Improving Supply Chain Performance and Managing Risk Under Weather-Related Demand Uncertainty

Frank Y. Chen
Department of System Engineering and Engineering Management
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong

Candace A. Yano
Haas School of Business and
Department of Industrial Engineering and Operations Research
University of California
Berkeley, CA 94720 USA

(Unabridged Version)

Abstract

This paper considers a manufacturer-retailer supply chain for a seasonal product whose demand is weather-sensitive. The retailer orders from the manufacturer (supplier) prior to the selling season and then sells to the market. We examine how a manufacturer can structure a weather-linked rebate to improve his expected profit. The proposed class of rebate contracts offers several advantages over many other contract structures, including no required verification of leftover inventory and/or markdown amounts, and no adverse effect on sales effort by the retailer.

We provide a thorough analysis of the manufacturer's and retailer's decisions in this context. We show that the weather-linked rebate can take many different forms, and this flexibility allows the supplier to design contracts that are Pareto-improving and/or limit his risk in offering the contract and the retailer's risk in accepting it. For weather rebates with certain characteristics, the manufacturer can fully hedge his risks of offering a weather rebate by paying a risk premium; we show how this can be accomplished. We also show that the basic structural results extend to more general settings, including situations where the two parties are concerned about risk and where they can set prices.

Keywords: Weather-linked rebate, weather risk, weather derivatives supply chain coordination

1 Introduction

Weather represents an important determinant of demand for many products. According to an estimate by the U.S. National Research Council, 46% of U.S. GDP is affected by weather. In the retail sector, Wal-Mart Stores Inc. reported in June 2005 that its inventory levels were higher than normal for the second straight quarter as below-normal temperatures crimped demand (Timberlake and Wiles 2005). The unfavorable weather conditions hurt demand not only in North America, but also in Europe. Cadbury Schweppes' beverage business was hit by cold summer weather in 2004, forcing the firm to lower its profit expectations. The company said the poor sales were "in line with the industry as a whole where cold and wet weather in 2004 was compared with record summer temperatures in 2003" (www.cadburyschwepps.com/EN/MediaCentre/PressReleases/sept4_market_update.htm). Coca-Cola and Unilever also blamed the weather for low sales of soft drink and ice cream products and issued profit warnings, and Nestle attributed its missing the half-year targets to the impact of poor weather on demand for ice cream and bottled water (Kleiderman 2004). USA Today (O'Donnell 2007) reported that warm weather in December 2006 in the northeastern part of the U.S. caused a dramatic fall-off in the demand for coats and sweaters, and subsequently, a cold April (2007) had the same effect on springtime clothing.

These examples are anecdotal evidence of the effect of weather on demand, but the overall impacts are pervasive. Niemira (2005) argues that weather influences sales primarily through its effect on economic activity. Broader and more systematic studies (e.g., Starr-McCluer 2000) document a significant impact of weather on retail sales at an aggregate level, although the primary effect may be that of shifting demand earlier or later. When the product of concern is a seasonal product, however, shifts in timing of aggregate retail demand translate into shifts of demand from one (type of) product to another (type of) product.

With this as a backdrop, we explore weather-linked rebates. Our modeling framework is motivated more specifically by the following situations. Weatherproof Garment Company, which designs and manufactures cold-weather apparel, including outerwear, was concerned that unseasonably warm fall weather would crimp demand for its products. The company purchased a weather derivative brokered by Storm Exchange, a firm that provides weather-related financial hedging solutions. (We discuss several types of weather derivatives and contracts later in this section.) The weather derivative would provide up to $10 million in coverage if weather in December (2007) turned out to be unseasonably warm. The CEO of Weatherproof, Eliot Peyser, indicated that the derivative would enable the company to offer rebate incentives to its customers for placing early orders (*Business Wire*, December 3, 2007). In reaction to a similar situation a European clothing manufacturer tried to encourage retailers to buy its winter collection early by offering a rebate if mild weather prevailed, and hedged its risk by purchasing a weather-linked contract (www.environmental-finance.com/2004/0403mar/hedge.htm).

With these motivating examples in mind, we take the vantage point of a manufacturer that, for a variety of reasons including capacity limitations and the concomitant long production lead times that are so common in the apparel and other industries, wishes to offer a weather rebate to retailers to encourage them to purchase (or otherwise commit to) a large quantity well in advance of the selling season. The manufacturer recognizes the difficult-to-quantify yet very significant side-benefits that the weather rebate offers (e.g., no auditing of leftover inventory at the retailer that would be required in the case of buy-back contracts or markdown allowances) and needs to structure the contract and choose (or design, if he has that choice) a weather derivative that appropriately hedges his risk of offering the rebate.

Generally speaking, we use "risk" to refer to the probability that the party loses more than a threshold amount (perhaps zero) and we use the phrase "risk tolerance" to capture the decision-maker's preferences with respect to the threshold and loss probability. Other types of risk considerations and metrics could be considered but some lead to messy mathematical expressions that obscure the main insights, so we have chosen to keep things simple. We also use the term "weather risk" to refer to the probability of low demand caused by unfavorable weather, which in turn results in an undesirable financial outcome (profit less than a threshold).

More specifically, we study a manufacturer-retailer supply chain for a seasonal product with weather-sensitive demand in a *newsvendor* context. The retailer chooses the order quantity and may take advantage of a manufacturer-offered weather rebate contract. Such a rebate linked to a weather index can be offered by a manufacturer (supplier) to protect the retailer from some potential financial effects (Malinow, 2002).

It is now possible to observe a wide range of atmospheric events in real time, to measure them with great precision and predict them with a good deal of accuracy. Extensive weather databases are now available (e.g., climetrix.com and weather-warehouse.com). Available data include high, low and average daily temperature, daily rainfall and snowfall, hourly data on humidity and cloud cover, and various indices such as heating and cooling degree days; these are available for thousands of locales around the world. Other governmental and private organizations (e.g., cpc.noaa.gov and longrangeweather.com) provide long term weather forecasts. Advances in information management and quantitative analysis tools facilitate the use of weather data for improved business decision-making (Dutton, 2002). Regnier (2008) describes recent advances in weather forecasting and applications of operations research models that utilize this information for improved decision-making. The vast majority of these applications are for short term decision-making, such as adjusting airline schedules in response to weather events. Improved weather prediction allows retail firms to make medium- and short-term adjustments in decisions which may include changes in order quantities or prices. Sophisticated firms such as Fedex, UPS and various agriculture and energy companies now more commonly employ meteorologists to improve their ability to forecast and to use those forecasts in making business decisions (Lustgarten 2005).

Not only can firms avail themselves of better weather information, but they can also use weather-risk-management products to reduce profit fluctuations caused by the weather. Weather derivatives were introduced about a decade ago to enable firms to hedge weather risks. Sophisticated firms now use weather contracts or derivatives, such as options and futures, or combinations of both, to hedge against the financial impact of adverse weather and to smooth out their weather-sensitive earnings. Companies using such weather-risk-management products represent an array of sectors, including electric utilities, natural gas, propane/heating oil, construction, agriculture, food/beverage, restaurants/hospitality, retailing, outdoor entertainment, transportation, manufacturing, and banking/insurance (Malinow, 2002). Weather risk derivatives and contracts are traded both on exchanges and over the counter. An active market exists on the Chicago Mercantile Exchange (CME), and the volume of transacted weather-hedge derivatives reached \$32 billion in 2007 (Davis 2008). (The volume declined in 2008 due to the global economic crisis but has started to recover.) Payouts are determined by weather records in a specific location for a given time period. Consequently, weather derivatives differ from weather insurance which requires proof of loss and usually includes additional conditions specifying when the insurer is (or is not) liable, deductibles, and so forth. Furthermore, weather insurance, once purchased, is not cancellable whereas a weather derivative can be sold at market value. (Some firms offer what they call weather insurance, but the vast majority are simply customized weather derivative contracts.)

Until fairly recently, the vast majority of weather derivatives were based on weather indices such as heating degree days or cooling degree days, and firms in the energy industry were key traders in these markets, as these indices are highly correlated with energy demands. However, within the past two years, derivatives have become available for rainfall (Colin 2008; http://www.rainprotection.net/about_us) and for other types of weather-related indices that are somewhat more customized for the needs of other industries (Wood 2007). Interestingly, even where customization has existed, the targeted industries have been primarily construction, agriculture, and other similar industries. Only recently have firms that offer weather derivatives targeted retail firms, despite the fact that apparel specialty stores cannot easily diversify product offerings to reduce their exposure to weather risk (O'Donnell 2007). Furthermore, despite the exponential growth in the market for weather derivatives and the use of more customized weather-risk-management products, details are often not made public because firms are reluctant to expose their points of vulnerability (Lustgarten 2005).

Retail firms now have access to weather derivatives based on indices that are more highly correlated with their demand than are traditional indices. These firms also can purchase highly customized weather contracts such as those available from Weatherbill.com (cf. http://www.weatherbill.com). A retailer's purchase of such a weather derivative or contract will mitigate his risk, but supply chain coordination is possible only if the manufacturer offers the contract; otherwise the effect of double marginalization remains. Of course, the retailer's purchase of a third-party derivative may reduce his need for a manufacturer-offered rebate (discussed in Section 3.1).

Although large retailers may be able to utilize (traded) weather derivatives, the typical small to medium-sized specialty retailer often lacks the financial prowess to do so. Moreover, procurement managers can opt for a weather rebate offered by a manufacturer but rarely have the authority to purchase weather derivatives.

To be concrete, we consider a scenario in which higher average seasonal temperatures lead to lower demand. As one example, studies by Storm Exchange, a weather-related risk manager, show that for every two degree (Fahrenheit) increase in the average temperature in September, sales at apparel specialty stores fall by 1% (Blumenthal 2007). Statistical estimates of the sensitivity of demand to weather are becoming more widely available. Weather Trends International has reported on the sensitivity of demand to temperature for over a dozen product categories. As an example, the demand for beer increases by 1.2% for each degree increase in temperature (http://www.wxtrends.com/content/images/ds2images/OneDegreeChange.jpg).

There are, of course, other examples for which demand is increasing or non-monotonic in the temperature (or other weather metric). Non-monotonic relationships typically arise when demand is either high or low for moderate temperature ranges and the reverse for extreme temperatures. Two examples are batteries, whose demands tend to be higher in extreme temperatures (cf. Shearer 1998), and soft drinks, whose demand is highest when the weather is warm but not exceedingly hot (see MSI Guaranteed Weather, undated). Our model allows for such relationships as well.

In the extreme case where demand and the weather index have a one-to-one correspondence, the contract that we propose is essentially contingent on the demand realization. Specifically, if the demand is below a specified level, then the rebate scheme is activated. Although researchers and practitioners have designed many contracts with the goal of increasing the supply chain's total profit, to the best of our knowledge, such a manufacturer-offered contract has not been studied in the literature. Several other types of contracts (see next section) can induce the retailer to order more than he would under a wholesale price contract by giving him some downside protection: for example, if demand is less than what he orders, the retailer receives a partial refund. However, the weather-linked contract departs from these standard supply-chain contracts, as it is based on the weather index rather than actual demand or leftovers.

We study scenarios in which the manufacturer provides the retailer an incentive to purchase more than he might otherwise at a point in time long before the selling season by offering a rebate if the actual seasonal average temperature is higher than a pre-determined threshold, with the rebate amount increasing in the deviation of the average temperature above the threshold. Such incentives can coordinate the supply chain by encouraging the retailer to order more, and may also make the manufacturer more competitive among risk-averse retailers. Manufacturers of products whose demands are highly weather–sensitive often sell in geographically distributed markets whose weather patterns are not highly correlated, and may offer products whose demands are countercyclical to one another. Thus, these manufacturers may be better able to bear some of the revenue uncertainty associated with weather-induced demand variability than, say, a specialty apparel retailer that cannot diversify easily. Large, diversified retailers can reduce the variability of their profits through product assortment choices. Yet, many products sold by these retailers have demands that are influenced by weather over short selling seasons, so these retailers can nevertheless benefit from weather-linked rebates.

It is clear that sophisticated manufacturing firms are now beginning to offer their customers weather-related contracts and are finding that they provide a competitive advantage. Meanwhile, financial institutions such as the CME are expanding the variety of weather derivatives being offered. Indeed, as we write this, the CME has just announced that it will offer snowfall contracts beginning in December 2009 (see http://online.wsj.com/article/SB125675206666013683.html). As the breadth of both standard (e.g., publicly-traded) and customized weather derivative contracts expands, manufacturers can more readily hedge their risk of offering weather-related contracts. These market trends beg the question of how manufacturing firms should structure weather-linked contracts. This represents the broader context of our study.

The remainder of the paper is organized as follows. Section 2 provides background on related research. In Section 3, we present a basic model in which prices are exogenous. The manufacturer decides the structure of the weather rebate and the retailer chooses the order quantity. We derive the structure of the supply-chain-coordinating weather rebate and show that it admits a variety of functional forms and allows for a good deal of flexibility in allocating profits and risks between the two parties while ensuring incentive compatibility (vis-a-vis the no-rebate scenario). Although prices could be decided in practice, this model is useful because it provides insight into the structure of the weather rebate unfettered by the algebraic complications of optimizing prices. In Appendix D, we extend the results to allow price setting by both parties.

The model also provides the basis for our exploration of how to take advantage of the contract's flexibility to design Pareto-improving, risk-free rebate structures. In Section 4, we show how the manufacturer can limit his risk via the choice of contract parameters, and also show that for some classes of contract structures, it is possible for the manufacturer to completely hedge his risk of offering a rebate by paying a risk premium for a weather derivative of a form that is commonly traded in practice. In Section 5, we discuss how the retailer's risk can be limited. Section 6 concludes the paper.

2 Literature Review

There are many ways in which retailers can mitigate the effect of demand uncertainty (caused by weather and other factors) on their overall profit. The most common mechanisms involve operational hedging, which can be achieved via choice of product assortment (cf. Devinney and Stewart 1988), accurate and/or quick response using more flexible production (or subcontractor) capacity (cf. Fisher and Raman 1996, Iyer and Bergen 1997), delayed product differentiation (Lee and Tang 1997), resource diversification and sharing (Van Mieghem 2007b), logistics technology such as electronic data interchange to support quick response, and the usual in-season and end-of season markdowns, among others. For a more comprehensive discussion of operational hedging, see Boyabatli and Toktay (2004).

When demand is uncertain, double marginalization leads to a loss of efficiency in supply chains. Supply contracts constitute a class of mechanisms to mitigate this efficiency loss. There is an extensive literature on supply contracts in a newsvendor context, much of it focusing on contract structures that coordinate the supply chain or improve performance vis-a-vis a scenario with decentralized decisions. In the interest of brevity, we refer the reader to surveys by Anupindi and Bassok (1999), Lariviere (1999), Corbett and Tang (1999), Tsay et al. (1999) and Cachon (2003).

