

Trade Credit and Bankruptcy Risk in Supply Chains: An Experimental Study

Andrew M. Davis^a, Rihuan Huang^b, Kyle Hyndman^c

^aSamuel Curtis Johnson Graduate School of Management, SC Johnson College of Business, Cornell University, Ithaca, NY 14853, USA, adavis@cornell.edu

^bSchool of Management and Economics and Shenzhen Finance Institute, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), Shenzhen 518172, China, huangrihuan@cuhk.edu.cn

^cNaveen Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080, USA, KyleB.Hyndman@utdallas.edu

In practice, supply chain parties often have limited capital, requiring them to seek financing and bear bankruptcy risk. In this paper, we behaviorally investigate a trade credit contract between a supplier and a capital-constrained retailer, the latter of which may face bankruptcy risk. After the supplier proposes a wholesale price, the retailer purchases a quantity through trade credit if its initial capital is insufficient and repays the supplier after demand is realized. If demand is too low, the retailer goes bankrupt. Through a controlled-laboratory experiment with human participants, we investigate how a retailer's exposure to bankruptcy risk, which we vary through its initial capital, affects supply chain decisions and outcomes. We find that the presence of such bankruptcy risk leads to decisions that systematically differ when compared to a setting without bankruptcy risk. Among others, the retailer significantly understocks when exposed to bankruptcy risk and the supplier attempts to offset this effect by offering a lower wholesale price. The resulting effect is that expected profits for the retailer, supplier, and supply chain, are all significantly different than the normative predictions. To account for these observed deviations, we show that a parsimonious reference-dependence model, with the retailer's initial capital as a reference point, organizes the data well. Overall, our work finds that the presence of bankruptcy risk for a retailer significantly alters supply chain decisions and outcomes, relative to an environment without bankruptcy risk.

Key words: Behavioral operations, supply chain contracting, bankruptcy, supply chain finance.

1. Introduction

Insufficient cash flow has always been a critical challenge for firms, especially for small-to-medium enterprises (SMEs). Therefore, firms often seek financing. According to the World Trade Organization (WTO 2019), up to 80% of trade is financed by credit or credit insurance. There are two prominent financing sources: financing provided by a third-party financial institution outside of the supply chain, such as a bank, and financing offered by supply chain members. Bank financing has been popular but can also be costly or even unavailable to SMEs with a limited credit history. As a result, providing financing within supply chains has become increasingly popular in recent years, not only because technology has made such financing much easier (Murphy 2018) but also because

it is frequently more accessible and affordable. Moreover, such supply chain financing can improve performance because the creditor, as a supply chain member, may have better information about the debtors and can make financial and operational decisions in concert (Kouvelis and Zhao 2012, Tang et al. 2018, Tunca and Zhu 2018).

One popular supply chain financing tool is trade credit, where a buyer purchases from a supplier through credit and repays the supplier after a period of time, usually after increasing its cash flow from selling products. A common example of trade credit is 2/10 net 30, where a buyer can either pay within ten days and receive a 2% discount, or pay within 30 days without receiving any discount. The full price can be treated as the discounted price plus the financing cost for an additional 20 days. Trade credit provides liquidity to buyers, allowing them to purchase a desired quantity with relatively low financing cost compared with bank financing. Indeed, it is an important financing source for retailers that are both short of bank credit (Giannetti et al. 2011) and retailers without capital constraints, who may use trade credit to purchase from smaller and weaker suppliers (Murfin and Njoroge 2015). Further, suppliers, as the financing provider, can take advantage of trade credit. By setting appropriate contract terms, suppliers may be able to indirectly affect the retailer's financing decisions or induce higher ordering quantities from retailers. Even suppliers with weak market power can leverage trade credit, primarily as a tool to compete with other suppliers (Lee et al. 2018).

Although trade credit brings benefits to the supply chain, it also comes with downsides. If retailers default due to low demand, which is driven by demand uncertainty, suppliers will incur losses. In this case, retailers who bear limited liability file for bankruptcy, but only partially repay suppliers (the retailer's initial capital and any revenue from demand). Therefore, by providing financing to retailers, such suppliers partially share bankruptcy risk with them. However, compared with a financing provider outside of the supply chain, suppliers can be better informed about a retailer's financial status, by making inferences from the received ordering quantity and better managing the risk of providing financing. In addition, suppliers may also be able to take advantage of retailers' exposure to bankruptcy risk and extract more profits.

Given its potential benefits and risks, trade credit has been well studied analytically. In a simple two-tier supply chain consisting of a supplier and a capital-constrained retailer, the normative theory predicts that a retailer facing higher bankruptcy risk benefits more from limited liability and should set a higher quantity (than if without bankruptcy risk), making the supplier charge a higher wholesale price to extract more profits. While the normative theory assumes rational decision makers, it is unclear whether behavioral biases can drive decisions from these optimal predictions. For example, retailers may be influenced by the endowment effect (Kahneman et al. 1991) and value their initial capital more. Moreover, if behavioral biases do play a role in decision-making,

the magnitude of the biases and how they will change with respect to the degree of bankruptcy risk is unknown. Because these decisions are often made by managers in practice, understanding the problem from a behavioral perspective is essential.

In this paper, we investigate how bankruptcy risk affects supply chain decisions and performance. We consider a two-tier supply chain consisting of a supplier (she) and a capital-constrained retailer (he). The supplier offers a per-unit wholesale price to the retailer, who purchases units from the supplier to satisfy random demand. We consider a trade credit contract. The retailer can purchase a quantity with his limited capital or purchase a quantity through trade credit provided by the supplier. In the latter scenario, the retailer repays the supplier after the demand realization. If demand is too low, the retailer bears limited liability and will transfer any remaining cash to the supplier, including any revenue earned from demand, and go bankrupt. To understand the impact of any potential behavioral biases, we adopt an experimental approach, while using existing analytical results as theoretical benchmarks. In particular, we conduct controlled lab experiments with human participants to test the normative theory. The experiments include three treatments with different levels of the retailer's initial capital, which directly affect the retailer's exposure to bankruptcy risk. With this design we aim to answer the following research questions. First, how does the presence of potential bankruptcy risk affect retailers' order-quantity decisions? Second, how does a trade credit contract, where retailers may face bankruptcy risk, affect suppliers' wholesale price decisions? And third, how do these decisions affect retailer, supplier, and supply chain profits? In answering these questions, if we find that decisions and outcomes do indeed deviate from the normative predictions, we are also interested in identifying what behavioral bias can account for such outcomes.

