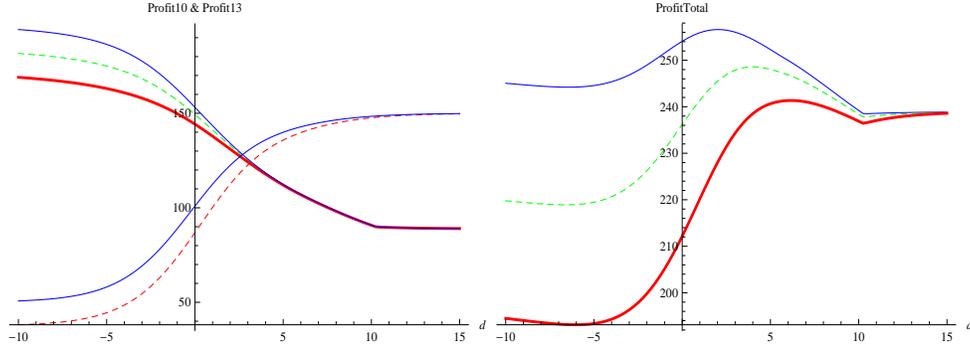


Figure 9: Impact of a low ($U_x = 0.005$, solid line), medium ($U_x = 0.01$, dashed line) and high value of the outside option ($U_x = 0.02$, bold line) on profits.

Finally, we study how the product architecture of the old product affects the optimal design decisions. As discussed in Section 3.1, a larger β_c captures a more modular design. This change yields a lower cost for producing the old product during the transition period, thus increasing the old product's profitability and d^* correspondingly (see Figure 10). There is one exception: Note the right side of the old products profit curve: If the optimal d was already very large, it is possible that a less modular old product would result in a higher optimal technology choice d^* . The intuition is as follows: With the less modular old product, the firm starts to produce everything up front at a lower level of d (no second order quantity); hence, it can focus on the new product when deciding d . Observation 5 summarizes this effect.

Observation 5 *If the old product's modularity increases (β_c), the optimal technology choice d^* increases, unless the original d^* was already very large, (e.g., very profitable new product).*

4.3 Different Market Scenarios and the Effect on Early Inventory Build-up

We ran various scenarios to confirm the generality of the above findings. Figure 11 summarizes the above results and presents three of the additional scenarios which capture different market assumptions. In particular, Scenario 2 captures a setting with a large market size of the old product (expected value of the proportion considering only the old product $E(s) = 0.7$ rather than $E(s) = 0.5$ in the base case). Scenario 3 captures a setting where customers that consider buying the new product do not consider buying the old, non-conform product ($U_o = 0$). And Scenario

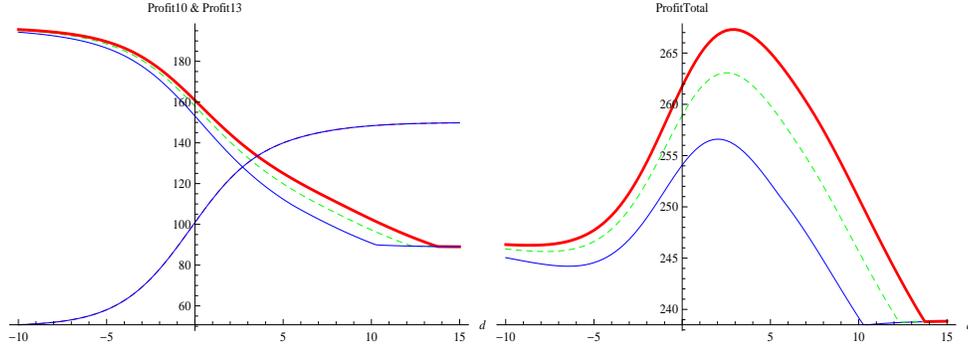


Figure 10: Impact of low modularity ($\beta_c = 0.5$, solid line), medium modularity ($\beta_c = 0.75$, dashed line) and high modularity ($\beta_c = 1$, bold line) on profits.

4 captures a setting where the old product is so profitable that a higher SL is desirable (70% as compared to 66% in the base case; note that a very large desired service level will collapse the Q_1 choice to a standard Newsvendor with no second period purchases and d being chosen based on the new product only.) While in almost all scenarios we found the same directional results for the design choice d , the four scenarios show how differences in markets can result in different effects of parameters on inventory build ups. The optimal inventory choice in all cases is summarized in the table, and we will next discuss some interesting differences.