A weather rebate is an alternative to other supply contracts that manufacturers might use to induce retailers not just to order greater (and perhaps coordinating) quantities, but also to order them well in advance of the selling season. Such inducements fall into two broad categories: (1) early-season incentives that reduce the retailer's financial obligation for any given purchase or commitment level and (2) end-of-season concessions paid by manufacturers when demand is weak. Early-season concessions in the form of advance purchase discounts are discussed by Gilbert and Ballou (1999), Cachon (2004) and McCardle et al. (2004), among others. Options contracts (see Barnes-Schuster et al. 2002 and Martinez-de-Albeniz and Simchi-Levi 2005 and references therein) also reduce the retailer's up-front obligation by requiring payment for only options, and not the full price, up front. The weather rebate falls into the second category because it is a type of end-of-season concession. Such concessions commonly come in the form of buy backs and markdown agreements (or "markdown money"). Early-season incentives include advance purchase discounts and reservation contracts. We briefly discuss each in turn.

Under buy-backs, the retailer returns some or all of the excess inventory to the manufacturer for a full or partial refund. Padmanabhan and Png (1995) provide a brief history of returns arrangements and discuss their advantages and disadvantages. Among the advantages are mitigating the retailer's risk, safeguarding the brand from deterioration of its image due to stale and/or discounted product, and facilitating collection of more accurate demand data. The disadvantages of such a policy include logistics costs and lessened retailer incentive to sell the product. The literature on buy-back contracts includes Pasternack (1985), Marvel and Peck (1995), Lau and Lau (1995), Kandel (1996), Emmons and Gilbert (1998), Lariviere (1999), Webster and Weng (2000), Taylor (2001, 2002), Glenn (2004), Krishnan et al. (2004) and Granot and Yin (2005).

Under markdown agreements, goods are not returned to the manufacturer, but the manufacturer fully or partially compensates the retailer for the lost margin on the inventory that must be discounted. Edelson (2005) provides a history of the use of markdown money and discusses how it affects incentives. Gottlieb (2005) argues that markdown money has disadvantages for both parties, such as blunting the incentive of retailers to forecast accurately and increasing risk and strains between manufacturers and customers. Tsay (2001) provides a comprehensive analysis of contracts with markdown provisions. Markdown and buy-back arrangements can, in principle, be made equivalent from a financial standpoint. However, markdown agreements usually protect the retailer's margins, and the retailer, not the manufacturer, has responsibility for the disposition of inventory that cannot be sold at full price. On the other hand, if buy-backs occur, the manufacturer has the burden of disposing excess goods. Tibben-Lemke (2004) describes various secondary markets available to both parties.

Our proposed rebate scheme shares one advantage of a markdown arrangement in that goods are not returned to the manufacturer, but it differs in that it is not designed to protect retail margins. Indeed, the rebate scheme need not limit the retailer's choices of prices or markdowns. In addition, no verification of leftover inventory or the number of units sold at each price is required. Thus, once the contract is negotiated, implementation is trivial. Moreover, the retailer's incentives (e.g., to forecast accurately and to invest in sales effort) are aligned with those of the supply chain as a whole, and because no verification is required, there are no incentives for dishonest reporting. Advance purchase discounts and reservation contracts, like a weather rebate, involve no administrative effort at the end of the season. Advance purchase discounts put the risk in the hands of the retailer with some financial compensation for doing so. Reservation contracts split the quantity risk and the financial risk between the manufacturer and retailer. On the other hand, the weather rebate puts the quantity and part of the financial risk in the hands of the retailer, and the manufacturer bears only financial risk. Finally, in contrast to contracts based on sharing (of profit, revenue, etc.), weather rebates require relatively little economic information to be shared between the parties.

In our motivating example, Weatherproof purchased a weather derivative to hedge its risk associated with offering weather rebates to retailers. We are not aware of any academic literature that has considered manufacturer-offered weather rebates, whether or not they are coupled with manufacturer-purchased weather derivatives to hedge the risk. However, numerous papers in the operations management literature have considered derivatives or similar financial instruments as hedges for risks faced by either the manufacturer or the retailer. In the interest of space, we list only a few. The papers include Gaur and Seshadri (2005) who analyze how best to use a financial asset whose value is correlated with demand as a hedging instrument, Chen et al. (2007) who consider financial hedging for a single-stage, multi-period inventory model, Ding et al. (2007) who consider both operational and financial hedging against exchange rate risks; and Caldentey and Haugh (2008), who consider financial hedges to counter periodic budget constraints. Chod et al. (2006) consider both operational flexibility and financial hedging in different contexts, and explore conditions in which they are complements or substitutes. For a review of hedging mechanisms when the capacity is the main operational decision, see Van Mieghem (2007a).

The literature on weather-linked supply contracts or weather derivatives used in connection with supply contracts is quite limited. Zhou and Rudi (2007) determine how the issuer should price a financial hedging contract being offered to a newsvendor who faces demand that is partially correlated with the index on which the hedging contract is based (e.g., a weather index). Chen et al. (2007) also study a scenario in which a retailer faces weather-sensitive demand and can purchase weather derivatives. They show that the risk-averse newsvendor orders more when he purchases weather derivatives and in doing so, increases his utility. Neither of these papers considers a two-party supply chain. As we argue in Appendix C, in a retailer-supplier supply chain, the retailer's use of weather derivatives cannot lead to supply-chain coordination because they do not eliminate the effect of double marginalization. On the other hand, a properly-designed weather rebate offered by the supplier will do so, with or without the supplier's purchase of weather derivatives.

We next analyze the retailer's and manufacturer's decisions and profits when prices are fixed.

3 The Basic Model

We consider a supply chain with two firms, a manufacturer and a retailer who sells a seasonal good with uncertain demand. In this basic model, we assume that both parties are risk neutral. In Sections 4 and 5, we show how the two parties can limit their respective risks (probability of an undesirable profit outcome).

The manufacturer has unit production cost c and as Stackelberg leader, chooses the contract terms and offers a wholesale price w. The retailer decides the order quantity, q and sells at a unit price p. Both w and p are exogenous here (but in Appendix D we extend the results to allow pricing decisions by both parties and show how the ability to postpone the pricing decisions affects the system). To ease the exposition, in this section, we use a generic concave expected revenue function for the retailer, $R(q)$.

Demand for the product depends on the weather, which, for the purposes of the contract, is encapsulated in a summary statistic such as the average temperature in a particular geographic area over a specified time interval. Throughout this paper, we use temperature as an example, although the problem can be similarly formulated with other weather indices, such as precipitation, rain-days, heating- or cooling-degree days, etc. Without loss of generality, suppose that higher temperatures have an adverse impact on demand. In the contracts that we discuss, there is a threshold or "strike" temperature at which the rebate is activated.

First, we introduce our key notation.

$c =$ unit production cost

$w =$ wholesale price per unit

$p =$ retail price per unit

$v =$ salvage value per unit leftover

$q =$ order quantity (decided by the retailer)

$T =$ temperature random variable

$t =$ observed temperature

$t^* =$ the strike temperature

$d(t)$ = expected demand at temperature t

$D(t, u)$ = demand (random variable)

$f(t)$ = probability density function of T

$F(t)$ = cumulative distribution function of T.

U = random variable influencing demand (unrelated to temperature)

$g(u)$ = probability density function of random variable U

$G(u)$ = cumulative distribution function of random variable U.

To avoid trivial solutions, we assume that $v < c < w < p$. We also assume that there is no shortage cost incurred by the retailer or manufacturer, apart from the lost gross margin.

Before the retailer chooses q, the manufacturer offers a weather rebate contract of the form:

$$K(t^*, q) = \begin{cases} 0 & \text{if } t \leq t^* \\ k(t, q) > 0 & \text{if } t > t^* \end{cases} \tag{1}$$

where $k(t, q)$ is non-decreasing in both t and q. In words, if the retailer orders q units at the beginning of selling season, then under the terms of the contract, the retailer receives compensation $k(t, q)$ from the manufacturer if the realized temperature, t, turns out to be greater than t^*, the strike temperature, or nothing otherwise. We will elaborate on the structure of $k(t, q)$ later. The strike temperature is fixed throughout the season. We initially assume it is set exogenously but discuss the choice of t^* later in the paper.

With this weather rebate contract, the retailer's expected profit is

$$\pi_r(q) = R(q) - wq + E_{t>t^*} k(t, q) \tag{2}$$

where $E_{t>t^*} k(t, q) = \int_{t^*}^{\infty} k(t, q) f(t) dt$ is the expected value of the rebate payment. A common expression for $R(q)$ is $E[p \max(q, d(t, u)) + v(q - d(t, u))^+]$ (where $z^+ = \max(0, z)$), i.e., the expected revenue from sales and salvaging. The manufacturer's expected profit is

$$\pi_s(q) = (w - c)q - E_{t>t^*} k(t, q) \tag{3}$$

i.e., the gross margin per unit multiplied by the quantity sold, less the expected rebate payment, and the expected profit for the supply chain is

$$\pi_c(q) = R(q) - cq.$$

Denote by q_c the quantity that maximizes the expected supply chain profit, $\pi_c(q)$, i.e., the system-coordinating solution. We assume that $\pi_c(q)$ is continuous in q, so q_c exists (although $\pi_c(\cdot)$ may have multiple maxima).

In the absence of a rebate, the retailer chooses a quantity q_r to maximize his expected profit. Let $\pi_r^d(w) = R(q_r) - wq_r$ and $\pi_s^d(w) = (w - c)q_r$ the retailer's and manufacturer's expected profit without a rebate, respectively. (The superscript d denotes the decentralized problem.)

In the remainder of this section, we first show how to structure a weather-linked rebate to achieve supply chain coordination when demand is monotonic in the weather index. In Appendix A1 we show that these results extend to situations in which demand is non-monotonic in the index. We subsequently derive more specific contract terms when the manufacturer wishes to achieve Pareto improvement.

3.1 Supply Chain Coordination

In this subsection, we investigate the possibility of achieving supply chain coordination with weather rebate contracts. Although a very wide range of contract structures can achieve the same result, our presentation here focuses on contracts with simple structures and for which the impact of the parameters is quite clear. Later in this section, we discuss specific contract structures that will coordinate the supply chain.

Theorem 1 shows that a class of weather rebate contracts can coordinate the supply chain. The structure of this class of rebates bears some similarity to the (target) sales rebates studied by Taylor (2002), but here the target specifies the minimum purchase quantity (which we call Λ) that qualifies the retailer to participate in the rebate program whereas in Taylor's model, the target is the minimum retail sales quantity above which the retailer earns a rebate. The class of rebates that we explore is extremely broad; we elaborate on this point later.

Theorem 1 *Consider a class of weather rebate contracts with the following form:*

$$\int_{t^*}^{\infty} k(t,q)f(t)dt = (w-c)(q-\Lambda)^+ \tag{4}$$

where the value of $\Lambda \leq q_c$ is prespecified and q_c is the optimal order quantity for the centralized supply chain. With such a contract, the retailer's expected profit is

$$\pi_r(q) = R(q) - cq - (w-c)\Lambda \tag{5}$$

and the manufacturer's expected profit is

$$\pi_s(q) = (w-c)\Lambda. \tag{6}$$

Moreover, the retailer chooses a value of q that maximizes $R(q) - cq$ (i.e., $q = q_c$) so contracts with this structure coordinate the supply chain.

Proof. Substituting equation (4) into (2) and (3) yields the desired expressions for $\pi_r(q)$ and $\pi_s(q)$, as asserted. □

Observe that the controllable portion of the retailer's profit, $R(q) - cq$, is such that the he will make his ordering decision as if his marginal per unit cost were c rather than w. Consequently, weather rebates of the form (4) eliminate the effect of double marginalization that would prevent supply chain coordination from being achieved.

The form of the coordinating contract in (4) is extremely flexible; the only requirements are that it satisfies (1) and (4). As such, both the specific form of $k(t,q)$ and the value of t^* may be selected by the manufacturer, or negotiated between the two parties provided that (4) is satisfied.

Our first example of a rebate scheme is a constant payout per unit (independent of t) to the retailer for each unit ordered in excess of Λ if the temperature metric (e.g., average temperature) exceeds t^*. That is,

$$K(t,q) = k(q-\Lambda)I_{t>t^*} \tag{7}$$

where $I_{\{\cdot\}} = 1$ (or 0) if the argument is true (or not true). This scheme has been implemented by companies in different settings. For example, several years ago, Bombardier Inc., a Canadian snowmobile manufacturer, offered an incentive that helped to protect itself against the lower sales and leftover inventory that accompany a mild winter. In the winter of 1998, the company offered buyers in the US Midwest a \$1,000 rebate on its snowmobiles if a pre-set amount of snow did not fall that season. (The pre-set amount was half the average snowfall of the past three years, and the price of its snowmobiles ranges from \$7,000 to \$9,000.) Sales increased 38% from the prior year! (Davis and Meyer 2000). In this case it appears that Bombardier set $\Lambda = 0$, i.e., it did not impose any minimum order quantity to qualify for the rebate, so the company effectively bore the risk of paying the rebates without any up-front requirements for the buyers. (Bombardier was able to hedge this risk via weather risk contracts but had to pay a risk premium to do so.)

This concept also applies to rebates that hinge upon extreme weather conditions—not aggregate weather metrics— that might occur during a short time horizon. One example would be snowfall exceeding a threshold (measured in a specific locale) during a weekend day before Christmas (cf. http://www.allweatherinsurance.com/products/income_stabilization.html). Another class of examples are weather contracts offered to golf courses where typically there is a threshold for rain for each day during the term of the contract (cf. Colin 2008). To represent this type of rebate, in (7), we simply replace $I_{\{t>t^*\}}$ by $I_\delta = 1$ if the extreme condition occurs and 0 otherwise. Such contracts are very simple and have been used for a number of years. In 2004, Martin Malinow of XL Weather & Energy was quoted as saying, "We are seeing end-users taking a more intuitive approach to the way weather impacts their business and that is reflected in the growing popularity of critical day contracts." (Environmental Finance 2004). Payouts on critical-day contracts are linked to relatively "extreme" weather on a number of days over the life of the contract, rather than average weather conditions over the entire period.

The second example is:

$$K(t, q) = k[(t - t^*)^+] \cdot (q - \Lambda)^+ \tag{8}$$

where $k(\cdot)$ is an increasing function, and Λ is a constant. Under this rebate, if the temperature exceeds t^*, the manufacturer pays the retailer a per unit rebate that depends upon the deviation of t above t^* for each unit that the retailer orders in excess of a base quantity, Λ. Unlike the rebate in (7), the per unit rebate amount depends upon t. This rebate can coordinate the chain if $k(\cdot)$

and t^* are chosen to satisfy (4). Contracts of this general form are often used in connection with heating (cooling) degree days or other aggregate weather metrics. Evolution Markets, a firm that sells customized weather derivatives, provides a case study on a derivative designed for a brewery whose demand falls in cool weather. The payout depended upon the shortfall of cooling degree days from the strike value (see http://new.evomarkets.com/pdf_documents/EvoWth_nyc_brewery.pdf). Because this derivative was offered by a third party and not by an upstream supplier, the payout was not a function of the order quantity, but the example shows that firms are considering weather contracts that depend upon the deviation of the observed temperature metric from a threshold.