Our experimental results indicate that retailers, who are capital constrained and thus exposed to higher levels of bankruptcy risk, significantly understock (which reduces their *actual* bankruptcy risk). Anticipating such understocking behavior by retailers, suppliers set significantly lower wholesale prices, in an attempt to induce higher quantities. Conversely, when retailers neglect to face any meaningful capital constraints (i.e., no bankruptcy risk), they set quantities that are slightly higher than the normative predictions, and suppliers set wholesale prices close to the optimal predictions. In short, the presence of bankruptcy risk leads retailers and suppliers to deviate from the normative predictions in systematic ways, relative to a setting without bankruptcy risk. As a consequence, the order of observed net retailer profit (excluding initial capital) and net supplier profit, across different levels of bankruptcy risk, is actually *reversed* relative to the normative predictions. For instance, the normative theory predicts that suppliers should be able to extract higher profits when interacting with capital-constrained retailers (i.e., when retailers face significant bankruptcy risk), but we observe that suppliers actually earn lower profits when interacting with such retailers.

In an effort to account for these deviations under bankruptcy risk, we find that incorporating a reference-dependent retailer can well predict the observed decisions, including wholesale prices and order quantities. Specifically, the retailer's initial capital, as a simple fixed reference point, is sufficient to fit the decisions well. To directly test this behavioral model, which assumes that the supplier is rational but that the retailer is reference dependent, we conduct additional experimental treatments which automate the retailer role (i.e., retailers make decisions in line with the normative theory). In contrast to our original experiments under bankruptcy risk, in these follow-up treatments with automated retailers, we observe that supplier decisions largely conform to the normative predictions. This provides some empirical validation of our behavioral hypothesis that reference-dependent retailers may be driving the deviations from the normative predictions, for both parties, in the presence of bankruptcy risk.

We believe our work contributes to both practitioners and researchers. Regarding the former, we show that retailers' exposure to bankruptcy risk leads them to understock and achieve a lower actual likelihood of bankruptcy. This indicates that retailers neglect to take advantage of their limited liability. It also implies that suppliers, when setting contract terms of trade credit, earn a higher profit when trading with retailers who are not overly capital constrained, which contradicts the normative theory. In terms of our contribution to research, to our knowledge, bankruptcy risk in a two-tier supply chain has not been studied experimentally, despite humans playing an integral role in practice. With human decision makers, we show that supply decisions can significantly differ when retailers face bankruptcy risk. Further, we show that behavioral factors, notably reference dependence, play an important role and can drive decisions away from the rational predictions. Overall, future behavioral research in the broader area of supply-chain finance is critical, especially given that many related topics remain unexplored (e.g., interest rates, hedging, factoring).

2. Literature Review

Inventory decisions with trade credit contracts have been studied for decades in the operations management literature, with analytical modeling and empirical research being the primary methodologies. Overall, Seifert et al. (2013) provide a comprehensive review of the trade credit literature in operations management. Beginning with analytical studies, Haley and Higgins (1973) examine the basic lot-size model with trade credit and find that order quantity and payment time have to be jointly decided to reach optimality. Dada and Hu (2008) find that a capital-constrained newsvendor will borrow funds from a bank and order less than would be ideal if the borrowing cost were not too high. They also show that the channel can be coordinated by a non-linear loan schedule. Lai et al. (2009) study the impact of financial constraints on sharing inventory risk in a supply chain. In their results, the existence of financial constraints (for both a supplier and retailer) will make

the supplier choose to share inventory risk with the retailer, whereas the supplier always prefers taking full inventory risk when without financial constraints. Lee and Rhee (2011) look at trade credit from a supply chain perspective and show that a supplier can fully coordinate the supply chain (with a retailer seeking financing) using both trade credit and markdown allowance.

Trade credit is also studied in more complex settings. For instance, Gupta and Wang (2009) study a multi-period stochastic inventory model and prove that credit terms only affect the value of parameters but not the structure of the optimal policy. Chod (2016) considers a multiple-item setting and finds that financing can distort a retailer's inventory decision. Further, Chod shows that such distortion can be mitigated through supplier financing, as the supplier can observe the actual order quantities prior to setting financing terms. Peura et al. (2017) investigate trade credit when competition between suppliers is considered. They prove that trade credit softens price competition and leads to higher equilibrium profits (i.e., a horizontal benefit), whereas most research focuses on vertical benefits of trade credit. On the other hand, Chod et al. (2019) point out that supplier competition may also lead to a free-rider problem: trade credit extended by a supplier increases the buyer's cash purchase from another supplier. They show that buyers with diverse suppliers receive less trade credit than those who have more concentrated suppliers. In a multi-period model, Luo and Shang (2019) consider a two-level trade credit structure (i.e., from supplier to buyer and from buyer to customer) and show that maximizing working capital at the end of the horizon is equivalent to minimizing the total cost within the horizon. Ning (2022) shows that when there is competition between buyers who have access to external capital, trade credit can incentivize buyers to order more and compete more aggressively, while buyers also benefit from lower wholesale prices.

In addition to supplier financing, other topics receive interest in the literature. For example, in some cases suppliers may be in a financially distressed position and seek financing from retailers. Tunca and Zhu (2018) show that such "buyer financing" can lead to lower financing costs and higher service levels and thus significantly benefit both parties. Tang et al. (2018) compare buyer financing with bank financing and find that the former brings more flexibility to the buyer under symmetric information or even better performance when asymmetric information exists. Cunat (2007) shows that the high implicit interest rate of trade credit is due to the combination of an insurance premium and a default premium. Yang et al. (2021) identify two main roles of trade credit insurance: a) smoothing the supplier's cash flows, and b) monitoring the buyer's continued creditworthiness after contracting.