	Scenario 1 (Base Case)			Scenario 2 ($s_1=0.5$; $E(s)=0.7$)			Scenario 3 ($U_0=0$)			Scenario 4 ($sv=0.15$, $SL=70\%$)			
	d^*	Q_1	Q_{Total}	d^*	Q_1	Q_{Total}	d^*	Q_1	Q_{Total}	d^*	Q_1	Q_{Total}	
(adj.) Base Case	1.8	97	455	2.5	279	453	2.8	57	450	1.8	100	455	
incr. new product margin	$m=0.75$	3.1	83	453	4	276	451	3.3	57	450	3.1	86	453
	$m=0.9$	4.1	77	451	5.1	275	450	3.8	58	449	4.6	112	451
decr. uncertainty	$s_1=0.2$	2.7	135	406	5.6 ($s_1=0.6$)	322	401	3.6	110	399	2.8	138	407
	$s_1=0.3$	4.6	178	353	d_{max} ($s_1=0.65$)	371	371	4.8	163	348	d_{max}	348	348
incr. delay risk	$\alpha_1=0.8$	1.1	109	456	2.0	287	455	2.5	57	450	1.1	113	457
	$\alpha_1=1$	0.9	113	457	1.9	290	455	2.4	57	450	0.9	118	457
incr. propensity to switch	$\beta_n=0.5$	2.2	92	454	2.8	276	453	2.4	57	450	2.2	94	454
	$\beta_n=0.01$	1.5	101	455	2	282	454	2.9	57	450	1.5	103	456
incr. value of outside options	$U_x=0.01$	3.7	78	452	4.7	274	451	4.3	58	449	3.8	84	452
	$U_x=0.02$	5.9	78	450	6.7	272	449	6.1	60	449	d_{max}	449	449
incr. modularity	$\beta_c=0.75$	2.3	89	454	2.9	275	454	3.0	56	450	2.3	91	455
	$\beta_c=1$	2.7	84	454	3.2	273	454	3.3	55	451	2.6	86	455

Figure 11: Summary of selected numerical results

Note first that for Scenarios 1, 2 and 4 the design and inventory parameters (d and Q_1) tend to

go into opposite directions, while in Scenario 3 they tend to go in the same direction. The reason is that d triggers two effects on Q_1 : The first effect of d increasing is that the market share of the old product decreases (demand effect), which drives Q_1 downwards. The second effect is the indirect effect of d on Q_1 via the cost c_2 to produce the old product close to the selling season (cost effect), which pushes Q_1 upward. Whether the demand or the cost effect dominate, depends on the exact parameter choices. Whenever a large enough section of the customers considers both the old and the new product, the demand effect tends to dominate. If the customers are however mostly segmented into those considering the old and those considering the new product, the cost effect tends to dominate. Note that in Scenario 3 the increase in d^* and decrease of Q_1 in response to increased modularity is no contradiction: due to the increased modularity, d^* can increase while still resulting in a lower c_2 and hence a lower Q_1 . Finally, across scenarios a decrease in uncertainty results in higher Q_1 . This might not be intuitive, since a lower variance would decrease the Newsvendor quantity. However, lower uncertainty about demand scenarios also allows the manufacturer to buy more or all of the product early, without exposing himself to as much risk of overage. This is confirmed by looking at Q_{Total} , the expected total quantity bought, which indeed decreases with variance.

Our model suggests that unless d^* is very large, manufacturers will build most of the inventory only toward the end of the selling season. Evidence presented in the introductory example suggests that this is what happened in the air-conditioning industry, and hence our model provides an explanation for the observed late inventory build up.

5 Managerial Discussion and Conclusion

In this paper, we studied the decision making of a manufacturer who faces a change in environmental regulations that requires a production stop by a pre-specified date. US regulations, however, typically allow the continued sales of already manufactured products. Given that some customers

may prefer the less energy efficient but also less costly equipment, the manufacturer faces two decisions for a transition period: the design choice for the new product that meets the new regulatory standards, and an inventory decision for the old product version. We explore the effects of various drivers on the design choice, resulting in a number of conjectures that could be tested empirically. We summarize the results in Figure 12.