From (4), we can see that the retailer makes an up-front payment of $(w-c)\Lambda$ to the manufacturer in exchange for a discount of $w-c$ per unit on his entire order quantity. Although this weather rebate contract appears to have the structure of a two-part tariff, the risks faced by the two parties are different under the two contracts. Under a two-part tariff, the retailer makes an advance purchase of the coordinating quantity and bears all of the risk. Under the weather rebate contract, the manufacturer receives a deterministic risk premium but bears an amount of risk that depends upon t^*, and thus, can be controlled via this parameter.

3.2 Pareto–Improving Rebates

Recall that the manufacturer's motive for offering a weather rebate is to induce the retailer to order more than he would othewise. We have shown in Section 3.1 that for any given w, a weather rebate of the form (4) is able to coordinate the supply chain. But some retailers may be unwilling to participate in a rebate scheme if it is accompanied by a higher wholesale price. Therefore, we first wish to determine whether it is possible to construct a coordinating weather rebate with $w = w_d$ where w_d is the manufacturer's chosen wholesale price in the absence of rebate (i.e., the "decentralized" solution). (Lariviere and Porteus 2001 has identified fairly general conditions under which the manufacturer's optimal wholesale price is unique, whereas here w_d need not be optimal.

If $w = w_d$, the retailer is better off *in expectation* under any rebate structure. If he does not like the terms of the rebate, he can simply order his decentralized order quantity q_r. Thus, if he agrees to the terms of the rebate and orders more than q_r, he is assuming additional inventory risk entirely of his own volition. On the other hand, if Λ is too large, then the retailer could be worse off, both in expectation and for a subset of demand outcomes, if he accepts this contract.

To find a range of Λ such that a Pareto-improving solution is obtained under the rebate scheme defined by (1) and (4), we consider two special cases with $\Lambda = q_r$ and $[\pi_c(q_c) - \pi_r^d(w)]/(w-c)$, respectively, where q_r and $\pi_r^d(w)$, as defined earlier, are the retailer's optimal order quantity and expected profit, respectively, when the manufacturer sets $w = w_d$ and does not offer a rebate.

When $\Lambda = q_r$, the retailer chooses $q = q_c$, so the manufacturer's profit remains equal to $\pi_s^d = (w-c)q_r$, while the retailer's profit improves, because

$$\begin{aligned} \pi_r(q_c) &= R(q_c) - cq_c - (w-c)q_r \\ &\geq R(q_r) - cq_r - (w-c)q_r = \pi_r^d(w). \end{aligned} \tag{9}$$

This means that the retailer gains all the incremental channel profit. A strict inequality holds when $w > c$ in (9). Also note that for all $\Lambda < q_r$, the manufacturer is worse off (in expectation) under the rebate scheme. Therefore, $\underline{\Lambda} = q_r$ is a lower bound on Pareto-improving values of Λ.

Now consider $\Lambda = [\pi_c(q_c) - \pi_r^d(w)]/(w-c)$. Note that $\pi_c(q_c) - \pi_r^d(w) \geq \pi_c(q_c) - \pi_c(q_c) + (w-c)q_r$. Thus, $q_r \leq (\pi_c(q_c) - \pi_r^d(w))/(w-c)$. In this case, $\pi_r(q_c) = \pi_r^d(w)$ and $\pi_s(q_c) = \pi_c(q_c) - \pi_r^d(w)$, meaning that the manufacturer takes all the incremental profit while the retailer's profit remains the same as that without a rebate scheme. For all

$$\Lambda > [\pi_c(q_c) - \pi_r^d(w)]/(w-c),$$

the retailer is worse off accepting the rebate scheme. Thus, this value represents an upper bound on the Pareto improving Λ. Let $\bar{\Lambda}$ denote this upper bound, which is clearly capped by q_c. Note that because $\pi_c(q_c - \pi_r^d(w) \geq (w-c)q_r$, we have that $\underline{\Lambda} \leq \bar{\Lambda}$.

To summarize, for both parties to be better off with the introduction of the rebate scheme defined by (4), the value of Λ must be set within the range $(\underline{\Lambda}, \bar{\Lambda})$.

We now investigate whether a weather rebate is guaranteed to make both parties better off if it is accompanied by a wholesale price higher than w_d. Let $\Lambda_w = (w_d - c)q_r/(w-c)$. We have the following result, the proof of which appears in Appendix A2.

Theorem 2 *A rebate that induces $q \geq \Lambda$ is Pareto-improving (in expectation) for both the manufacturer and retailer if for $w \geq w_d$, condition (4) holds and $\Lambda \in [\Lambda_w, \bar{\Lambda}]$, where w is the wholesale* *price that accompanies the rebate.*

Note that $\Lambda_w \leq q_r = \underline{\Lambda}$. Therefore $[\Lambda_w, \bar{\Lambda}]$ is non-empty.

4 Limiting the Manufacturer's Risk

In this section, we discuss two ways in which the manufacturer may limit his risk: (i) via his choice of parameters for the coordinating contract described in the previous section and (ii) via the purchase of a weather derivative.

4.1 Limiting Risk via the Choice of t^* and $K(t^*, q)$

In the previous section, we assumed that the value of t^* was set exogenously. It is, however, one feature of the weather rebate that makes it distinctive. If the manufacturer sets t^* to a very high value, then he will pay the rebate infrequently but the payout for each such instance will be large. On the other hand, if the manufacturer sets t^* to a low value, the structure of the contract approaches that of a full-value markdown arrangement on units purchased in excess of q_r. Provided that $K(t^*, q)$ is defined so that (4) is satisfied, both the manufacturer's and retailer's expected profits remain the same for all values of t^*. Thus, t^* can be adjusted—or perhaps negotiated—according to the risk tolerances of the two parties.

Suppose the manufacturer wants to impose a constraint on the probability that he is worse off in the presence of the rebate, i.e.,

$$\text{Prob}\ [(w - c)q - K(t^*, q) \leq (w - c)q_r] \leq \gamma$$

where $\gamma \in (0, 1)$ is pre-specified. Note that without the rebate, the manufacturer earns a risk-free profit $(w-c)q_r$. Here, we use a specific form of the manufacturer's chance constraint, but $(w-c)q_r$ can be replaced by any constant value and the same types of conclusions can be drawn from the analysis.

We consider the rebate scheme $K(t^*, q) = k \cdot (t - t^*)^+ \cdot (q - \Lambda)^+$. If $q \geq \Lambda$, then the constraint can be reinterpreted as

$$t^* \geq F^{-1}(1 - \gamma) - (w - c)(q - q_r)/[k \cdot (q - \Lambda)].$$

In the special case where $\Lambda = q_r$, this becomes

$$t^* \geq F^{-1}(1-\gamma) - (w-c)/k.$$

In this case, the manufacturer only needs to set t^* so that his downside risk is limited.

We note that other types of risk constraints can be handled, albeit with more complex algebra, provided that the risk metric is monotonic in t^*. As such, metrics based on conditional value at risk (CVaR; expected shortfall from a target) may also be employed. The business press suggests that manufacturers are not necessarily trying to offset all risks, but are instead attempting to hedge against moderately bad (or worse) outcomes, as was true in the case of Weatherproof. As such, chance constraints similar to those mentioned above and constraints on CVaR can reflect these concerns, which mirror those expressed by retailers. (For further details, see the web sites of the weather risk management firms cited in the Introduction.)

Note that so long as the rebate contract satisfies (4), the supply chain will be coordinated, irrespective of t^*. However, in a non-coordinating rebate contract, the choice of t^* can have an impact on the allocation of profit between the parties. To see this effect, consider what happens when all contract parameters are held constant but the strike temperature is reduced to $t^{**} < t^*$, where t^* satisfies (4) but t^{**} does not. Then,

$$\int_{t^{**}}^{\infty} k(t,q)f(t)dt > \int_{t^*}^{\infty} k(t,q)f(t)dt = (w-c)(q-\Lambda)^+,$$

implying that the retailer will earn more than he would with the coordinated contract.

The impact of k in (8) is similar: holding all else constant, the retailer gains more in expectation with a larger k. However, with a larger k, the lower limit on t^* increases, so the retailer experiences greater variability in his profit: he will be paid larger rebate payments less frequently. (The variability of the manufacturer's profit also increases for the same reason.)

4.2 Weather Derivatives to Hedge the Risk of Rebate Offers

In this subsection, we show that a weather derivative–with appropriate characteristics–gives the manufacturer a riskless means to offer a rebate (by paying a fixed premium). So whether the manufacturer should purchase the weather derivative depends upon the amount of the premium, his own risk attitude and other means available for mitigating risk.

Here we take the average temperature as an example and consider a call option with the following characteristics. The agreed-upon strike (average) temperature is $\hat{t}$, and the manufacturer (who buys the option) pays a premium B. If the realized average temperature is greater than $\hat{t}$, then the manufacturer receives b for each unit of deviation $(t-\hat{t})^+$ and nothing otherwise, but the maximum payoff is capped by $\hat{b}$. Then, the payoff per option to the manufacturer is $\min(\hat{b}, b(t-\hat{t})^+) - B$ (here we ignore the time value of money). This is a standard weather option (see Malinow 2002). Naturally, $\hat{b} > B$ and $E_t[\min(\hat{b}, b(t-\hat{t})^+)] \leq B$, i.e., the expected payoff from the option, $\pi(K, \hat{t})$, is equal to $E_t[\min(\hat{b}, b(t-\hat{t})^+)] - B \leq 0$.

Suppose that the manufacturer can choose $t^* = \hat{t}$. By aligning the form of the option and the strike temperature in the option with those in the rebate, the manufacturer can transfer all of the weather risks to the derivative writer. This can be shown as follows.

Consider the rebate scheme $K(t,q) = k(t-t^*)^+(q-\Lambda)^+$, as given by (8). Let $k(t-t^*) = (b/L)(t-t^*)$ for $t \leq t^* + \hat{b}/b$ and $k(t-t^*) = \hat{b}/L$ for $t > t^* + \hat{b}/b$, where $L > 0$ is a parameter to be determined. By carefully choosing the value of L and the number of options, n, that the manufacturer buys, it is possible to have

$$n \int_{t^*}^{\infty} \min(b(t-t^*)^+, \hat{b}) f(t) dt = (w-c)L$$

so that the chain will be coordinated. In particular, letting $L = q_c - \Lambda$ and $K(t,q) = nk(t-t^*)^+(q-\Lambda)^+$, the chain can be coordinated.

However, with a coordinating rebate contract combined with weather call options, the manufacturer's and retailer's profits become:

$$\begin{aligned}
\pi_s^o(q_c) &= (w-c)q_c - \int_{t^*}^{\infty} k(t-t^*)(q_c-\Lambda)f(t)dt + n[\int_{-\infty}^{\infty} b(t-t^*)f(t)dt - B] \\
&= (w-c)\Lambda + n[\int_{-\infty}^{\infty} b(t-t^*)f(t)dt - B], \text{ or } (w-c)q_c - nB \qquad (10) \\
\pi_r^o(q_c) &= -wq_c + R(q_c) + n\int_{t^*}^{\infty} k(t-t^*)(q_c-\Lambda)f(t)dt \\
&= \pi_c(q_c) - (w-c)\Lambda. \qquad (11)
\end{aligned}$$

Note that hereafter we use the superscript o to represent the case with weather options.

If there is no risk premium, then $\int_{-\infty}^{\infty} b(t-t^*)f(t)dt \approx B$. Thus, $\pi_s^o(q_c) = (w-c)\Lambda$, so the manufacturer earns the same amount as his expected profit in a coordinated supply chain, and the additional risk due to the weather contract offered to the retailer is transferred to the weather derivative writer. If $\Lambda = q_r$, then $\pi_s^o(q_c) = (w-c)q_r$ and $\pi_r^o(q_c) = \pi_c(q_c) - (w-c)q_r$.

When a risk premium has to be paid for the weather options, i.e., $\int_{-\infty}^{\infty} b(t-t^*)f(t)dt < B$, then $\pi_s^o(q_c) = (w-c)q_c - nB$. The supply chain profit becomes

$$\pi_c^o(q_c) = \pi_c(q_c) - n(B - \int_{-\infty}^{\infty} b(t-t^*)f(t)dt).$$

Thus, as long as $\pi_c^o(q_c) > \pi_c(q_r)$, the options can increase the supply chain profit above what it would be without the use of options.

To illustrate these points, we present a simple example in which the product is an inexpensive article of seasonal clothing. The parameters are: $p = 15$, $w = 10$, $c = 5$, $v = 0$ (cost parameters are normalized so that $v = 0$), demand function $D(t,u) = (15000 - 100t) * u$, $T \sim N(70, 5^2)$ (for the month of September in the southern half of the U.S.; measured in Fahrenheit), and U is Uniform on $[0.5, 1.5]$. At the average temperature, expected demand is 8,000 and for each two degree increase in temperature, demand falls about 1.3%. There is great deal of inherent uncertainty in the demand, even apart from the weather, which is captured in the distribution of U. The optimal centralized order quantity is $q_c = 9,263$, while in the decentralized supply chain (without a rebate) the retailer's optimal order quantity is $q_r = 6,612$. The decentralized solution yields a total supply chain profit of \$59,466 (\$26,406 for the retailer and \$33,060 for the manufacturer), whereas the coordinated solution achieves a total profit of \$66,120, or an increase of \$6,654 for the supply chain as a whole.

Suppose the weather option is defined as

$$B = \$2,860, \; b = \$8,486 \text{ (per degree)}, \; \hat{b} = \$33,945, \text{ and } t^* = 75$$

where the strike temperature is set at one standard deviation above the mean. Now let $L = 2651 (= q_c - q_r = 9,263 - 6,612)$. Then, one way the rebate can be structured is: $k(t-t^*)^+ = 3.2 * (t - t^*)$ for $75 \leq t < 79$ and $k(t-t^*)^+ = 12.8$ for $t \geq 79$. Then, 5 options are needed (i.e., $n = 5$). Note that the risk premium for the option is

$$B - \int_{t^*}^{\infty} b(t - t^*) f(t) dt = 209 \text{ (8\% of the expected payoff).}$$

Thus, as long as the total risk premium for the five options is less than \$6,654, the supply chain profit can be improved when the manufacturer uses the weather options.

Note that the manufacturer must pay a premium up front for each option. However, as mentioned above, he can increase w so as to earn a higher risk-free profit than what is possible without the use of options and a rebate contract. In the above example, if

$$(w - c)q_c - nB > (w_d - c)q_r$$

(see (10)), the manufacturer earns a larger risk-free profit. For example, suppose $w_d = 10$ without the rebate and $w = \$10.40$ with the rebate, and all else remains the same. Then the above inequality implies:

$$(w - c)q_c - nB = \$35,720 > (w_d - c)q_r = \$33,060,$$

so the manufacturer earns \$2,660 (or 8%) more risk-free. Alternatively, the manufacturer can keep $w = w_d = 10$ but set $\Lambda = q_r + 532$. Then, he also earns \$2,660 more (after deducting the cost of the derivative). However, note that although the retailer also earns an 11% (\$2,923) higher expected profit, the standard deviation of his profit increases from 11,496 (at an expected profit of \$26,406) to 16,214 (at an expected profit of \$29,354). This raises the question of how to limit the retailer's risk, which we discuss in the next section.