Along with analytical results, there are also an abundance of empirical findings for trade credit in recent years. Although it is common for powerful suppliers to provide trade credit to small buyers, Murfin and Njoroge (2015) find that even retailers without capital constraints purchase via trade credit from smaller and weaker suppliers. Lee et al. (2018) study how, when competition

exists, trade credit affects firm performance and find that when trade credit offered by suppliers exceeds industry-average levels, buyers' performance is negatively associated with the amount of trade credit. Although it is well documented that diversification can be more profitable in many circumstances, it can also increase the recovery cost for a financially distressed buyer upon default. Taking advantage of a quasi-natural experimental setting provided by a regulatory shock, Aral et al. (2021) study how a distressed buyer's sourcing strategy is affected by bankruptcy risk. They show that capital-constrained buyers are forced to under-diversify compared with those without capital constraints. Astvansh and Jindal (2022) find that provided and received trade credit have different impacts on firm value: the former has a negative direct and a positive indirect effect, and the latter has the opposite effect. They argue that the main reason is the disparate nature of dependence in the supply chain.

Among the existing literature, two papers are closely related to the normative model in our work: Kouvelis and Zhao (2012) and Yang and Birge (2018). Both papers study a supply chain of a supplier and a capital-constrained retailer, with endogenous contract terms. Kouvelis and Zhao (2012) compare bank-financing only and supplier-financing only scenarios. In the former, bank interest is assumed to be competitively priced. In the latter, the supplier provides an early-payment-discount contract (i.e., the retailer can either pay upon delivery of the order at a discounted price or repay the supplier after the demand realization at a regular price). The retailer's financing decision is fully determined by the inventory decision. They show that a capital-constrained retailer always prefers supplier financing, and that the supplier always sets its interest rate at a risk-free rate. In a recent paper, Kouvelis and Zhao (2018) further study the impact of credit ratings on inventory, pricing, and financial terms decisions in the same framework. Yang and Birge (2018) then extend the model by allowing the retailer to take advantage of both financing sources. In particular, given an early-payment-discount contract, the retailer can determine the quantity purchased at the discounted price and at the regular price, respectively. He can also use a bank loan if his cash is insufficient to cover the discounted-price purchase. In this way, the retailer has to decide how much to borrow from the bank and how much of trade credit should be used. They also validate the model using empirical data.

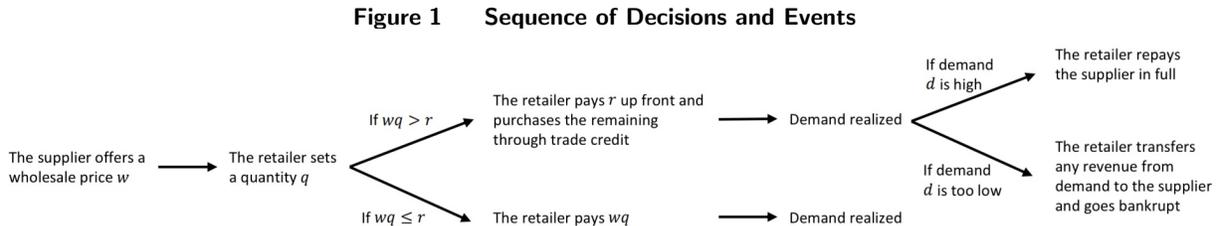
To our knowledge, none of the existing studies in the operations management literature investigate trade credit and bankruptcy risk from a behavioral perspective (Donohue et al. 2019). Therefore, to better understand the role of behavioral factors in a clean environment, in this paper we simplify the model studied in Kouvelis and Zhao (2012) and Yang and Birge (2018) and focus on supplier financing. This allows us to contribute to the literature in two key ways: a) we experimentally test existing analytical models on supply chain finance, which show that supplier financing has an advantage over bank financing (Kouvelis and Zhao 2012, Chod 2016) and b) we examine

how the presence of bankruptcy risk impacts decisions and supply chain outcomes, with human decision makers.

3. Normative Model

Consider a retailer (he) ordering a product from a supplier (she) at wholesale price w per unit, and selling to the market at unit price p . Both the retailer and the supplier are risk-neutral. The supplier produces the product at a per-unit cost c . The retailer is capital-constrained and starts with initial capital r . Before the selling season begins, the retailer determines a quantity q based on the wholesale price w to satisfy a random demand d .

Because the retailer is capital-constrained, he may lack sufficient cash and require financing. Depending on w, q and r , there are two scenarios for the retailer's financing outcome. If $wq \leq r$, the retailer's capital is sufficient to cover his purchasing cost. The retailer pays the supplier wq up front. If $wq > r$, the retailer will need financing from the supplier. Specifically, the retailer pays r to the supplier up front, and the rest after the demand realization. If demand is sufficiently high, the retailer can repay the supplier in full, $wq - r$. If demand is low, the retailer transfers all of his revenue from realized demand to the supplier, and earns zero profit, thus going bankrupt. By bankruptcy, we assume that the retailer bears limited liability. Figure 1 summarizes the sequence of events.



For the normative model introduced in this section, we make some underlying assumptions related to the financial aspects. As such, the model is parsimonious but captures key trade-offs. First, we assume a perfect market without tax or bankruptcy costs. Second, both parties are creditworthy and the retailer will repay any loan obligations to the extent possible. Finally, there is no information asymmetry (i.e., the price, cost, and demand distribution are common knowledge).

3.1. Retailer's Decision

Given a wholesale price w , the retailer determines q (and the choice of financing) to maximize his expected profit:

$$\mathbb{E}[\pi_r] = \mathbb{E}[p \min(q, d) - wq + r]^+. \quad (1)$$

Define $k = (wq - r)^+/p$ as the retailer's bankruptcy threshold, which represents the minimum demand such that the retailer does not go bankrupt (k will be zero if no financing is needed). Equation (1) can be rewritten as

$$\mathbb{E}[\pi_r] = \begin{cases} p\mathbb{E}[\min(d, q)] - p\mathbb{E}[\min(d, k)] & \text{if } wq - r > 0, \\ p\mathbb{E}[\min(d, q)] - wq + r & \text{Otherwise.} \end{cases} \quad (2a)$$

Let $f(\cdot)$ and $F(\cdot)$ be the probability density function (PDF) and cumulative distribution function (CDF) of demand d , respectively. Define $\bar{F}(q) = 1 - F(q)$; the retailer's optimal quantity derived from Equation (2a) and (2b) is

$$q^*(w) = \begin{cases} \bar{F}^{-1}((w/p)\bar{F}(k)) & \text{if } wq - r > 0, \\ \bar{F}^{-1}(w/p) & \text{Otherwise.} \end{cases} \quad (3a)$$