	Parameter	Effect of an increase on d^*
New product profitability	m	↑
Market uncertainty	gap of s_L vs s_H	↓
Effect of product delays	α_i or β_n	↓ (but ↑ for large d^*)
Outside Options	U_x	↑
Modularity (of product architecture)	β_c	↑ (but ↓ for large d^*)

Figure 12: Summary of Results

While this research was motivated by the specific case of a regulatory change in the air-conditioning industry, we believe the insights are useful for similar regulatory changes being imposed on a number of industries, as for example the minimum energy efficiency regulation for refrigerators. Furthermore, the results provide not only direct recommendations for the optimal design choice of manufacturers, but they also have implications for regulators. First, it suggests the importance of providing sufficient time for implementing a change in the regulation. A short time horizon increases the risk of delays in introducing new products, pushing manufacturers to make minor changes during the transition period rather than going for a thorough redesign. Second, it has implications for attempts by regulators to influence customer preferences for higher energy efficiency equipment. For example, the DOE provided customers with information about long term savings with higher energy efficiency equipment, and several public utility companies offered financial incentives for customers buying the higher energy efficiency equipment. However, much of this was put in place or confirmed only towards the end of 2005, when both the SEER 13 design and the inventory levels

for SEER 10 were already decided. Our results suggest that a reduction of the uncertainty (e.g., by forcing counties to commit early whether or not to allow SEER 10 installations), could have resulted in a better SEER 13 design for the transition period.

An important limitation of our research is the recourse to numerical analysis. While we could show some basic properties analytically, the overall problem cannot be solved in close form. Our numerical analysis suggests, however, that the original profit functions could be represented by logistics functions, and that the effects of the various influence factors on the profit curves could be captured by changing the logistic function parameters accordingly. Let $\Pi_{old} = k_{old} + \frac{\alpha_{old}}{1 + \beta_{old}e^{\gamma_{old}d}}$ represent the old product's expected profit curve, and $\Pi_{new} = k_{new} + \frac{\alpha_{new}}{1 + \beta_{new}e^{\gamma_{new}d}}$ the new product's expected profit curve. Under certain conditions, we can then provide a simple closed form expression for the optimal design of the new product $d^* = \frac{1}{\gamma} \ln \left[\frac{b\beta_{old} - \sqrt{ab}}{\beta_{old}\sqrt{ab} - 1} \right]$.⁵

Using this approximation in combination with the above observed shifts of the profit curves in response to parameter changes, we can indeed confirm analytically the above numerical findings. More importantly, the graphs of the logistics curve provide a visual tool as to how changes in the expected profit curves drive the optimal design choice. We believe that this can then serve as the basis to discuss managerial or regulatory changes in a more intuitive manner, for example in a classroom setting.

6 References

- Baldwin, C.Y., K.B. Clark. 2000. *Design Rules: The Power of Modularity*. Cambridge, MA: The MIT Press.
- Banerjee, S., M. Sarvary. 2009. How incumbent firms foster consumer expectations, delay launch but still win the markets for next generation products. *Journal of Quantitative Marketing and Economics* 7(4) 445-481.
- Bhaskaran, S.R., A. Goel, K. Ramachandran. 2010. Product Transition under Development Uncertainty: Managing End-of-Life Inventory and Launch of New Products. *working paper*.

⁵The expression is obtained setting $a = \frac{\alpha_{new}}{\alpha_{old}}$ and $b = \frac{\beta_{new}}{\beta_{old}}$, under the condition that $\gamma_{old} = -\gamma_{new} = \gamma$ and $\beta_{new} < \frac{\min\{\frac{1}{a}, a\}}{\beta_{old}}$. The latter condition is only required to ensure the non-trivial case of a (unique) interior maximum; otherwise, the firm profits are maximized at the extreme points, i.e., $d^* \in \{d_{min}, d_{max}\}$.