5 Limiting the Retailer's Risk

Weather contracts typically require the retailer to make a financial commitment up front in exchange for risk mitigation at a later date. A retailer may be unwilling to make the additional up-front commitment if the subsequent risk mitigation does not meet his expectations. In this section, we explore one approach for limiting the risk incurred by the retailer.

We consider an extension of our basic model in which both parties seek to maximize their own expected profit and the retailer imposes a constraint specifying that the probability that his profit falls below a threshold, α, should be no greater than β. Constraints of this type are popular in the finance literature to capture bankruptcy risk, and in broader contexts to capture participation or risk constraints in stochastic settings. We show how to structure the weather rebate so that it provides (weak) Pareto improvement while the retailer's risk constraint is satisfied.

Here, we assume the manufacturer is risk neutral, but note that the risk mitigation mechanisms described in Subsection 3.1 and Section 4 may be used to limit the manufacturer's risk. For ease of exposition, we only consider a rebate of the form given in (8). Below we provide an overview of the results; Appendix B contains further details.

Without a Rebate The retailer's problem is

$$\max \pi_r^{NR}(q) = E\{p \min[q, D(t,u)] + v(q - D(t,u))^+ - wq\}, \tag{12}$$

$$s.t. \;\; \mathsf{Pr}\{\Pi_r^{NR}(q) \le \alpha\} \le \beta. \tag{13}$$

where NR denotes "no rebate," Π is the random variable representing profit and π is the expected profit. Such a constraint has been analyzed previously in the literature. As shown by Gan et al. (2005) and others, when q is too small (i.e., $q < q_\alpha = \frac{\alpha}{p-w}$), the constraint is violated; and for q above a threshold,, the probability that the retailer's profit achieves the level α decreases as q increases. Therefore, the optimal order quantity, $\hat{q}_r^{NR} = \min\{q_r, \bar{q}_r^{NR}\}$, where q_r is the unconstrained optimal order quantity in the absence of a rebate (as defined in Section 3), and $\bar{q}_r^{NR} = \max\{q : \mathsf{Pr}\{\Pi_r^{NR}(q) \le \alpha\} \le \beta\}$.

With a rebate

Here, we set $\Lambda = \hat{q}_r^{NR}$, the (possibly constrained) optimal order quantity when there is no rebate, in a manner analogous to what we did in the risk-neutral setting. The retailer's problem becomes

$$\max_q \pi_r^R(q) \;=\; E\{p \min\{q, D(t,u)\} + v(q - D(t,u))^+ - wq + k(t-t^*)^+(q - \hat{q}_r^{NR})^+\} \tag{14}$$

$$s.t. \;\; \mathsf{Pr}\{\Pi_r^R(q) \le \alpha\} \le \beta; \tag{15}$$

and the manufacturer's profit is

$$\pi_s(q) = \mathsf{E}\{(w-c)q - k(t-t^*)^+(q - \hat{q}_r^{NR})^+\}. \tag{16}$$

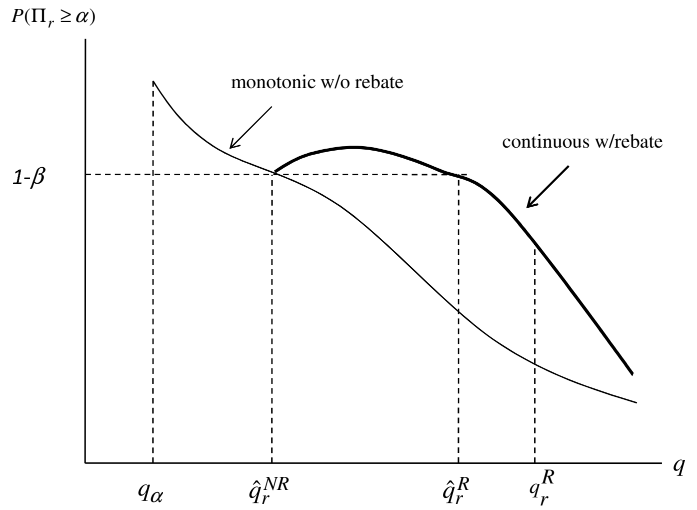

Figure 1: Monotonicity or Continuity of the Probability that the Profit Exceeds α

Figure 1 shows a stylized example of $\mathsf{Pr}\{\Pi_r \geq \alpha\}$ (with $\Lambda = \hat{q}_r^{NR}$). As noted above, in the absence of a rebate, for q values above some threshold, the probability that the profit exceeds α is monotonically non-increasing in q. In the presence of the rebate, it is not possible to show in general that the probability is monotonic in the relevant range, but we have observed that the function is relatively well behaved. In Appendix B, we show that it is continuous, which is the only condition needed for the analysis that follows.

We now turn to question of the existence of feasible solution for the retailer in the presence of a rebate. Let q_r^R denote the (unconstrained) order quantity that maximizes the retailer's expected profit when the manufacturer offers a rebate. (Note that it is identical to q_c if the rebate coordinates the supply chain in the risk-neutral setting.) It can be shown that $q_r^R > \hat{q}_r^{NR}$, because a rebate always benefits the retailer. If q_r^R satisfies the retailer's risk constraint, then we are done. Otherwise, we need to identify the best constrained solution. By the monotonicity of $\mathsf{Pr}\{\Pi_r^{NR}(q) \geq \alpha\}$ for $q \in [\hat{q}_r^{NR}, q_r^R]$, the continuity of $\mathsf{Pr}\{\Pi_r^R(q) \geq \alpha\}$ and the fact that $\mathsf{Pr}\{\Pi_r^{NR}(\hat{q}_r^{NR}) \geq \alpha\} = \mathsf{Pr}\{\Pi_r^R(\hat{q}_r^{NR}) \geq \alpha\}$ when $\Lambda = q_r^{NR}$, we can infer the existence of some q such that the risk constraint is satisfied and that the value of q lies in the interval $[\hat{q}_r^{NR}, q_r^R]$. These results imply that the upper envelope shown in the figure passes through $1 - \beta$ for at least one value of $q \in [\hat{q}_r^{NR}, q_r^R]$. Thus, because the retailer's objective under the rebate is increasing for $q \leq q_r^R$, the optimal constrained order quantity, $\hat{q}_r^R$, is equal to $\bar{q}_r^R$, the largest value of q satisfying the retailer's risk constraint in the presence of the rebate, when q_r^R does not satisfy the retailer's risk constraint.

So far, we have assumed that $\Lambda = \hat{q}_r^{NR}$. If $\Lambda > \hat{q}_r^{NR}$, it is possible that the retailer chooses not to purchase enough to make him eligible for the rebate. Referring again to Figure 1, the curve for the case with a rebate is shown for $\Lambda = \hat{q}_r^{NR}$. As Λ increases, this curve moves downward. (Note that the value of q_r^R also depends on the value of Λ. In the figure, q_r^R is the unconstrained solution under a rebate with $\Lambda = \hat{q}_r^{NR}$.) If Λ is too large, the retailer will not accept the rebate offer and his solution defaults to the same one as in the case of no rebate. This illustrates one danger of the manufacturer choosing too large a value of Λ.

Although it is easy to determine whether the retailer's risk constraint is binding in the absence of a rebate, this is more difficult to ascertain in the presence of a rebate because it depends upon Λ. Thus, to determine whether the rebate is Pareto-improving (in expectation) for the retailer, we consider all possible types of outcomes with and without a rebate, as shown in Table 1. In the table, "unconstrained" ("constrained") means that the unconstrained (constrained) solution is applicable and "greater" means weakly greater. The results in the first row of the table follow from the fact that any solution that is feasible in the absence of a rebate is also feasible when the manufacturer offers a rebate. The results in the second row follow directly from the foregoing analysis. The cell in the lower left is not applicable because it is not possible for the retailer to be constrained in the absence of a rebate and unconstrained when the rebate is offered but he does not order enough to be eligible for it. In the lower right cell, the retailer is unconstrained and cannot or does not avail himself of the rebate offer, so his solution is the same in both cases. Because q is never smaller when a rebate is offered and the retailer is never forced to order more than his unconstrained solution (with or without a rebate offer), the retailer is always at least as well off in expectation when the rebate is offered.

Now we turn to the manufacturer's profit in (16). Taking the expectation over t yields

$$\pi_s^R(q) = (w - c)q - (q - \hat{q}_r^{NR})^+ \int_{t^*}^{\infty} k(t - t^*)^+ f(t)dt.$$

Under the rebate scheme, we have $\int_{t^*}^{\infty} k(t-t^*)^+ f(t)dt = (w-c)$, which implies that for each unit that the retailer orders in excess of $\hat{q}_r^{NR}$, the expected rebate payout by the manufacturer is equal to his gross margin. Thus, for any $q \geq \hat{q}_r^{NR}$, the manufacturer is not strictly better off in expectation if he offers the rebate. We have already shown that the retailer is strictly better off in expectation

		Without Rebate	
		Constrained	Unconstrained
With Rebate Offer	Constrained	q is larger w/rebate	q is larger w/rebate
	Unconstrained	q is larger w/rebate	q is larger w/rebate
	Not constrained but no rebate	--	Same solution

Table 1: Comparison of Solutions with and without a Rebate

if he chooses an order quantity that makes him eligible for a rebate. Thus, the manufacturer can take part of the incremental profit by setting a larger "threshold" quantity (i.e., $\Lambda > \hat{q}_r^{NR}$). As such, even when the retailer imposes a downside risk constraint, the weather rebate leaves both parties at least as well off in expectation as they were in the absence of a rebate. Observe that the manufacturer still has the option to choose t^* to limit his risk. Furthermore, as discussed in Section 4, any additional risk borne by the manufacturer can often be hedged using a weather derivative.

The retailer could instead purchase a weather derivative directly. In Appendix C, we explore such a scenario and compare it with a manufacturer-offered rebate. The results suggest that the retailer and the supply chain are better off with a manufacturer-offered rebate instead of a retailer-purchased derivative. The manufacturer bears more risk but can completely hedge his risk by purchasing an equivalent weather derivative, and the additional supply chain profit can cover the risk premium for the weather derivative. On the other hand, if the retailer purchases the derivative, he pays a risk premium but with no increase in his expected revenue.

6 Conclusions

We have introduced and analyzed weather rebate contracts for newsvendor settings that can achieve supply chain coordination and allow an arbitrary allocation of profits between the two parties. The proposed class of rebates also provides Pareto improvement without the need to increase the existing wholesale price. More importantly, unlike other rebates designed to address excess end-of-season inventory, no inventory or markdown audits are necessary for enforcement of truth-telling and the contract has no adverse effect on sales effort. As such, the contract is easy to implement.

The class of coordinating weather contracts is also extremely flexible, allowing a wide range of functional forms and parameter values (such as the strike "temperature" at which the rebate is activated). This, in turn, allows the risk tolerances of the two parties to be reflected more easily. Interestingly, the flexibility poses technical challenges because decisions now need to be made where no analogous choice exists in other types of supply contracts. Moreover, not only do more decisions need to be made, but the associated analysis is more complicated because the presence of thresholds that must be decided, such as the strike temperature, causes the distribution of the rebate payments to take on more complicated forms. Our analysis illustrates how these challenges can be handled in cases where the retailer and manufacturer are concerned about limiting their downside risk. However, more work needs to be done to understand how rebate parameters should be structured when the retailer and manufacturer have other types of risk preferences. In practical implementation, the manufacturer needs a good understanding of which weather indices are the best predictors of demands, as well as the functional relationship between the selected weather index and the demand quantities of his many diverse customers. We expect that more statistical information along these lines will become available as weather derivative markets expand.

The rebate also can be incorporated when the parties can choose prices, and this is true whether the retailer must choose his price when he orders or closer to the selling season when improved weather information is available. We analyze both cases in Appendix D. Results for the case of power function demand reveal that the retailer's ability to price late and the weather rebate both have multiplicative (> 1) effects on profits in a compounding fashion, but the multiplicative effect of the weather rebate is not as strong when the retailer prices late, so late pricing and the weather rebate are partial substitutes. Both provide risk mitigation to the retailer while increasing his expected profit, and the manufacturer still obtains incremental benefits from the contract when the retailer prices late.

Financial services firms are beginning to offer business insurance policies to hedge against weather risk. Financial executives may be in a position to take advantage of these offerings. However, inventory managers rarely have the authority to purchase weather derivatives but they can accept a weather rebate offer from a manufacturer in the same way as they can agree to a buy-back contract or a markdown agreement. As such, forward-thinking manufacturers may be well-positioned to design and offer weather rebate contracts that would be attractive to their customers, thereby gaining a competitive advantage in the marketplace and simultaneously increasing their own profits.

Acknowledgements

This research was supported in part by the Hong Kong Research Grants Council under grant CUHK411105. The authors benefited from early discussions with Mr. Jing Zhu and assistance by Minghui Xu. The authors also acknowledge the helpful comments of two Associate Editors and three anonymous referees.

References

[1] Anupindi, R. and Y. Bassok. 1999. Supply Contracts with Quantity Commitments and Stochastic Demand, in **Quantitative Models for Supply Chain Management**, S. Tayur, R. Ganeshan and M. Magazine (eds.), Kluwer Academic Publishers, Boston, 197-232.

[2] Barnes-Schuster, D., Y. Bassok and R. Anupindi. 2002. Coordination and Flexibility in Supply Contracts with Options. *Manufacturing and Service Operations Management* **4** (3), 171-207.

[3] Blumenthal, R.G. 2007. Warm Weather Blues. *Barron's* **87** (40), 17.

[4] Boyabatli, O. and L.B. Toktay. 2004. Operational Hedging: A Review with Discussion. Working paper, Technology Management Department INSEAD, Fontainebleau, France.

[5] Business Wire, Weatherproof Adopts Apparel Industry's First-Ever Weather Hedge. 2007. *Business Wire*, December 3, 2007.

[6] Byron, E and T. Agins. Style and Substance: Probing Price Tags: A U.S. Attorney and the SEC Shine a Light on Stores' Murky 'Markdown Money': Is It Shared Risk-Or Uneven Playing Field? *Wall Street Journal* (Eastern edition), May 13, 2005, B1.

[7] Cachon, G. 2003. Supply Chain Coordination with Contracts, in Handbooks in Operations Research and Management Science: Supply Chain Management, S.C. Graves and A.G. de Kok (eds.), Elsevier, New York, 229-340.

[8] Cachon, G.P. 2004. The Allocation of Inventory Risk in a Supply Chain: Push, Pull, and Advance-Purchase Discount Contracts, *Management Science*. **50** (2), 222-238.

[9] Caldentey, R. and M.B. Haugh. 2009. Supply Contracts with Financial Hedging, *Operations Research*, **57** (1), 47-65.

[10] Chen, F.Y., F. Gao and X. Chao. 2007. Joint Optimal Ordering and Weather Hedging Decisions. Working paper, Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, Shatin, Hong Kong.