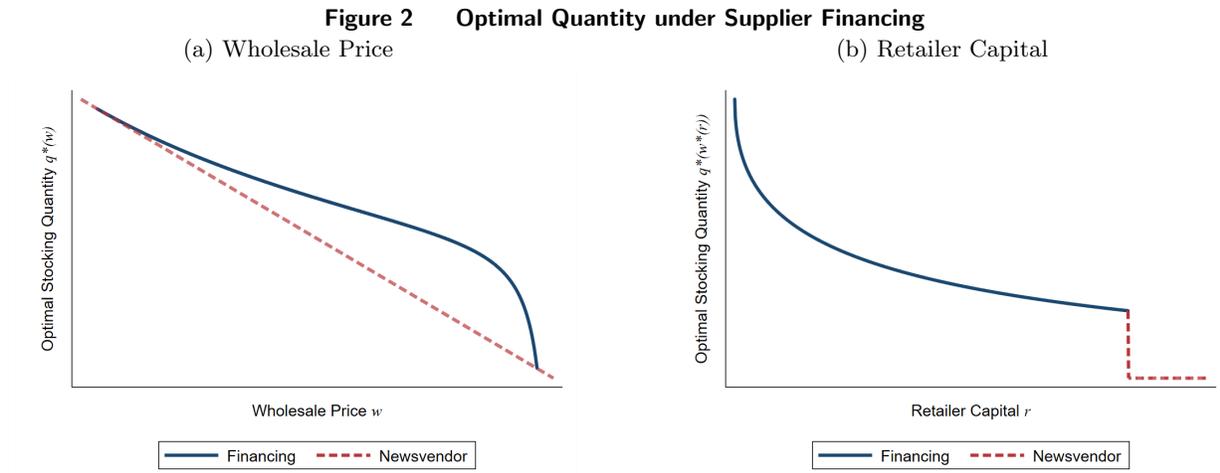
Kouvelis and Zhao (2012) show that there is a one-to-one mapping between w and $q^*(w)$. Let q_l, q_u be the two solutions for

$$q\bar{F}(q) = \frac{r}{p}, \quad (4)$$

with $q_l \leq q_u$, and define $w_l = p\bar{F}(q_l)$, $w_u = p\bar{F}(q_u)$. The "financing region" corresponding to Equation (3a) is equivalent to $w \in (w_u, w_l)$, and the "no-financing region" of Equation (3b) is equivalent to $w \in [0, w_u] \cup [w_l, p]$, respectively.

Equation (3a) shows that the retailer requires financing when the wholesale price is intermediate. If the wholesale price is too low, the unit purchasing cost is so low that the retailer has enough cash to cover the purchased quantity. If the wholesale price is too high, the retailer will order a quantity that is so small that no financing is needed. Denote the interval (w_u, w_l) as the financing region, and observe that it depends on r . In other words, the region is larger when the retailer has less capital and therefore is exposed to higher bankruptcy risk.

Figure 2 shows how q^* varies with respect to w and r . In Figure 2a, supplier financing with limited liability leads to a higher quantity compared with the standard newsvendor quantity. The intuition is that when demand is low, the retailer is actually protected by limited liability (otherwise he would have earned negative profits). Therefore, the retailer will choose a higher quantity in general, which leads to a higher profit (driven by those instances when demand is high). Figure 2b shows how q^* varies with respect to retailer capital r given the corresponding optimal wholesale price. Again, the retailer will order more when he has less capital (i.e., more exposed to bankruptcy risk), as he is more protected by limited liability. The optimal quantity decreases in r and will become the newsvendor quantity when the retailer has a sufficiently high amount of initial capital.



3.2. Supplier's Decision

We assume that the supplier is capital-constrained starting with initial capital r_s and bears limited liability as well. However, Kouvelis and Zhao (2012) show that the supplier's optimal decision is independent of r_s . Given the retailer's optimal quantity, the supplier's expected profit function can be written as

$$\mathbb{E}[\pi_s(w)] = \begin{cases} p\mathbb{E}[\min(d, k)] - cq + r + r_s & \text{if } q_l < q^*(w) < q_u, \\ (w - c)q + r_s & \text{Otherwise.} \end{cases} \quad (5a)$$

$$(5b)$$

To illustrate the supplier's optimal decision, define the following equations:

$$\bar{F}(q) - qf(q) = \frac{c}{p}. \quad (6)$$

$$\frac{\bar{F}(q) - qf(q)}{1 - (wqf(k))/(p\bar{F}(k))} = \frac{c}{p}, \quad (7)$$

Let \bar{q} and \hat{q} be the solutions to Equations (6) and (7), respectively. Kouvelis and Zhao (2012) prove that the supplier's optimal wholesale price w^* and corresponding retailer's decision are as follows:

- (No financing) If $\bar{q} \leq q_l$, $w^* = p\bar{F}(\bar{q})$ and the retailer does not need financing.
- (Financing) If $\bar{q} > q_l$, w^* is derived from $\hat{q} = \bar{F}^{-1}(w/p)\bar{F}(k)$ and the retailer needs financing.

To intuitively understand these results, recall that there is a one-to-one mapping between w and $q^*(w)$. Therefore, the supplier can indirectly determine the retailer's financing decision by setting a wholesale price. The “no financing” case is when it is optimal for the supplier to not finance the retailer. The retailer's optimal decision is \bar{q} , which is the standard newsvendor quantity and independent of the retailer's capital r . In this case, \bar{q} is smaller than q_l and is not in the financing region (q_u, q_l) . The supplier's optimal wholesale price is $w^* = p\bar{F}(\bar{q})$. Such a “no financing” case occurs when the retailer has a high level of initial capital. When the retailer has a low level of initial capital, the financing region becomes larger and $\bar{q} > q_l$. In this case, $w = p\bar{F}(\bar{q})$ cannot be

optimal, and the supplier will set the price w^* following the “financing” case. We refer interested readers to Kouvelis and Zhao (2012) for a general solution.