- Billington, C., H. L. Lee, C. S. Tang. 1998. Successful strategies for product rollovers. *Sloan Management Review*. Spring. 23-30.
- Cachon, G.P., C. Terwiesch, Y. Xu. 2008. On the Effects of Consumer Search and Firm Entry in a Multiproduct Competitive Market. *Marketing Science* 27(3) 461-473.
- Cohen, M.A., J. Eliashberg, T. Ho. 1996. New Product Development: The Performance and Time-to-Market Tradeoff. *Management Science* 42(2) 173-186.
- Corbett, C. J., P. R. Kleindorfer. 2001a. Introduction to the special issue on environmental management and operations (Part 1: Manufacturing and Eco-Logistics). *Production and Operations Management* 10(2) 107-111.
- Corbett, C. J., P. R. Kleindorfer. 2001b. Introduction to the special issue on environmental management and operations (Part 2: Integrating Operations and Environmental Management Systems). *Production and Operations Management* 10(3) 225-227.
- Corbett, C.J., P.R. Kleindorfer. 2003. Environmental management and operations management: introduction to the third special issue. *Production and Operations Management* 12 (3) 287-289.
- Corbett, C.J., P.R. Kleindorfer, L.N. Van Wassenhove. 2008. Call for Papers Special Issue of Production and Operations Management: Measuring the Impact of Sustainable Operations. *Production and Operations Management* 17(6) 653.
- Corstjens, M.L., D.A. Gautschi 1983. Formal choice models in marketing. *Marketing Science* 2(1) 19-56.
- Donohue, K.L. 2000. Efficient Supply Contracts for Fashion Goods with Forecast Updating and Two Production Modes. *Management Science* 46(11) 1397-1411.
- El Khoury, H., C. van Delft. 2011. Optimal Strategy for Stochastic Product Rollover. *HEC working paper*.
- Greenley, G.E., B.L. Bayus. 1994. A Comparative Study of Product Launch and Elimination Decisions in UK and US Companies. *European Journal of Marketing* 28(2) 5-29.
- Gurnani, H., C. S. Tang. 1999. Optimal ordering decisions with uncertain cost and demand forecast updating. *Management Science* 45(10) 1456-1462.
- Ho, T., S. Savin, C. Terwiesch. 2002. Managing Demand and Sales Dynamics in New Product Diffusion Under Supply Constraint. *Management Science* 48(2) 187-206.
- Iyer, A.V., V. Deshpande, A. Mishra. 2008. Value of Operational Flexibility During Product Upgrade in Closed Loop Supply Chains. *working paper*.
- Journal of Cleaner Production. 2008. Editorial: Sustainability and supply chain management - An introduction to the special issue. *Journal of Cleaner Production* 16, 1545-1551.
- Kleindorfer, P.R., K. Singhal, L.N. Van Wassenhove. 2005. Sustainable Operations Management. *Production and Operations Management* 14(4) 482-492.
- Kornish, L.J., 2001. Pricing for a Durable-Goods Monopolist under Rapid Sequential Innovation. *Management Science* 47(11) 1552-1561.
- Kumar, S., J.M. Swaminathan. 2003. Diffusion of Innovations Under Supply Constraints. *Opera-*

tions *Research* 51(6) 866-879.

Lee, H., C. S. Tang. 1997. Modeling the costs and benefits of delayed product differentiation. *Management Science* 43(1) 40-3.

Lim, W. S., C. S. Tang. 2006. Optimal product rollover strategies. *European Journal of Opns. Res.* 174, 905-22.

McGrory, L.V.W., M. McNamara, M. Suozzo. 2000. Residential Market Transformation: National and Regional Indicators. *Conference Paper, Proceedings: 2000 ACEEE Summer Study on Energy Efficiency in Buildings*, August 20-25, 2000-06-01.

McFadden, D. 1986. The choice theory approach to market research. *Marketing Science* 5(4) 275-297.

Plambeck and Wang (2009) Effects of E-Waste Regulation on New Product Introduction. *Management Science* 55(3) 333-347.

Robertson, T.S., H. Gatignon. 1986. Competitive Effects on Technology Diffusion. *The Journal of Marketing* 50(3) 1-12.

Saunders, J., D. Jobber. 1994. Product Replacement: Strategies for Simultaneous Product Deletion and Launch. *Journal of Product Innovations Management* 11, 433-450.

Ulrich, K. 1995. The role of product architecture in the manufacturing firm. *Research Policy* 24, 419-441.

Wassenaar, H.J., W. Chen, J. Cheng, A. Sudjianto. 2005. Enhancing discrete choice demand modeling for decision-based design. *ASME Journal of Mechanical Design* 127(4) 514-523.