[11] Chen, X., M. Sim, D. Simchi-Levi and P. Sun. 2007. Risk Aversion in Inventory Management. *Operations Research* **55** (5), 828-842,

[12] Chod, J., N. Rudi and J. Van Mieghem. 2006. Operational Flexibility and Financial Hedging: Complements or Substitutes? Working paper, Boston College (revised 2009).

[13] Colin, C. 2008. Doing Something About the Weather, Financially, At Least. *Portfolio.com*, August 7, 2008. Available at: http://www. wired.com/ techbiz/ people/ news/ 2008/08/ portfolio_0807.

[14] Corbett, C. and C.S. Tang. 1999. Designing Supply Contracts: Contract Type and Information Asymmetry, in **Quantitative Models for Supply Chain Management**, S. Tayur, R. Ganeshan and M. Magazine (eds.), Kluwer Academic Publishers, Boston, 269-297.

[15] Davis, G.S. 2008. Knowing It's Going to Rain is Better than Guessing, available at http:// community.investopedia.com/news/ IA/2008/ Knowing_Its_Going _To_Rain_Is_Better_Than_Guessing _CMECS.aspx? partner=aol

[16] Davis, S.M. and C. Meyer. 2000. **Future Wealth**. Harvard Business School Press,Boston, MA.

[17] Devinney, T.M. and D.W. Stewart (1988). Rethinking the Product Portfolio: A Generalized Investment Model. *Management Science* **34** (9), 1080-1094.

[18] Dutton, J.A. 2002. Opportunities and Priorities in a New Era for Weather and Climate Services. *Bulletin of the American Meteorological Society*, **83** (9), 1303-1311.

[19] Edelson, S. 2005. A Business Practice's Evolution, *WWD: Women's Wear Daily* 189(120), 17.

[20] Emmons, H. and S. Gilbert. 1998. Return Policies in Pricing and Inventory Decisions for Catalogue Goods. *Management Science*, **44**(2), 276-283.

[21] Environmental Finance. 2004. Looking Beyond the Energy Sector. Available at: http://www.environmental-finance.com/2004/0403mar/hedge.htm.

[22] Fisher, M. and A, Raman. 1996. Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales. *Operations Research* **44** (1), 87-99.

[23] Gan, X., S. Sethi, H. Yan. 2005. Channel Coordination with a Risk- Neutral Supplier and a Downside-Risk-Averse Retailer. *Production and Operations Management* **14**(1), 80-89.

[24] Gaur, V. and S. Seshadri. 2005. Hedging Inventory Risks Through Market Instruments. *Manufacturing and Service Operations Management 7* (2), 103-120.

[25] Gilbert, S.M. and R.H. Ballou. 1999. Supply Chain Benefits from Advanced Customer Commitments. *Journal of Operations Management* **18** (1), 61-73.

[26] Glenn, D. W. 2004. Returns Policies for a Pessimistic Retailer. Available at http://www.business.uiuc.edu/Working_Papers/papers/04-0111.pdf

[27] Gottlieb, R. Markdown Money: It's Just a Bad Idea. *Playthings*, October 2005.

[28] Granot, D. and S. Yin. 2006. On the Effectiveness of Returns Policies in the Price-Dependent Newsvendor Model. *Naval Research Logistics* **52** 8, 765-779.

[29] Iyer, A. and M.E. Bergen. 1997. Quick Response in Manufacturer-Retailer Channels. *Management Science* **43** (4), 599-570.

[30] Kandel, E. 1996. The Right to Return. *Journal of Law and Economics*, **39**, 329-356.

[31] Kleiderman, A. 2004. Soggy Summer Spells Boardroom Gloom, Available at http://news.bbc.co.uk.

[32] Krishnan, H., R. Kapuscinski and D.A Butz. 2004. Coordinating Contracts for Decentralized Supply Chains with Retailer Promotional Efforts. *Management Science*, 50 (1), 48-63.

[33] Lariviere, M.A. 1999. Supply Chain Contracting and Coordination with Stochastic Demand. in **Quantitative Models for Supply Chain Management** (Chapter 8), eds. S. Tayur, R. Ganeshan and M. Magazine. Boston, Kluwer Publisher. 233-268.

[34] Lariviere, M.A. and E. L. Porteus. 2001. Selling to the Newsvendor: An Analysis of Price-Only Contracts. *Manufacturing & Service Operations Management*, **3**(4), 293-305.

[35] Lau, H.S. and A. H.L. Lau. 1999. Manufacturer's Strategy and Return Policy for a Single-Period Commodity. *European Journal of Operational Research*, **116**, 291-304.

[36] Lee, H.L. and C.S. Tang. 1997. Modeling the Costs and Benefits of Delayed Product Differentiation. *Management Science* **43** (1), 40-53.

[37] Lustgarden, A. 2005. Getting Ahead of the Weather. *Fortune* **151** (3), 87-94.

[38] Malinow, M. 2002. End-users. In Banks, E. (Ed.), **Weather Risk Management**. Palgrave, NY.

[39] Martinez-de-Albeniz, V. and D. Simchi-Levi. 2005. A Portfolio Approach to Procurement Contracts. *Production and Operations Management* **14** (1), 90-114.

[40] Marvel, H.P. and J. Peck. 1995. Demand Uncertainty and Returns Policies. *International Economic Review*, 36, 691-713.

[41] McCardle, K., K. Rajaram and C.S. Tang. 2004. Advance Booking Discount Programs under Retail Competition. *Management Science* **50** (5), 701-708.

[42] MSI Guaranteed Weather. Undated. Temperature Effects on Beverage Sales. Available at: http://www.guaranteedweather.com/content_page.aspx?content_id=36.

[43] Niemira, M. P. 2005. Weather Matters. *International Council of Shopping Centers Research Review* **12** (2), 23-27.

[44] O'Donnell, J. 2007. 'Wrong' Weather Can Foul Up Retailers' Sales. *USA Today*, November 22, 2007, 1B.

[45] Padmanabhan, V. and P.L. Png. 1995. Returns Policies: Making Money by Making Good. *Sloan Management Review*, 37 (1), 65-72.

[46] Pasternack, B.A. 1985. Optimal Pricing And Return Policies For Perishable Commodities. *Marketing Science*, **4**(2), 166-176.

[47] Regnier, E. 2008. Doing Something About the Weather. *Omega* **36** 22-32.

[48] Starr-McCluer, M. 2000. The Effects of Weather on Retail Sales, Federal Reserve Board of Governors Report.

[49] Taylor, T. A. 2001. Channel Coordination Under Price Protection, Midlife Returns, and End-of-Life Returns in Dynamic Markets. *Management Science*, **47**(9), 1220-1234.

[50] Taylor, T. A. 2002. Supply Chain Coordination Under Channel Rebates with Sales Effort Effects. *Management Science*, **48**(8), 992-1007.

[51] Timberlake, C. and G. Wiles. 2005. Retailers' May Sales Rise 2.9%; Fashion Stores Lead (Update5), Available at www.bloomberg.com

[52] Tsay, A. 2001. Managing Retail Channel Overstock: Markdown Money and Returns Policies. *J. Retailing*, 77 (4), 457-492.

[53] Tsay, A.A., S. Nahmias and N. Agrawal. 1999. Modeling Supply Chain Contracts: A Review, in **Quantitative Models for Supply Chain Management**, S. Tayur, R. Ganeshan and M. Magazine (eds.), Kluwer Academic Publishers, Boston, 299-336.

[54] Van Mieghem, J.M. 2007a. Capacity Management, Investment and Hedging: Review and Recent Developments. *Manufacturing and Service Operations Management* **5** (4), 269-302.

[55] Van Mieghem, J.M. 2007b. Risk Mitigation in Newsvendor Networks: Resource Diversification, Flexibility, Sharing and Hedging. *Management Science***53** (8), 1269-1288.

[56] Webster, S. and Z. K. Weng. 2002. A Risk-free Perishable Item Returns Policy. *Manufacturing & Service Operations Mgt*, **2**(1), 100-106.

[57] Wood, D. 2007. Weathering the Odds, available at http:// www.treasuryandrisk.com/topic/risk/987.

[58] Zhou, J. and N. Rudi. 2007. Pricing of Over-the-Counter Hedging Contracts for Real Investments. Working paper, Simon Graduate School of Business Administration, University of Rochester, Rochester, NY.

Appendices

Appendix A1: Model with Non-monotonic $d(p, t)$

Here we assume that the retail price is exogenously given, thus $d(t)$ is a deterministic function of t, but the temperature itself is a random variable. We model a situation in which $d(t)$ is non-monotonic but unimodal in t, first strictly increasing as t increases, then strictly decreasing. (Generalization to accommodate general unimodal functions is straightforward.) With such a representation, for each demand $\hat{d}$, there is a unique pair of temperatures, $\underline{t}(\hat{d})$ and $\bar{t}(\hat{d})$ such that for $t \in [\underline{t}(\hat{d}), \bar{t}(\hat{d})]$, $d(t) \geq \hat{d}$.

We can then write the retailer's expected profit as:

$$\begin{aligned}
\pi_r(q) &= -wq + pq[F(\bar{t}(q)) - F(\underline{t}(q))] \\
&+ p\int_{t=0}^{\underline{t}(q)} x(t)f(t)dt + p\int_{\bar{t}(q)}^{\infty} x(t)f(t)dt \\
&+ \int_{t=0}^{\underline{t}^*} K(t, q(t^*))f(t)dt + \int_{\bar{t}^*}^{\infty} K(t, q(t^*))f(t)dt
\end{aligned}$$

Taking the derivative with respect to q and simplifying, we obtain:

$$\begin{aligned}
-w &+ p[F(\bar{t}(q)) - F(\underline{t}(q))] \\
&+ \int_{t=0}^{\underline{t}^*} \frac{\partial K(t,q)}{\partial q} f(t)dt + \int_{t=\bar{t}^*}^{\infty} \frac{\partial K(t,q)}{\partial q} f(t)dt
\end{aligned}$$

Although this expression is slightly more complicated that in the case of a monotonic $d(t)$, it has the same basic structure. Consequently, the results for monotonic $d(t)$ also extend to the case of unimodal $d(t)$.

Appendix A2: Proof of Theorem 2

First consider the manufacturer. To make the rebate attractive to the manufacturer, it must be such that

$$[(w-c)q - \int_{t^*}^{\infty} k(t-t^*)(q-\Lambda)^+] - (w_d - c)q_r \geq 0, \tag{B-1}$$

where the term in brackets on the left hand side is the expected profit of the manufacturer when the rebate is offered with the wholesale price being set at $w \geq w_d$, and the second term is that without the rebate. By (4), a Pareto–improving rebate scheme exists for the manufacturer if

$$[(w-c)q - (w-c)(q-\Lambda)^+] - (w_d - c)q_r \geq 0.$$

If the retailer orders $q \geq \Lambda$, then the condition becomes

$$(w-c)\Lambda \geq (w_d - c)q_r.$$

Clearly, if $\Lambda \geq \Lambda_w$, the manufacturer is better off with the rebate.

We now consider the retailer. From Theorem 1, especially (5), it can be seen that under the rebate contract, the retailer is better off with the decision q where $q > \Lambda$ if

$$R(q) - cq - (w-c)\Lambda - \pi_r^d(w_d) > 0$$

where $\pi_r^d(w)$ is the retailer's expected profit in the absence of a rebate. Because

$$R(q_c) - cq_c - (w-c)\Lambda - \pi_r^d(w_d) \geq R(q) - cq - (w-c)\Lambda - \pi_r^d(w_d) > 0,$$

for any q, the retailer will choose q_c with the rebate. Hence,

$$R(q_c) - cq_c - (w-c)\Lambda - \pi_r^d(w_d) > R(q_c) - cq_c - (w-c)\bar{\Lambda} - \pi_r^d(w_d) = 0$$

as $R(q_c) - cq_c = \pi_c(q_c)$ and $\bar{\Lambda} = (\pi_c(q_c) - \pi_r^d(w_d))/(w-c)$. This completes the proof. □

Appendix B: Limiting the Retailer's Risk – Detailed Analysis

Without a Rebate

For ease of reference, we repeat some of the equations appearing earlier here. The retailer's problem is

$$\max \pi_r^{NR}(q) = E\{p \min[q, D(t,u)] + v(q - D(t,u))^+ - wq\}, \tag{B-1}$$

$$s.t. \ \ \mathsf{Pr}\{\Pi_r^{NR}(q) \leq \alpha\} \leq \beta. \tag{B-2}$$

where NR denotes "no rebate." For the value of q selected by the retailer, the manufacturer's profit is $\pi_m^{NR}(q) = (w - c)q$.

If $q < \frac{\alpha}{p-w}$ then it is impossible to satisfy the retailer's risk constraint. Let $q_\alpha = \frac{\alpha}{p-w}$, i.e., the threshold value of q above which there is a positive probability that the retailer will achieve a profit α (or higher). Then the probability that the retailer's profit fails to reach α for an order quantity, q, is

$$\mathsf{Pr}\{\Pi_r^{NR}(q) \leq \alpha\} = \begin{cases} 1, & \text{if } q \leq q_\alpha, \\ \mathsf{Pr}\{D(t,u) \leq \frac{\alpha+(w-v)q}{p-v}\}, & \text{if } q > q_\alpha. \end{cases} \tag{B-3}$$

Note that the probability of failing to achieve a profit of α is **increasing** with q, so there is an inherent tradeoff between increasing expected profit (which increases with q up to the unconstrained optimum) and keeping the probability of failing to achieve a profit of α low.

Let q_r be the unconstrained optimal order quantity in the absence of a rebate, i.e., $q_r = \arg\max_q \pi_r^{NR}(q)$ (or see Section 3). Then for given values of q and d, $d \leq q$, the retailer's profit is:

$$dp + (q - d)v - wq \tag{B-4}$$

which is less than or equal to α if $d \leq \frac{\alpha+(w-v)q}{p-v}$. Therefore, we have:

(1) if $\beta \leq \mathsf{Pr}\{D(t,u) \leq q_\alpha\}$, there is no feasible solution;

(2) if $\mathsf{Pr}\{D(t,u) \leq q_\alpha\} < \beta \leq \mathsf{Pr}\{D(t,u) \leq \frac{\alpha+(w-v)q_r}{p-v}\}$, then the optimal order quantity is $\bar{q}_r$ where $\bar{q}_r$ is the maximum value of q that satisfies the downside risk constraint, i.e., $\bar{q}_r^{NR} = \max\{q : \mathsf{Pr}\{\Pi_r^{NR}(q) \leq \alpha\} \leq \beta\}$;

(3) if $\mathsf{Pr}\{D(t,u) \leq \frac{\alpha+(w-v)q_r}{p-v}\} < \beta$, then the optimal order quantity is the unconstrained solution, q_r.

From the above, we can conclude that if $\mathsf{Pr}\{D(t,u) \le q_\alpha\} < \beta$, the retailer's optimal (constrained) order quantity when there is no rebate is $\hat{q}_r^{NR} = \min\{q_r, \bar{q}_r^{NR}\}$.

The rest of the analysis follows the main text.