Figure 3 Optimal Bankruptcy Threshold with Respect to Wholesale Price

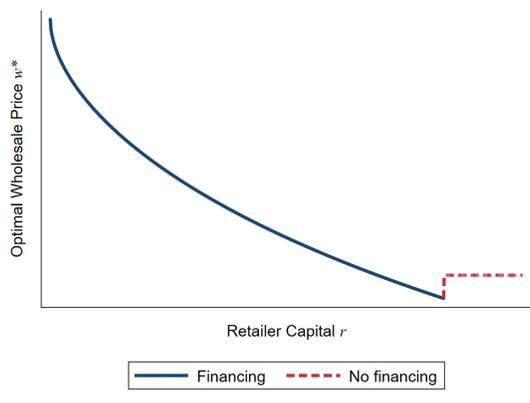


Assuming the retailer optimally responds to the supplier’s wholesale price, the retailer’s degree of bankruptcy risk does not vary linearly with respect to the wholesale price. Figure 3 shows that the retailer’s bankruptcy threshold $k(q^*(w))$ exhibits an inverted U-shape in the financing region. Note that this $k(q^*(w))$ can be regarded as the bankruptcy risk that the supplier offers to the retailer. In Figure 3, a certain degree of bankruptcy risk can be achieved through a low or a high wholesale price. This is because the retailer will order more given a lower wholesale price, which results in a similar degree of bankruptcy risk. In other words, the supplier may not be able to reduce the retailer’s risk by offering a lower price. We discuss this further when we propose our behavioral hypotheses in Section 4.2.

Turning to how w^* varies with respect to r , Figure 4 provides a numerical example. Qualitatively, when in the financing case, w^* decreases in the retailer’s capital r . In other words, because a retailer with less capital prefers to set a higher quantity, the supplier will set a higher wholesale price to extract more profit. Although, it is possible for the financing w^* to be lower than the no-financing $w^* = p\bar{F}(\bar{q})$ (when the retailer’s initial capital is relatively high, but not so high that the supplier prefers not to finance him).

While theory predicts that a retailer exposed to higher bankruptcy risk should take advantage of limited liability and order a higher quantity (relative to lower bankruptcy risk), it is unclear whether human decision makers would behave in this way. For instance, a human retailer with high potential bankruptcy risk may fail to realize the benefits of limited liability and understock to reduce the actual risk. Moreover, if behavioral biases are driving decisions, how the impact of any biases varies with respect to the degree of bankruptcy risk is unknown (along with their effect on supply chain performance). To investigate the impact of behavioral factors, in the next section we design controlled laboratory experiments with human participants and discuss behavioral hypotheses.

Figure 4 Optimal Wholesale Price with Respect to Retailer Capital



4. Behavioral Experiment

4.1. Experimental Design

We conduct controlled laboratory experiments to test the theory introduced in Section 3. As a baseline, we first include a treatment where the retailer has no bankruptcy risk predicted by theory, denoted as NR. Note that NR is essentially a standard two-tier supply chain wholesale price contract scenario under optimal decisions (bankruptcy risk is only possible when retailers over-order by extreme amounts). Because we are mainly interested in those settings where the retailer is exposed to bankruptcy risk, we then consider two additional treatments, low bankruptcy risk (LR) and high bankruptcy risk (HR), by varying the retailer’s initial capital r .

Turning to our parameters, we set the unit retailer selling price $p = 30$ and the supplier production cost $c = 10$. The retailer faces an integer demand drawn from a uniform distribution between 0 and 100. Regarding the supplier’s initial capital, we set $r_s = 400$ across all treatments. The reason is twofold: a) the normative theory shows that the supplier’s optimal decision is unaffected by r_s , and b) $r_s = 400$ makes supplier bankruptcy rarely happen,¹ which means that we can focus on the retailer’s bankruptcy risk as a key treatment variable. Note that if a supplier goes bankrupt in the game, she is also protected by limited liability (i.e., earns at least a profit of zero).

We set r as 700 (NR), 400 (LR), and 100 (HR). The normative predictions of all three treatments are shown in Table 1. Understanding that the retailer is more protected by limited liability when he faces a higher potential bankruptcy risk, the supplier chooses a higher wholesale price to extract more profit, resulting in a predicted wholesale price of 20 in NR, 21.47 in LR, and 25.87 in HR. The predicted order quantity exhibits a similar pattern, 33.33 in NR, 38.73 in LR, and 42.49 in HR. The retailer’s exposure to bankruptcy risk is reflected by the bankruptcy threshold k . In NR, the retailer has no bankruptcy risk under the optimal decisions ($k = 0$), while it is $k = 14.38$ in LR

¹ The average supplier bankruptcy rate in our data is 0.39%.

and $k = 33.31$ in HR. Last, we again note that the bankruptcy threshold in NR is 0.00, but it is possible that the retailer may face some bankruptcy risk if they significantly overstock.

Regarding predicted profits, supplier expected profit increases from 733.33 in NR, to 813.16 in LR, to 907.87 in HR, as the supplier is better able to extract profit from the retailer as he becomes more exposed to bankruptcy risk (i.e., as r decreases). On the other hand, the retailer's predicted expected profit decreases from NR to HR, for both total profit and net profit (excluding initial capital, reported in square brackets). For example, the predicted net retailer profit is 166.67 in NR, 136.44 in LR, and 71.13 in HR. This is due to the increasing wholesale price set by the supplier. Finally, the last row of Table 1 shows predicted supply chain profit and net profit. Note that we do not show predicted supply chain efficiency because the benchmark varies from treatment to treatment as total initial capital $r + r_s$ varies, which may lead to inconsistent comparisons.

Table 1 Experimental Parameter Settings and Predictions

Treatment	NR	LR	HR
Retailer Initial Capital r	700	400	100
Wholesale Price w	20.00	21.47	25.87
Quantity q	33.33	38.73	42.49
Bankruptcy Threshold k	0.00	14.38	33.31
Supplier Expected Profit	733.33 [333.33]	813.16 [413.16]	907.87 [507.87]
Retailer Expected Profit	866.67 [166.67]	536.44 [136.44]	171.13 [71.13]
Supply Chain Expected Profit	1600.00 [500.00]	1349.60 [549.60]	1079.00 [579.00]

Note: the supplier's initial capital is $r_s = 400$ across all treatments. Net expected profits (excluding any initial capital) are reported in square brackets.

Each treatment included 60 participants recruited from a large university where cash was the only incentive offered. Each participant was randomly assigned a role and stayed in that role throughout the game. In the experiment, suppliers and retailers were placed in cohorts of six (three of each type), which participants were unaware of. Within each cohort, in each round, pairs of one retailer and one supplier were randomly formed. Both parties were provided with decision support. Specifically, participants could use a sliding bar to test their decisions and observe a plot of the realized profit of both parties for each possible demand realization. In addition, both parties were shown the probability of retailer bankruptcy. For suppliers, the retailer's realized profits and bankruptcy probability were calculated by assuming an optimal quantity response (which participants were aware of). Note that the word "bankruptcy" was not used in the experiments.