7 Appendix

Proof of Proposition 1: Given any signal s , we have the following profit as a function of Q_1 if $Q_1 \geq Q^*(s)$:

$$E[\Pi_O[Q_1|s]] = p \int_0^{Q_1} \xi f_{D|s}(\xi) d\xi + sv \int_0^{Q_1} (Q_1 - \xi) f_{D|s}(\xi) d\xi + p(1 - F_{D|s}(Q_1))Q_1 - c_1Q_1.$$

Thus, we have:

$$\partial_{Q_1} E[\Pi_O[Q_1|s]] = p - c_1 - (p - sv)F_{D|s}(Q_1)$$

Note that if $Q_1 < Q^*(s)$, then we have $\partial_{Q_1} E[\Pi_O[Q_1|s]] = (c_2 - c_1)$. Thus, the second derivative has a value zero or $(sv - p)f_{D|s}(Q_1)$ when $Q_1 < Q^*(s)$ and $Q_1 > Q^*(s)$ respectively (for the first case it is zero because Q_1 is replaced with $Q^*(s)$ which is independent of Q_1 value). Notice that the derivative is continuous at $Q^*(s)$ and thus the function is concave in Q_1 . Then, since $E[\Pi_O[Q_1]]$ is a weighted average of $E[\Pi_O[Q_1|s]]$ over different realizations of s , we conclude that $E[\Pi_O[Q_1]]$ is also concave in Q_1 .

Proof of Proposition 2: Since Q_1 is unique from proposition 1, Q_1^* must satisfy the first-order condition:

$$\int_0^1 \partial_{Q_1} E[\Pi_O[Q_1|s]] h_S(s) ds = 0.$$

Now, from proposition 1's proof, we know: $\partial_{Q_1} E[\Pi_O[Q_1|s]] = p - c_1 - (p - sv)F_{D|s}(Q_1)$

But $F_{D|s}(Q_1)$ depends on the signal value. Define $s(Q_1)$ to be the smallest signal s at which the conditional service level (i.e. $F_{D|s}(Q_1)$) is equal to $\frac{p-c_2}{p-sv}$. Thus, we divide the integral above according to the signals to the left and right of $s(Q_1)$ as shown below:

$$\int_0^{s(Q_1)} h_S(s) (-c_1 + p - (p - sv)F_{D|s}(Q_1)) ds + \int_{s(Q_1)}^1 h_S(s) (-c_1 + c_2) ds = 0.$$

Rewriting the second integral term as $-(c_1 - p) - (p - c_2)$, transferring the integral terms to the right, and dividing both sides by $(p - sv)$, we get:

$$\int_0^{s(Q_1)} h_S(s) F_{D|s}(Q_1) ds + (1 - H_S(s(Q_1))) F_{D|s \geq s(Q_1)}(Q^*(s)) = \frac{p-c_1}{p-sv}.$$

Notice that the left hand side is equal to $\int_0^1 F_{D|s}[\max[Q_1, Q^*[s]]] h_S[s] ds$ which is the expected service level given Q_1 commitment. Also the right hand side is equal to the unconditional service level $F_D(Q_1^*)$. Thus the proposition follows.

Proof of Corollary 1: First, notice that $s(Q_1)$ is nondecreasing in c_2 (this is because $\frac{p-c_2}{p-sv}$ decreases in c_2 , thus fixing Q_1 , the signal $s(Q_1)$ increases in c_2). Observing the left hand side of $\int_0^{s(Q_1)} h_S(s) F_{D|s}(Q_1) ds + (1 - H_S(s(Q_1))) F_{D|s \geq s(Q_1)}(Q^*(s)) = \frac{p-c_1}{p-sv}$ (i.e. the unconditional service level expression), let's divide the signals into three intervals which are split by the two points: $s(Q_1^*(c_2))$ and $s(Q_1^*(c'_2))$ where $c'_2 \geq c_2$. Clearly, for s values greater than $s(Q_1^*(c'_2))$, we have the conditional service level greater for the c_2 scenario (i.e. since $\frac{p-c_2}{p-sv} \geq \frac{p-c'_2}{p-sv}$). Furthermore, if $Q_1^*(c'_2) \leq Q_1^*(c_2)$ (i.e. Q_1^* is not increased for the case of c'_2) then clearly the conditional service level is smaller for the c'_2 case for the remaining two intervals (i.e. $s \leq s(Q_1^*(c'_2))$) since Q_2 always = 0 for such s values under c'_2 's scenario. Thus Q_1^* has to increase in c_2 to maintain an unconditional service level of $\frac{p-c_1}{p-sv}$ (See figure 4).

Proof of Lemma 1: Here we can either use the previous corollary or the following logic. Imagine that for a given c_2 and $Q_1^*(c_2)$, c_2 is reduced to a lower value c but the $Q_1^*(c_2)$ commitment is not changed. It should be apparent that the expected revenue is unchanged since changing c_2 does not affect the underlying demand. But the total cost is clearly reduced. Since by definition $Q_1^*(c)$ optimizes the problem for c , it never results in a profit worse than $Q_1^*(c_2)$'s.