With a rebate

Here, we set $\Lambda = \hat{q}_r^{NR}$ (the constrained optimal order quantity when there is no rebate) in a manner analogous to what we did in the risk-neutral setting. Later in this subsection, we discuss how the contract parameters affect risks in a broader sense.

The retailer's problem becomes

$$\max_q \pi_r^R(q) = E\{p\min\{q, D(t,u)\} + v(q - D(t,u))^+ - wq + k(t-t^*)^+(q - \hat{q}_r^{NR})^+\} \quad \text{(B-5)}$$

$$s.t.\ \ \mathsf{Pr}\{\Pi_r^R(q) \le \alpha\} \le \beta. \quad \text{(B-6)}$$

and the manufacturer's profit is

$$\pi_s(q) = \mathsf{E}\{(w-c)q - k(t-t^*)^+(q - \hat{q}_r^{NR})^+\}. \quad \text{(B-7)}$$

Let q_r^R denote the (unconstrained) order quantity that maximizes the retailer's expected profit when the manufacturer offers a rebate. (With a coordinating rebate, $q_r^R = q_c$.) Since the rebate satisfies $\int_{t^*}^{\infty} k(t-t^*)f(t)dt = w - c$ and $c > v$ by assumption, it is always true that $\int_{t^*}^{\infty} k(t - t^*)f(t)dt < w - v$, i.e., the expected rebate compensation per unit ordered in excess of q_r is less than the retailer's net overage cost, $w - v$. It can be seen that q_r^R is finite and unique.

We next show that $q_r^R > \hat{q}_r^{NR}$. We do so by showing that $q_r^R > q_r$ which we know is greater than or equal to $\hat{q}_r^{NR}$. We know that $\hat{q}_r^{NR}$ is feasible for the scenario with a rebate, so this result will allow us to restrict our attention to solutions with $q \ge \hat{q}_r^{NR}$. For notational simplicity, let $H(q) = \int_{t^*}^{\infty} D(t,q)dt$. Taking the derivative of π_r^R with respect to q, we have:

$$\frac{\partial \pi_r(q)}{\partial q} = \begin{cases} p - w - (p-v)H(q) & \text{if } q \le \hat{q}_r^{NR}, \\ p - w - (p-v)H(q) + \int_{t^*}^{\infty} k(t-t^*)f(t)dt, & \text{if } q > \hat{q}_r^{NR}. \end{cases} \quad \text{(B-8)}$$

We know that $H(q_r) = \frac{p-w}{p-v}$. Therefore, at $q = \hat{q}_r^{NR}$ the partial derivative is strictly positive, and because the profit function is concave, the partial derivative is strictly positive at $q = \hat{q}_r^{NR} \le q_r$

as well. So the value of q that equates the second expression in (B-8) to zero, i.e., $\bar{q}_r^R$, is strictly greater than q_r, which in turn is greater than $\hat{q}_r^{NR}$.

We now proceed to derive $\mathsf{Pr}\{\Pi_r(q) \le \alpha\}$ for the scenario with a rebate. For $q > \hat{q}_r^{NR}$ and a given demand d, the retailer's profit is:

$$p \min(d, q) + v(q - d)^+ - wq + k(t - t^*)(q - \hat{q}_r^{NR}) \tag{B-9}$$

if $t > t^*$, so the retailer's profit is α or less if

$$d \le \frac{\alpha + (w - v)q + k(t - t^*)(q - \hat{q}_r^{NR})}{p - v} \text{ if } d < q \tag{B-10}$$

and if $d \ge q$ the profit is

$$(p - w)q + k(t - t^*)(q - \hat{q}_r^{NR}), \tag{B-11}$$

which is a deterministic profit and is always greater than α for $q > q_\alpha = \alpha/(p - w)$.

Similarly, for $q < \hat{q}_r^{NR}$, the retailer's profit is

$$p \min(d, q) + v(q - d)^+ - wq \tag{B-12}$$

if $t \le t^*$, so the retailer's profit is α or less if

$$d \le \frac{\alpha + (w - v)q}{p - v} \text{ if } d < q \tag{B-13}$$

and if $d \ge q$ the profit is

$$(p - w)q \tag{B-14}$$

which again is a deterministic quantity that is less (greater) than α for q less (greater) than $\alpha/(p - w)$. This is the same as in the no-rebate case.

From the foregoing analysis, we can write

$$\Pr\{\Pi_r(q) \leq \alpha\} = \begin{cases} 1, & \text{if } q \leq q_\alpha, \\ \Pr\{D(t,u) \leq \frac{\alpha+(w-v)q}{p-v}\}, & \text{if } q_\alpha < q \leq \hat{q}_r^{NR} \\ \Pr\{D(t,u) \leq \frac{\alpha+(w-v)q}{p-v} - k(t-t^*)\frac{q-\hat{q}_r^{NR}}{p-v}, t > t^*\} & \\ +\Pr\{D(t,u) \leq \frac{\alpha+(w-v)q}{p-v}, t \leq t^*\}, & \text{if } q > \hat{q}_r^{NR}. \end{cases} \tag{B-15}$$

Notice that $\Pr\{\Pi_r(q) \leq \alpha\}$ is a continuous function.

If q_r^R satisfies the risk constraint, then we are done. Otherwise, we need to identify the best constrained solution. Under the assumption that $k(t - t^*) > 0$ for all $t > t^*$, for any $q > \hat{q}_r^{NR}$, the expression on the right hand side of the third entry in (B-15) is strictly less than $\frac{\alpha+(w-v)q}{p-v}$. In other words, due to the potential for receiving a rebate, not surprisingly, the probability that the retailer's profit falls short of the threshold α declines for any fixed q. Thus, there exists some $q > \hat{q}_r^{NR}$ such that the risk constraint is still satisfied. Let $\bar{q}_r^R$ be the largest value of q for which the third entry on the right hand side of (B-15) is less than or equal to β. Then the constrained optimal order quantity, $\hat{q}_r^R$ is equal to min $(\bar{q}_r^R, q_r^R)$.

So far, we have assumed $\Lambda = \hat{q}_r^{NR}$. For $\Lambda > \hat{q}_r^{NR}$, all entries in (B-15) remain unchanged except with $\hat{q}_r^{NR}$ being replaced by Λ. Therefore, the foregoing analysis also applies. However, it is easy to see that if Λ is too large, then the retailer's risk constraint may not be satisfied for any q.

Appendix C: Retailer's Choice between a Third-Party Weather Derivative and the Manufacturer's Weather Rebate

In this Appendix, we explore differences between the retailer purchasing a third-party derivative and utilizing a manufacturer-offered weather rebate. First, we consider a coordinating rebate where coordinating means that the rebate satisifies (4), so the maximum overall (first-best) supply chain profit is achieved. Under this coordinating weather rebate, the retailer gains all of the incremental expected profit and the manufacturer faces a greater variability of profit, so a risk-averse manufacturer would not offer the rebate. We show that the manufacturer can structure non-coordinating weather rebates such that the retailer satisfies his risk constraint and achieves a strictly greater expected profit while the manufacturer increases his expected profit by more than enough to cover the risk premium for a weather derivative to completely hedge his risk. (In a non-coordinating rebate (4) is satisfied as a "<" relationship rather than an equality.) In the absence of retailer risk constraints, one could use a side payment from the retailer to the manufacturer to cover the risk premium for the weather derivative, but when the retailer's risk constraint must be considered, such side payments also affect the retailer's decision when the risk constraint is binding. This gives rise to the possible need for a non-coordinating contract *without a side payment that is related to the risk premium* to enable the manufacturer to extract enough additional profit to cover the risk premium.

We illustrate our results via a numerical example (the same one discussed in Section 4.2), but the patterns hold more generally for reasons that we will explain as the discussion proceeds. Table A1 shows the retailer's expected profit under no derivatives or rebate, with retailer-purchased derivatives and with a manufacturer-offered rebate that has exactly the same payout as the derivatives. Here, we set $\alpha =$ (mean demand - 2 standard deviations)*$(p - w)$ which is an approximation of the retailer's "guaranteed" profit if he orders a very low percentile of the demand distribution and sells it all. The β values are shown in the first column in the table. We also assume the retailer pays a zero risk premium for the derivatives when computing his expected profit; this gives the derivatives their maximum advantage, as risk premia are typically positive.

Table A1: Retailer's Expected Profit, Order Quantities and Expected Rebate Payouts

β	No Deriv. or Rebate	Exp. Deriv. or Rebate Payout	With Deriv.	With Rebate	% Profit Increase[3]
0.05	23,789[1] (4,950)[2]	145	23,805 (4,955)	24,028 (4,980)	1.0%
0.10	25,343 (5,550)	239	25,359 (5,558)	25,677 (5598)	1.3%
0.15	26,187 (6,135)	314	26,196 (6,145)	26,554 (6,198)	1.4%
0.20	26,406 (6,612)	1,297	26,406 (6,612)	27,640 (6,872)	4.7%
0.25	26,406 (6,612)	6,601	26,406 (6,612)	31,354 (7,935)	18.7%

[1] retailer's expected profit; [2] order quantity; [3] with rebate versus no derivative or rebate.

For this particular example, the retailer's risk constraint is binding for $\beta \leq 0.195$. Notice that for β in this range, the retailer achieves no increase in expected profit from the derivatives. (The small differences in the table are due to unavoidable numerical imprecision in the calculations. It is easy to show that the retailer has no expected gain from a derivative when the risk premium is zero.) The variance of the retailer's profit including any payout from the derivative (values not reported here) does decline, so the derivatives are beneficial if the retailer is concerned only with risk mitigation and not with increasing his expected profit. Indeed, his expected profit would decline if he must pay a risk premium for the derivatives.

The manufacturer-offered rebate provides for a small (1% to 1.5%) increase in the retailer's expected profit when $\beta \leq 0.195$. This arises because the retailer orders slightly more, which shifts the mean of his profit distribution upward, but the variance of his profits also increases slightly. The retailer's risk constraint limits how much he is willing to order because he cannot gain much upside potential (from increased sales) unless he orders more, and ordering more also increases his probability of not satisfying his risk constraint. When the retailer's risk constraint is stringent, a manufacturer-offered rebate is limited in the degree to which it can induce larger orders, but the expected benefit is always positive, whereas for the retailer-purchased derivatives, the expected benefit is, at best, zero.

For $\beta > 0.195$, the retailer's risk constraint is not binding, and the rebate offers sizable advantages over the derivatives for the retailer. The retailer will not purchase a derivative because he has no need for further risk reduction and any risk premium for the derivative will reduce his expected profit. On the other hand, a manufacturer-offered rebate provides value to the retailer even though his risk constraint is not binding.

Observe that the retailer's order quantity in the scenarios with retailer-purchased weather derivative never exceeds 6,612, which is the retailer's *decentralized* order quantity when he is risk neutral. Thus, the retailer-purchased weather derivatives do not counteract the effect of double marginalization; they only (partially) counteract the retailer's risk aversion. On the other hand, the manufacturer-offered rebates do partially but markedly counteract the effect of double marginalization when the retailer's risk constraint is not binding, and even offer some benefits along these lines when the retailer's risk constraint is binding.

As noted earlier, under the manufacturer-offered rebate, all of the incremental supply chain profit goes to the retailer (in expectation) and the manufacturer's profit variance increases due to the possibility of having to pay a rebate. So a risk averse manufacturer may be unwilling to offer the rebate unless he can find a source of profit to cover the risk premium for weather derivatives to hedge his risk from the rebate. In this case, particularly when the retailer's risk constraint is binding, one cannot simply use a side payment from the retailer to the manufacturer to cover this risk premium because the side payment changes the retailer's profit distribution.

One option is to use a non-coordinating contract that enables the manufacturer to extract some of the incremental profit. Results are shown in Table A2. Here, we have selected $E[k(t,t^*)] < (w-c)(q-\Lambda)^+$ where Λ is the retailer's order quantity in the absence of a rebate. Notice that in a coordinating rebate, $E[k(t,t^*)] = (w-c)(q-\Lambda)^+$. Thus, the difference $(w-c)(q-\Lambda)^+ - E[k(t,t^*)]$ represents the manufacturer's profit increase. Observe that even though the manufacturer has designed the contract to be generous to the retailer and the rebate is not a coordinating one, the manufacturer's profit increases by about 8.9% of the standard deviation of the rebate payout, which is likely to be enough to cover the typical risk premia for the weather derivative that he would purchase to completely hedge his risk of offering the weather rebate. (In our example, the retailer is strictly better off so the manufacturer could extract more of the profit for himself.) Researchers have suggested several different methods for calculating premia (expected payout plus the risk premium) for weather derivatives, noting that the market for weather derivatives is incomplete because there is no underlying asset. For example, Barrieu and Scaillet (2009) discuss methods based on Black-Scholes or capital asset pricing model (CAPM) concepts, methods based on the expected utility of terminal wealth, and methods based on equilibrium analysis considering the utility functions of the buyer and seller, while Davis (2001) offers a method based on marginal substitution value.

Jewson (2004) indicates that utility-based approaches are not used in practice and suggests a CAPM-based model in which the risk premium is a fraction of the standard deviation of the payout. Although the risk premium is quite naturally a function of the variability of the payout, it also depends upon the probability of a positive payout. For example, Jewson and Brix (2005) give an example with a risk premium equal to 20% of the standard deviation of the payout for a rebate with a strike set at 0.25 standard deviation from the mean; such rebate has a relatively low coefficient of variation of the payout and a relatively high payout probability. For the rebate considered here, the probability of a positive payout is less than 15%, which is typical for a weather rebate or derivative, as the intent is to hedge against severe outcomes, not against small deviations from the norm. As such, the standard deviation of the rebate payout is large relative to the mean. Thus, a factor far less than 20% would be appropriate here. Indeed, if we set the strike temperature at 0.25 standard deviation above the mean (i.e., at 71.25 degrees), leading to a rebate with a more frequent but smaller average payout, we obtain the results in Table A3: the retailer still benefits and the manufacturer's profit increases by more than 20% of the standard deviation of the payout, which would comfortably cover the risk premium for a weather derivative.

Table A2: Outcomes from Non-Coordinated Contracts: Typical Strike Temperature

β	Order Quant.	Retailer's Profit	Rebate Payout Mean (Std. Dev.)	Mfr's Add'l Profit
0.05	4,975	23,975 (0.7%)[1]	101 (270)	23.9 (8.85%)[2]
0.10	5,590	25,594 (0.9%)	162 (431)	38.2 (8.86%)
0.15	6,185	26,454 (0.9%)	202 (539)	47.8 (8.87%)
0.20	6,860	27,461 (3.6%)	1002 (2880)	237.4 (8.87%)
0.25	7,856	29,974 (13.5%)	5030 (13,411)	1189.0 (8.87%)

[1] percent increase from the case with no derivatives or rebate
[2] percentage of the standard deviation of the payout

Table A3: Outcomes from Non-Coordinated Contracts: Low Strike Temperature

β	Order Quant.	Retailer's Profit	Rebate Payout Mean (Std. Dev.)	Mfr's Add'l Profit
0.05	4,970	23,928 (0.6%)[1]	76 (111)	23.56 (21.23%)[2]
0.10	5,568	25,447 (0.4%)	69 (99.9)	20.2 (21.23%)
0.15	6,162	26,314 (0.5%)	103 (150)	31.8 (21.23%)
0.20	6,806	27,112 (2.0%)	741 (1,077)	228.5 (21.23%)
0.25	7,816	29,638 (12.2%)	4,602 (6,682)	1,418 (21.23%)

[1] percent increase from the case with no derivatives or rebate
[2] percentage of the standard deviation of the payout

In summary, a manufacturer-offered weather rebate has advantages over retailer-purchased weather derivatives except in instances where the retailer is only interested in reducing his variability of profits and is willing to sacrifice expected profit to achieve this reduction. Furthermore, as mentioned earlier, a buyer at a retail store can easily avail him/herself of a weather rebate in much the same way as he/she can take advantage of a markdown agreement, but purchasing a weather derivative would be beyond the bounds of his/her usual authority.