Instead, we used the phrases the “possibility that the retailer cannot pay you in full” for suppliers and the “possibility that you cannot repay the supplier in full” for retailers. Finally, retailers were provided with a rejection button, if they disliked a received offer. If a contract was rejected, both parties kept their initial capital.

Our experiment was implemented through oTree (Chen et al. 2016).² Before each session started, a researcher read the instructions aloud and answered any questions. Participants were then required to answer several multiple-choice comprehension questions about the game. They received cash based on their profits from a randomly chosen round in the game plus a \$7 show-up fee. Average earnings were \$22.34 across all treatments. Each session lasted for 70 minutes on average.

4.2. Behavioral Hypotheses

Although the normative theory provides predictions for a rational decision maker, existing experimental evidence suggests that human participants can exhibit behavioral biases and deviate from rationality. In this subsection, we discuss three behavioral hypotheses related to the order quantity, wholesale price, and bankruptcy risk, followed by how such potential deviations can affect profit outcomes. Also, recall that our baseline NR treatment resembles a classic two-tier supply chain with a wholesale price contract, where the retailer makes the order quantity decision. In such a setting, it has been shown that decisions generally conform to the normative predictions (Davis et al. 2014). Therefore, we focus our hypotheses on the two treatments where retailers are exposed to bankruptcy risk, LR and HR, but will indeed make comparisons across all three treatments in the subsequent results section.

Beginning with the order quantity, because retailers are endowed with initial capital r in each round, they may exhibit the endowment effect and over-value the initial capital, thus becoming loss averse (Kahneman et al. 1991). In other words, retailers will attempt to face an actual bankruptcy risk that is lower than the normative theory predicts. This can be accomplished by setting an order quantity that is below the normative prediction. Moreover, recall that a retailer who is exposed to more bankruptcy risk (e.g., HR) is actually more protected by limited liability and should order a higher quantity, compared to a setting where the retailer is exposed to a relatively lower bankruptcy risk (e.g., LR, recall Figure 2b). A human retailer may fail to recognize the benefits of limited liability. As a result, retailers may have less incentive to order the high predicted quantity in HR, and to a lesser extent LR (and loss aversion should be more salient when the bankruptcy risk is high). Therefore, we offer the following Hypothesis 1:

HYPOTHESIS 1. Retailers in LR and HR will set lower quantities than the normative theory predicts, in order to achieve a lower bankruptcy risk. The magnitude of deviation will be larger in HR than in LR.

² Because of the COVID-19 pandemic, we ran all sessions synchronously online through Zoom.

The supplier provides financing to the retailer and, when retailers are capital constrained, may anticipate understocking behavior. If so, they will set a more generous wholesale price to induce a higher quantity. We also expect the magnitude of deviation to be larger in HR than in LR, for three reasons. First, as stated in Hypothesis 1, we expect to observe more understocking in HR. Suppliers will need to lower wholesale prices significantly to partially offset this effect and achieve higher profits. Second, the retailer begins at a significant disadvantageous position in HR, with less initial capital, and the supplier may prefer a more fair outcome (e.g., Cui et al. 2007, Kalkanici et al. 2014). Third, the predicted wholesale price in HR is relatively extreme (25.87), and a human supplier may be reluctant to set extreme values based on past lab evidence (e.g., Gurnani et al. 2014). Thus, we have Hypothesis 2:

HYPOTHESIS 2. Suppliers in LR and HR will set lower wholesale prices than the normative theory predicts. The magnitude of deviation will be larger in HR than in LR.

Although we expect to see lower wholesale prices in LR and HR, this does not necessarily lead to less actual bankruptcy risk for the retailer. Recall that Figure 3 shows that $k(q^*(w))$ exhibits an inverted-U shape with respect to w . In LR and HR, the predicted wholesale prices are on the right side of the peak, and a lower wholesale price can lead to even greater bankruptcy risk. Therefore, suppliers may be facing a trade-off between offering between a lower wholesale price with more bankruptcy risk and a higher price but with less bankruptcy risk. While it is unclear which will be the case, it is unlikely that suppliers offer a higher degree of bankruptcy risk than the normative predictions. Here, we provide Hypothesis 3:

HYPOTHESIS 3. Assuming that retailers set order quantities optimally, the bankruptcy threshold in LR and HR will not be higher than the normative theory predicts.

Even if we can anticipate these deviations in decisions directionally, a controlled experiment is still valuable. For instance, how the resultant profit will change is not fully clear. Even if both Hypotheses 1 and 2 are supported, we can anticipate that supplier profit will be lower than predicted. However, how the retailer's profit will change is unknown; although the retailer can be hurt by his suboptimal ordering, he may also benefit from the lower wholesale price set by the supplier. Attempting to extrapolate these effects to the total supply chain profit is even more challenging. Therefore, we aim to investigate these questions by conducting a controlled laboratory experiment.

5. Results

We first present observed decisions and any deviations, and then show how these decisions translate into profits. Recall that retailers were allowed to reject a supplier's offer if he found it unfavorable. From the experiment, the rejection rates for the three treatments are 8.67% (HR), 9.00% (LR),

and 6.00% (NR). Given the similar rates, we exclude the rejected data in the following analyses. Unless otherwise noted, we use regressions with random effects for hypothesis tests (Hyndman and Embrey 2019).

Before showing the results, we first check whether any significant learning took place. Overall, there is some learning for the wholesale price but not for the order quantity. By comparing observed wholesale prices of the first half and the second half of the experiment (rejected data excluded), suppliers set slightly higher wholesale prices in the second half in NR (on average 20.36 vs. 21.07, $p < 0.01$) and LR (18.84 vs. 19.30, $p < 0.01$), but not in HR (19.80 vs. 20.00). Although the difference in LR and HR is significant, the absolute difference is less than 1. Turning to order quantities, we do not observe significant learning in terms of the quantity deviation ($q - q^*(w)$, observed quantity minus conditionally optimal quantity) in any treatment (4.46 vs. 5.21 in NR, -11.51 vs. -11.34 in LR, -15.82 vs. -14.71 in HR). Given that the learning is both slight and inconsistent across treatments, we keep the full data for all analyses (rejected data are still excluded).

5.1. Decisions

Table 2 presents the average observed decisions (left side) and the theoretical predictions (right side). For the latter, we provide order quantity and bankruptcy threshold predictions conditioning on any previous decisions. Unconditional predictions are in square brackets. Beginning with Hypothesis 1, we first look at the observed quantity and find opposite deviations in HR and LR, versus NR, relative to the normative predictions. Specifically, retailers significantly understock when they are exposed to bankruptcy risk (HR 40.55 observed vs. 55.76 predicted, and LR 35.14 vs. 46.46, both $p < 0.01$), but weakly significantly overstock when they are not exposed to bankruptcy risk (NR 36.38 observed vs. 31.46 predicted, $p < 0.1$). As a result, the actual bankruptcy threshold is significantly lower than predicted in HR and LR (both $p < 0.01$). Also, recall that there should not be any bankruptcy risk in NR, but because of the slight overstocking behavior of retailers in NR, the observed bankruptcy risk is actually higher than zero ($p < 0.01$). Coming back to Hypothesis 1, although the absolute deviation is larger in HR than in LR, we do not find a significant difference in terms of the relative deviation (i.e., $q/q^*(w)$). Therefore, Hypothesis 1 is partially supported. To summarize, retailers in HR and LR significantly understock (and thus face lower actual bankruptcy risk), whereas retailers in NR tend to slightly overstock (and thus face some bankruptcy risk). Therefore, we have the first result:

RESULT 1. Compared with the normative predictions, when retailers are exposed to bankruptcy risk (HR and LR), they set order quantities that are too low, which translates into lower actual bankruptcy risk. Conversely, when retailers are not exposed to bankruptcy risk (NR), they set order quantities that are too high, which translates into an actual non-zero bankruptcy risk.

Table 2 Observed Decisions and Theoretical Predictions

	Observed			Predicted		
	NR	LR	HR	NR	LR	HR
Wholesale Price	20.70 [†] (0.34)	19.05 [‡] (0.28)	19.91 [‡] (0.47)	20.00	21.47	25.87
Quantity	36.38* (2.68)	35.14 [‡] (2.17)	40.55 [‡] (1.77)	31.46 (0.68) [33.33]	46.46 (0.60) [38.73]	55.76 (0.68) [42.49]
Bankruptcy Threshold	5.53 [‡] (1.22)	8.98 [‡] (1.25)	22.66 [‡] (1.11)	0.30 (0.05) [0.00 [‡]]	15.48 (0.13) [14.38 [‡]]	32.72 (0.32) [33.31]

Note: Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Predicted quantities and bankruptcy thresholds are conditioning on observed wholesale prices. Unconditional normative predictions, when applicable, are reported in square brackets. Significance of regressions with random effects compared with observed versus conditional normative predictions are given by [‡] $p < 0.01$, [†] $p < 0.05$ and * $p < 0.1$.

Turning to wholesale prices, the average observed value is significantly lower than predicted in HR, 19.91 versus 25.87, and in LR, 19.05 versus 21.47 (both $p < 0.01$), but slightly higher in NR, 20.70 versus 20.00 ($p < 0.05$). In other words, similar to retailers, suppliers exhibit a different direction of deviation in response to retailers' degree of bankruptcy risk. Hypothesis 2 is fully supported. Although suppliers in NR set prices slightly higher than the normative prediction, note that the absolute difference is relatively small. Combining these observations around wholesale prices yields the following result:

RESULT 2. *Compared with the normative predictions, when retailers are exposed to bankruptcy risk (HR and LR), suppliers set wholesale prices which are too low. When retailers are not exposed to bankruptcy risk, suppliers set wholesale prices slightly higher than optimal.*

Table 2 also indicates that suppliers do not offer wholesale prices that, in theory, lead to lower bankruptcy risk for retailers, relative to the normative predictions. In particular, the conditional predictions of the bankruptcy threshold at the bottom right of Table 2 are calculated based on observed wholesale prices and conditionally optimal quantities. Therefore, they can be regarded as the degree of bankruptcy risk offered by suppliers. By comparing conditional bankruptcy thresholds with unconditional ones, we find that suppliers do not offer significantly lower degrees of risk. In LR, the average bankruptcy threshold offered is 15.48, which is higher than the theoretical prediction, 14.38 ($p < 0.01$). Similarly, in NR, the average bankruptcy threshold is 0.30, higher than the predicted value, 0.00 ($p < 0.01$). Such findings indicate that Hypothesis 3 is not supported. Although, we emphasize that the absolute differences are relatively small. Thus we have Result 3.

RESULT 3. *Compared with the normative predictions, suppliers do not offer wholesale prices that result in significantly lower levels of bankruptcy risk to retailers.*

Table 3 Regressions of Deviation of Retailer Decisions

Dependent variable	$q - q^*(w)$			$k - k^*(w)$		
	NR	LR	HR	NR	LR	HR
Wholesale Price	0.813 [‡] (0.194)	-0.038 (0.273)	-0.430 [†] (0.214)	0.157 (0.116)	-0.241 (0.157)	-0.734 [‡] (0.146)
Lagged Demand	0.009 (0.019)	0.033 (0.022)	0.065 [†] (0.026)	0.012 (0.012)	0.021* (0.013)	0.043 [†] (0.018)
Lagged Bankruptcy	3.412 (2.526)	-0.578 (2.154)	3.187* (1.768)	7.545 [‡] (1.542)	-0.323 (1.234)	1.985 (1.208)
Constant	-12.483 [‡] (4.319)	-12.438 [†] (5.433)	-10.613 [†] (4.749)	1.147 (2.463)	-3.171 (3.124)	1.841 (3.223)
N	538	520	524			

Note: regressions are run with random effects. Standard deviations of estimated coefficients are reported in parentheses. “Lagged bankruptcy” is a binary outcome of retailer bankruptcy in the previous round.