In principle, a retailer can choose to utilize both a weather rebate and a weather derivative, but the best way to structure a weather derivative to complement a weather rebate is likely to depend upon subtle details of the retailer's risk preferences. This remains a topic for further research.

Appendix D: The Price-Setting Supply Chain with Possible Retail Price Postponement

We consider generalizations of our basic model in which the manufacturer chooses the wholesale price in addition to the contract terms. The retailer places an order well before the beginning of the selling season, and either (i) postpones his pricing decision until after he has observed a signal regarding the weather; or (ii) decides price at the same time as he makes his production decision (e.g., in the case of catalog sales).

We derive the decentralized equilibria for these Stackelberg games and show how a weather–linked contract can be incorporated within this framework. For the case of a base demand that has a power function form, we illustrate the joint effects of early versus late pricing and the presence or absence of the weather contract and show that although the weather contract and late pricing are partial substitutes, both have compounding effects on profit levels.

The literature on price postponement is relatively young and much of it is focused on its impact on capacity decisions. Van Mieghem and Dada (1999) were among the first to examine the relative benefits of postponing pricing and production decisions on capacity decisions in a newsvendor framework. They found that postponement of pricing decisions is usually more beneficial and it limits the benefits of production postponement. On the other hand, Granot and Yin (2008) conclude that in a newsvendor setting with a buy-back option in place, neither order postponement nor pricing postponement has an effect on equilibrium outcomes. Researchers have considered price postponement in settings where there is resource flexibility of some type (e.g., Bish 2006, Chod and Rudi 2005, Biller et al. 2006, Anupindi and Jiang 2005). Some of the above papers consider competitive effects. Bish and Wang (2004) and Bish and Suwandechochai (2005) extend this line of research to multi-product settings. Chan et al. (2006) consider a multi-period model with production and inventory (but not capacity) decisions and find that production postponement is usually more beneficial than price postponement. Other publications that consider price postponement include Ruiz-Benitez (2007), Lenk (2008) and Tang and Tomlin (2008).

Model with Price Postponement

We assume that events unfold as follows. The manufacturer, as the Stackelberg leader, sets the wholesale price. The retailer then decides his purchase quantity, q, based upon a temperature distribution that may be a Bayesian prior or based on historical data–or any other source, long before the selling season starts. (For notational simplicity, we omit the subscript r.) Then, the retailer observes an unbiased signal of the weather (e.g., an accurate forecast of the average temperature during the selling season) at the beginning of the selling season, which can be viewed as a draw from his distribution at the beginning of the season, and sets the selling price accordingly. In reality, the weather signal may not be unbiased or accurate and the retailer can change the price during the season to adapt to actual market and weather conditions. Our reason for adopting these particular assumptions is to reflect the facts that the ordering decision must be made before accurate weather information is available, and that the retailer has the flexibility to choose a price (or a trajectory of prices) in view of the (more accurate) weather information at the beginning of and/or during the selling season. We recognize that this is simply an approximation, but we believe it is a more reasonable representation than simply assuming that the retailer must set the price long before the selling season starts.

We assume that any shortage cost beyond the loss of profit is zero. This assumption is not without loss of generality but it simplifies the technical analysis and exposition.

We consider a retail demand function $D(p, t, u)$ consisting of two parts: (1) a riskless part, $d(p, t)$ which is a deterministic function that is decreasing in the retail price p and is non-increasing in t, and (2) a random part, U, a non-negative random variable defined over a finite support $[A, B)$ with mean μ (> 0), distribution $G(u)$ and density $g(u)$. We assume that the demand uncertainty takes the multiplicative form:

$$D(p, t, u) = d(p, t)U.$$

We assume T and U are independent random variables and $d(p, t)$ is decreasing in both p and t. Without loss of generality, we also assume that $E(U) = \mu = 1$.

From a technical standpoint, the multiplicative model of demand uncertainty is more amenable to analysis, but we believe that it is also a more realistic representation of the demand uncertainty caused by weather variability, namely that the demand uncertainty is higher when the mean is higher. In other words, it is more difficult to predict extreme (high) demands in conditions when the weather is expected to increase the (average) forecasted demand.

Before proceeding with the analysis, we present two assumptions regarding the deterministic and random components of the demand that apply in the remainder of this section.

Assumption 1 *The function $d(p,t)$ is decreasing in both p and t and has increasing price elasticity (IPE) for any given t, where the price elasticity is defined as*

$$\eta(p,t) = -p\frac{d'}{d} \ (\geq 1).$$

Assumption 2 *The random variable U has an increasing generalized failure rate (IGFR).*

Discussion of Assumptions 1 and 2. The IPE property is intuitive: as the price increases, the demand decreases by a larger percentage, which eventually makes it less desirable to raise the price further. Many commonly-used demand functions in the literature satisfy the IPE assumption. (For examples, see Chen et al. 2004). If $\eta(p) < 1$ for all p, i.e., the product is price-insensitive, then the price should be set to the maximum possible level. IGFR is a generalization of IFR, so the set of distributions satisfying the IGFR property is a superset of those satisfying the IFR property (Lariviere, 1999). The IGFR property is satisfied by almost all theoretical distributions used in the operations management literature (see Petruzzi and Dada, 1999; and Chen et al. 2004).

As in the previous section, we model the situation as a manufacturer-Stackelberg game. The difference here is that the retailer's problem is now a two-stage stochastic program. In the second stage (after the temperature signal is observed), the retailer optimizes p given the q that was decided before the beginning of the selling season and the observed t. In the first stage (before the temperature signal is observed), the retailer optimizes q in view of the distribution of T and his optimal price for each outcome t.

The Retailer's Problem

Because the retailer's problem is a two-stage stochastic program, we solve the problem by backward induction. At the beginning of the selling season (the second stage of the stochastic program), the retailer chooses the optimal price $p^*(q,t)$ with q and t given.

The retailer's profit is given by

$$\Pi_r(p|q,t) = pd(p,t)[1 - \Theta(\frac{q}{d(p,t)})] - wq \tag{D-1}$$

where $\Theta(z) = \int_z^B (u-z)g(u)du$, $z \in [A,B)$, and w is the wholesale price offered by the manufacturer.

The expression $[1-\Theta(\cdot)]$ can be interpreted as the fill-rate (fraction of demand satisfied from stock) given q.

The partial derivative of the profit function in (D-1) with respect to p for fixed q and t is

$$\frac{\partial \Pi_r(p|q,t)}{\partial p} = d[1-\Theta(\frac{q}{d})] + p\frac{\partial d}{\partial p}\lambda(\frac{q}{d}), \tag{D-2}$$

where $\lambda(\frac{q}{d}) = \int_0^{\frac{q}{d}} ug(u)du$.

In the following, we will use d and $p(q)$ to represent $d(p,t)$ and $p(q,t)$, respectively, to simplify the exposition, until it is necessary to make the dependence on t explicit.

We characterize the retailer's profit function in the following theorem.

Theorem 3 *Under Assumptions 1 and 2, for any given order quantity q (≥ 0) and temperature t: (a) the retailer's expected profit function $\Pi_r(p|q,t)$ is unimodal in p. (b) $p^*(q,t)$ is non-increasing in q for any fixed t, (c) $p^*(q,t)$ is decreasing in t for any fixed q and (d) $p^*(q,t)$ satisfies*

$$v(\frac{q}{d(p,t)}) \quad = \quad \eta(p(q,t)) \textit{ for all } t. \tag{D-3}$$

and does not depend on w.

Proof. To simplify the exposition, in the proof we use d to represent $d(p,t)$.

(a) We first show that $\Pi_r(p|q,t)$ is unimodal in p for fixed q and t. The partial derivative of the profit function in (D-1) with respect to p for fixed q and t is given by (D-2).

Rearranging the right hand side of (D-2), we have

$$\frac{\partial \Pi_r(p|q,t)}{\partial p} = d\lambda(\frac{q}{d})[v(\frac{q}{d}) - \eta(p)] \tag{D-4}$$

where $v(z) = \frac{1-\Theta(z)}{\lambda(z)}$. Because of Assumption 2, $v(z)$ is decreasing in z (Song et al., 2008). Given q and t, as p increases, so does $\frac{q}{d}$ because d decreases with p. Together with the first order condition, it is fairly easy to show the rest of part (a).

(b) We next show that $p^*(q,t)$ is non-increasing in q for fixed t. Notice that the first two terms in (D-4), namely d and $\lambda(\frac{q}{d})$, cannot be zero or negative in an optimal solution with positive profit. Thus, in an optimal solution, the price must equate the expression in square brackets to zero, i.e., the optimal price satisfies

$$v(\frac{q}{d}) - \eta(p) = 0. \tag{D-5}$$

Taking the derivatives of both sides with respect to q gives

$$v'(\frac{q}{d})\frac{1}{d} - [v'(\frac{q}{d})\frac{qd'}{d^2} + \eta'(p)]p' = 0, \tag{D-6}$$

where $p' = \frac{dp}{dq}$. In the above expression, the first term on the left hand side is non-positive and the term in square brackets is non-negative (because $v' < 0, d' < 0$ and $\eta' > 0$). Hence, we must have $p' \leq 0$, i.e., $p^*(q,t)$ is non-increasing in q.

(c) We next show that $p^*(q,t)$ is increasing in t for fixed q. Note that d is non-decreasing in t if p is fixed, so q/d is increasing in t if q and p are fixed. Consider $t_1 < t_2$ and the optimal price for t_1, $p(t_1|q)$. Then, at price $p(t_1|q)$, $v(q/d(p(t_1|q), t_2)) \leq v(q/d(p(t_1|q), t_1)) = \eta(p(t_1|q)$ (by (D-5)). As a result, the optimal price is lower if the temperature increases.

(d) Finally, we show that $p^*(q,t)$ is independent of w. This is evident from the fact that p for a fixed q and t is determined by (D-3) which is independent of w.

It should be noted that part (c) only requires monotonicity of $d(p,t)$ with respect to t. The other results do not require any conditions on how $d(p,t)$ depends on t. □

It is worth mentioning that our method of analysis is similar to that of Song et al. (2008), but in our model, demand is influenced by the weather (temperature) and the deterministic part of the demand function is also slightly more general.

We now turn to the first stage of the retailer's stochastic program. The retailer optimizes q in view of the distribution of T and his optimal pricing strategy discussed above. His expected profit prior to the temperature observation is:

$$\pi_r(q) = \int_{-\infty}^{+\infty} \{p(q,t)d(p(q),t)[1 - \Theta(\frac{q}{d(p(q,t),t)})] - wq\}f(t)dt. \tag{D-7}$$

The first derivative w.r.t. q yields the following first-order condition (after simplification):

$$w \quad = \quad \int_{-\infty}^{+\infty} p(q,t)[1 - G(\frac{q}{d(p(q,t),t)})]f(t)dt \tag{D-8}$$

The following theorem shows that the optimal solution is defined by the first order condition and provides a monotonicity result that is useful in solving the manufacturer's problem.

Theorem 4 $\Pi_r(p(q), q|t)$ *is concave in* q*, i.e.,* $\frac{\partial^2 \Pi_r(p(q),q|t)}{\partial q^2} < 0$ *for all* t*. Therefore, the optimal* q *can be determined from the first order condition in (D-8) for a given* w*. Moreover, the optimal order quantity is strictly decreasing in* w*.*

Proof. We first show that $\Pi_r(q)$ is concave in q. Rearranging and simplifying the terms in (D-6) and making appropriate substitutions, we have that

$$\frac{d(q/d)}{dq} = \frac{d - qd'p'}{d^2} \geq 0, \tag{D-9}$$

because $p' \leq 0$. We use this result below.

Using the shorthand notation $p = p(q, t)$, the first and second derivatives w.r.t. q are:

$$\begin{aligned} \frac{d\Pi_r(q|t)}{dq} &= p[1 - G(\frac{q}{d})] - w \\ \frac{d^2\Pi_r(q|t)}{dq^2} &= p'(1 - G(\frac{q}{d}) - pg(\frac{q}{d})\frac{d - qd'p'}{d^2} > 0 \end{aligned}$$

where the last inequality follows because $\frac{d-qd'p'}{d^2} \geq 0$ and $p' \leq 0$.

Before the temperature is observed, the expected profit of the retailer is $\pi_r(q) = \int_{-\infty}^{\infty} \Pi_r(p, q|t)|_{p=p(q)} f(t)dt$. Therefore,

$$\frac{d^2\pi_r(q)}{dq^2} = \int_{-\infty}^{\infty} \frac{d^2\Pi_r(q|t)}{dq^2} f(t)dt < 0,$$

and thus $\pi_r(q)$ is concave in q.

We now show that the retailer's optimal order quantity is strictly decreasing in w. Suppose to the contrary that $q(w_2) \geq q(w_1)$ where $w_1 < w_2$. By (D-9) we know that $\frac{d(q/d)}{dq} \geq 0$. Hence, $0 \leq 1 - G(\frac{q(w_2)}{d}) \leq 1 - G(\frac{q(w_1)}{d})$. By Theorem 3, we have $p(q(w_1), t) \geq p(q(w_2), t) > 0$. Then, (D-8) leads to $w_1 \geq w_2$, a contradiction.

Finally, note that in the above proof, there are no conditions on the dependence of $d(p, t)$ on t.

□

We now turn to the manufacturer's problem.

The Manufacturer's Problem

Following Lariviere and Porteus (2001) among others, we express the manufacturer's maximization problem in terms of the quantity q. Although this may be unusual, as we showed in Theorem 4, the retailer's first order condition (D-8) defines a strictly monotonic mapping between w and q, so the manufacturer can solve his optimization in the domain of q values (knowing the retailer's optimal pricing strategy $p^*(q,t)$) and then map the solution back to an optimal solution for w.

The manufacturer's profit maximization problem, $\pi_s(w) = wq - cq$, can be expressed as:

$$\begin{aligned} \pi_s(q) &= [\int_{-\infty}^{\infty} p(1 - G(\frac{q}{d}))f(t)dt - c]q \qquad \text{(D-10)} \\ \text{s.t.} & \quad (D-3) \end{aligned}$$

where the constraint implicitly defines the optimal price for each (q,t) pair.

The manufacturer's profit function in (D-10) is not concave in q in general. (See examples in Raz and Porteus (2003) who treat a special case of our model with deterministic temperature.) Nevertheless, due to the strict monotonicity of (D-8), we know that the retailer's solution q^* is unique for fixed w. So, to optimize $\pi_s(q)$, a unidimensional search on w is sufficient, although, in general we cannot solve for w^* in closed form. Alternatively, in view of the above results, we can solve the problem as follows: (1) for each q and t, compute $p^*(q,t)$ using (D-3), with discretization of q and t to any desired precision; (2) conduct a unidimensional search on q to optimize (D-10), with a resultant optimal q^*; and (3) finally, use (D-8) to obtain w^*.

Model without Price Postponement

We argued in the previous subsection that retailers can usually postpone their pricing decisions until better information is known about the weather. We refer to this as the "price late" case. There are instances, however, where the retailer has restricted latitude in changing prices, such as catalog sales, where the catalog must be printed before accurate weather information is available. There are other situations in which the retailer does not have (or wish to commit) the resources to adjust prices in response to the weather. We consider such situations in this subsection. Here, the retailer needs to decide both the price and order quantity before observing the temperature. We refer to this as the "price early" case.

Given w, the retailer's expected profit is

$$\pi_r(p,q) = p\int_{-\infty}^{\infty} d(p,t)[1-\Theta(\frac{q}{d(p,t)})]f(t)dt - wq. \tag{D-11}$$

Taking the first derivatives of the retailer's profit with respect to q and p, respectively and setting them equal to zero, we obtain:

$$\int_{-\infty}^{\infty}(1-G(\frac{q}{d}))f(t)dt = \frac{w}{p}. \tag{D-12}$$

$$\int_{-\infty}^{\infty} d\lambda(\frac{q}{d})[v(\frac{q}{d}) - \eta(p)]f(t)dt = 0 \tag{D-13}$$

Although (D-12) is relatively well-behaved under certain conditions, (D-13) is much more complicated, and the retailer's objective function may not be jointly unimodal in p and q. However, if $d(p,t)$ and the distribution of U have relatively simple analytic forms, it is possible to obtain closed-form solutions. Thus, the results are situation-specific. We present an example later in this section.

We now explain how to incorporate a weather rebate into the "price early" and "price late" models.

Incorporating a Weather Rebate

For the "price late" scenario, the system-optimal order quantity can be obtained by substituting c for w on the left hand side of (D-8). For the "price early" scenario, under the conditions discussed in the previous section, one can similarly substitute c for w in (D-12) to obtain the system-coordinating quantity. Let q_c denote this quantity and q_d denote the decentralized solution obtained as described above. (These values will, in general, differ for the "price early" and "price late" scenarios.) To achieve supply chain coordination, we can again use a rebate of the form (4) in either case. To ensure that the rebate is Pareto- improving, we only need to impose the condition $\Lambda \geq q_d$. (Both of these results are straightforward.) The solution method remains similar to that for the decentralized model.

Example: Power Demand Function and Uniform U

Recall that in the "price-late" model, the optimal solution satisfies $v(q/d) = \eta(p)$ for each t. When $\eta(p)$ is a constant for all p, the solution is significantly simplified: q^*/d is constant. For this reason, in our examples we use a power (constant elasticity) demand function, which satisfies this property. This demand function has a variety of limitations, but it has been used often in the economics literature and is supported by some empirical evidence (e.g., Baltas 2005). We use it here primarily for analytic tractability. We consider the demand function

$$D(p,t,u) = a(t)p^{-\beta}u,\ \beta > 1,$$

where $d(p,t) = a(t)p^{-\beta}$. The value ranges of the two random variables are properly defined so that $u \in [A, +\infty)$ where $A > 0$ and $D \geq 0$ for all possible values of u and t. We assume the distribution of the U is uniform on $[0, 2]$.

Table 2 provides a comparison of price early and price late cases (before and after the retailer observes the weather condition, respectively), with and without a weather rebate. We first make a few general observations. First, the expected retail price (averaged across demand levels as influenced by the temperature) is exactly the same whether the retailer prices early or late assuming that a contract either is, or is not, in force. Second, given the presence or absence of a contract, the retailer's order quantity when he can price late is $D_1^\beta \cdot D_2$ times as large as when he can price early, where $D_1 = E_t[a(t)^{(1/\beta)}]$ and $D_2 = \int_{-\infty}^{\infty} \frac{f(t)}{a(t)}dt$. Using Jensen's inequality and the Cauchy-Schwarz inequality, it can be shown that this factor is greater than 1. Thus, the retailer not only has more units to sell but because of his pricing flexibility, he can extract additional gains when demand is high that more than offset his losses when demand is low. Third, in the absence of a contract, the supply chain profit in the "price late" case is $D_1^\beta \cdot D_2 \cdot \frac{2\beta-1}{\beta}$ times as large as it is in the "price early" case and the magnitude of this ratio is determined in an intricate way by the impact of temperature on demand. But in the presence of a contract, it is $D_1^\beta \cdot D_2$ times as large. Because $\beta > 1$, $(2\beta - 1)/(\beta - 1) > 1$, and hence the proportional increase in the supply chain profit is smaller with a weather rebate than without one. In particular, the introduction of the weather rebate magnifies the order quantity by a multiplicative factor of $(w/c)^\beta$. These changes in the order quantities eventually lead to the supply chain profit being $(\frac{\beta}{\beta-1})^{\beta-1}$ times as large with a contract than without when the retailer prices early and $\frac{\beta}{2\beta-1} \cdot (\frac{\beta}{\beta-1})^{\beta-1}$ times as large with a contract than without when the retailer prices late. Since $\frac{\beta}{2\beta-1}$ is convex and decreasing in $\beta > 1$, with a maximum value of 1 when $\beta = 1$, the (multiplicative) benefit of the contract is smaller when the retailer prices late, as might be expected. Thus, pricing late and the weather rebate are partial substitutes under the assumptions of this model, as might be expected.

Table 2: Summary of Results for Decentralized and Coordinated Solutions

	Wh'le Price	Ret. Price	Order Quant.	SC Profit	Suppl. Share††
Price Early	$\frac{\beta c}{\beta-1}$	$\frac{\beta+1}{\beta-1}w$	$\frac{4(\beta-1)^\beta}{D_2(1+\beta)^{\beta+1}w^\beta}$	$\frac{\beta}{\beta-1}\pi_s^*(E)$	$\frac{\beta-1}{\beta}$
Price Late	$\frac{\beta c}{\beta-1}$	$\frac{\beta+1}{\beta-1}\frac{(a(t))^{\frac{1}{\beta}}}{D_1}w$	$\frac{4((\beta-1)D_1)^\beta}{(1+\beta)^{\beta+1}w^\beta}$	$\frac{2\beta-1}{\beta-1}\pi_s^*(L)$	$\frac{\beta-1}{2\beta-1}$
Coord'd/Price Early	$c^\dagger$	$\frac{\beta+1}{\beta-1}c$	$\frac{4(\beta-1)^\beta}{D_2(1+\beta)^{\beta+1}c^\beta}$	$\frac{\pi_{SC}^*}{D_2}$	$[(\frac{\beta-1}{\beta})^\beta, 1-(\frac{(\beta-1)^{\beta-1}}{\beta^\beta}]$
Coord'd/Price Late	$c^\dagger$	$\frac{\beta+1}{\beta-1}\frac{(a(t))^{\frac{1}{\beta}}}{D_1}c$	$\frac{4((\beta-1)D_1)^\beta}{(1+\beta)^{\beta+1}c^\beta}$	$\pi_{SC}^* D_1^\beta$	$[(\frac{\beta-1}{\beta})^\beta, 1-(\frac{\beta-1}{\beta})^{\beta-1}]$

Note: $D_1 = E_t[a(t)^{(1/\beta)}]$, $D_2 = \int_{-\infty}^{\infty} \frac{f(t)}{a(t)}dt$, $\pi_s^*(E) = \frac{4c^{1-\beta}(\beta-1)^{2\beta-1}}{(\beta+1)^{\beta+1}\beta^\beta D_2}$, $\pi_s^*(L) = \frac{4c^{1-\beta}(\beta-1)^{2\beta-1}D_1^\beta}{(\beta+1)^{\beta+1}\beta^\beta}$, and $\pi_{SC}^* = \frac{4c^{1-\beta}(\beta-1)^{\beta-1}}{(\beta+1)^{\beta+1}}$.

† Although the manufacturer does not actually charge a wholesale price of c, the optimal coordinated solution is the same as if he were charging this price.

†† In the coordinated supply chain, the manufacturer's profit depends on the value of Λ.

The manufacturer's *share* of the supply chain profit is larger in the "price early" setting. However, because $\pi_s(L) = D_1^\beta \cdot D_2 \cdot \pi_s(E)$ and $D_1^\beta \cdot D_2 > 1$, the manufacturer's profit is greater in the "price late" case. Thus, even though the manufacturer earns a larger share of the profit in the "price early" case, he prefers that the retailer employ the "price late" strategy. The retailer, of course, always benefits from using the "price late" strategy.

Additional References

References

[1] Anupindi, R. and L. Jiang. 2005. Capacity Investment under Postponement Strategies, Market Competition and Demand Uncertainty. Working Paper, Ross School of Business, University of Michigan, Ann Arbor, MI.

[2] Baltas, G. 2005. Modelling category demand in retail chains. *Journal of the Operational Research Society.* **56** (11) 1258-1264.

Barrieu, P. and O. Scaillet (2009), A Primer on Weather Derivatives, Chapter 5 in J.A. Filar and A. Haurie (eds.), **Uncertainty and Environmental Decision Making: A Handbook of Research and Best Practice**, Springer, New York, New York.

[3] Biller, S., A. Muriel and Y. Zhang. 2006. Impact of Price Postponement on Capacity and Flexibility Investment Decisions. *Production and Operations Management*, **15** (2), 198-214.

[4] Bish, E. Optimal Investment Strategies for Flexible Resources. 2006. Chapter 4 in **Supply Chain Optimization** (J. Geunes and P. Pardalos, eds.), Springer Series in Optimization. Vol. 98. Springer Sciences+Business Media, Inc., New York.

[5] Bish, E.K. and R. Suwandechochai. 2005. Impact of Demand Substitution/Complementarity and Price/Quantity Postponement on the Optimal Capacity of the Flexible Resource. *Proceedings of the 2005 Manufacturing and Service Operations Management Conference*, Kellogg School of Management, Northwestern University, Evanston, IL.

[6] Bish, E.K. and Q. Wang. 2004. Optimal Investment Strategies for Flexible Resources, Considering Pricing and Correlated Demands. *Operations Research*, **52** (6), 954-964.

[7] Chan, L.M.A., D. Simchi-Levi and J. Swann. 2006. Pricing, Production and Inventory Policies for Manufacturing with Stochastic Demand and Discretionary Sales. *Manufacturing and Service Operations Management*, **8** (2), 149-168.

[8] Chen, F.Y., H. Yan and L. Yao. 2004. A Newsvendor Pricing Game. *IEEE Tran. on Systems, Man, and Cybernetics - Part A: Systems and Humans*, **34**(4), 450-456.

[9] Chod, J. and N. Rudi. 2005. Resource Flexibility with Responsive Pricing. *Operations Research*, **53** (3), 532-548.

Davis, M. (2001), Pricing weather derivatives by marginal value, **Quantitative Finance** 1, 305-308.

[10] Granot, D. and S. Yin. 2008. Price and Order Postponement in a Decentralized Newsvendor Model with Multiplicative and Price-Dependent Demand. *Operations Research* **56** (1) 121-139.

Jewson, S. (2004), Weather derivatives and weather derivatives pricing. Available at stephen-jewson.com.

Jewson, S. and Brix, A. (2005), **Weather Derivative Evaluation**, Cambridge University Press, Cambridge, U.K. (Chapter 3).

[11] Lenk, T. 2008. The Effect of Price Postponement on the Coordination of a Two-Stage Supply Chain Facing Customer Returns. Ph.D. Dissertation, University of Massachuetts, Amherst, Massachusetts.

[12] Petruzzi, N. and M. Dada. 1999. Pricing and the Newsvendor Problem: A Review with Extensions. *Operations Research* **47**(2), 183-194.

[13] Raz, G., and E. Porteus. 2003. Selling to the Newsvendor Under Endogenous Retail Pricing. Working paper, Australian Graduate School of Management, Sydney, Australia.

[14] Ruiz-Benitez, R. (2007). Supply Chain Coordination in the Presence of Customer Returns. Ph.D. Dissertation, University of Massachusetts, Amherst, Massachusetts.

[15] Song, Y., S. Ray and S. Li. 2008. Structural Properties of Buy-back Contracts for Price-setting Newsvendors. *Manufacturing and Service Operations Management*, **10**(1), 1-18.

[16] Tang, C. and B. Tomlin. 2008. How Much Flexibility Does It Take to Mitigate Supply Chain Risk? Chapter 10 in **Supply Chain Risk** (G.A. Zsidisin and B. Ritchie, eds.) Springer International Series in Operations Research and Management Science, Vol. 124. Springer, New York.

[17] van Mieghem, J.A. and M. Dada. 1999. Price versus Production Postponement: Capacity and Competition. *Management Science*, **45** (12), 1631-1649.

Partial List of Notation to Assist the Referees

We provide a partial list of notation (beyond those listed on page 10) to assist the referees.

$K(t^*, q)$	&	$k(t, q)$: functional forms for the weather rebate contract (see (1))
$\pi_r(q)$	:	retailer's expected profit for a given quantity q (see (2))
$\pi_s(q, w)$	:	manufacturer's expected profit for a given quantity-price pair (q, w) (see (3))
$\pi_c(q)$	:	supply chain's expected profit for a given quantity q
$E_{t>t^*}\{\}$	:	expectation taken over all outcomes $t > t^*$
$R(q)$	:	expected net revenue when the quantity q
q_c	:	quantity that maximizes $\pi_c(q)$
q_r	:	quantity that maximizes $\pi_r(q)$
$\pi_r^d(w)$	:	the retailer's profit without a rebate for given q_r
Λ	:	a contract parameter that specifies a threshold order quantity, with lower and upper bounds $\underline{\Lambda}$ and $\overline{\Lambda}$, resp.
t^*	:	another contract parameter, the strike temperature
w_d	:	manufacturer's chosen wholesale price in the absence of a rebate
Λ_w	=	$(w_d - c)q_r/(w - c)$
$I_{\{\cdot\}}$	:	an indicator function; $I = 1$ (or 0) if the argument $\{\cdot\}$ is true (or false)
γ	$\in$	$(0, 1)$, a preset probability value
$g(u)\&G(u)$	:	density and prob. distribution of random variable U
$d(t)$	:	deterministic part of demand function; often abbreviated as d
$\Pi_r(q\|t)$	:	retailer's expected profit when the quantity is q) and temperature is t
$\pi_r(q)$	:	retailer's expected profit before the weather metric is observed and the order quantity is q
$\pi_s(q)$	:	manufacturer's expected profit before the weather metric is observed and the order quantity is q
β	=	the parameter, a constant, for the power function demand model