Next we explore these observed decisions in more depth. Specifically, one might be interested in knowing whether retailers react to wholesale prices differently, across treatments. To address this, we conduct further regression analysis of retailer decisions, presented in Table 3. The three leftmost columns show treatment-specific results with the observed quantity deviation, conditioning on the observed wholesale price $q - q^*(w)$, as the dependent variable. Estimated coefficients indicate that retailers in NR tend to overstock when facing high wholesale prices whereas retailers in HR understock. There are two possible reasons for this deviation pattern: to reduce the difference in realized profits between the two parties or to reduce the bankruptcy risk. Results in the right side of the table show that only the first reason applies to NR and that both apply to HR. The dependent variable of the right three columns is the observed deviation in the bankruptcy threshold $k - k^*(w)$. The coefficient of the wholesale price is not significant in NR, indicating that retailers change little of their bankruptcy risk by responding differently to wholesale prices. However, the coefficient is significant and negative in HR, meaning that retailers do choose lower bankruptcy risk when suppliers set higher wholesale prices. Regarding the coefficients of lagged demand, retailers with a high bankruptcy risk seem to be more sensitive to realized demand, which is understandable. Finally, we do not find that retailers deviated more in quantity or tended to reject offers more after experiencing bankruptcy in any treatment (not depicted). In addition, there is no significant evidence that suppliers set wholesale prices differently after their paired retailers went bankrupt in the previous round.

5.2. Profits

Next we examine how the observed decisions translate into profits, first comparing to the normative predictions, and second, comparing across treatments. To this end, Table 4 shows the observed

average profits on the left side and the normative predictions on the right side. Beginning with suppliers, they earn profits that are lower than predicted, when retailers are exposed to bankruptcy risk (HR and LR, 682.35 vs. 907.87 and 672.41 vs. 813.16, respectively, both $p < 0.01$). Given our earlier results on decisions, this is expected, as we observe under-pricing and understocking in HR and LR. As for NR, although suppliers achieve slightly higher profits than predicted on average, 757.49 versus 733.33, the difference is not significant.

Table 4 Observed Profits and Theoretical Predictions

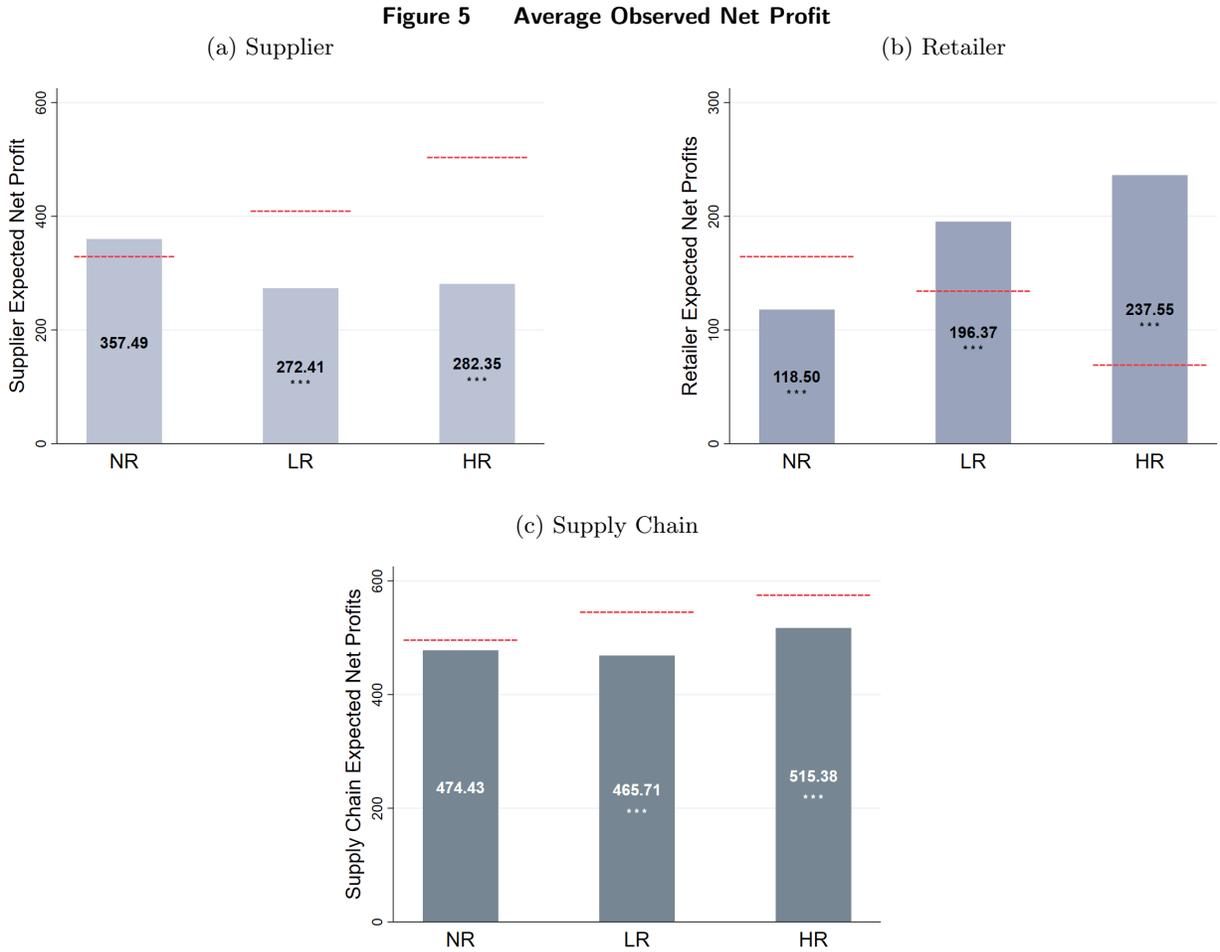
	Observed			Predicted		
	NR	LR	HR	NR	LR	HR
Supplier Profit	757.49 (19.75)	672.41 [‡] (11.24)	682.35 [‡] (9.29)	733.33	813.16	907.87
Retailer Profit	818.50 [‡] (11.07)	596.37 [‡] (7.68)	337.55 [‡] (9.54)	866.67	536.44	171.13
Supply Chain Profit	1574.43 (16.72)	1265.71 [‡] (13.75)	1015.38 [‡] (14.07)	1600.00	1349.60	1079.00

Note: Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Significance of regressions with random effects comparing observed versus conditional normative predictions are given by [‡] $p < 0.01$.

In contrast to supplier profit, we observe higher average retailer profit in HR and LR, relative to the normative predictions. Specifically, the observed average retailer profit in HR is 337.55, much higher than its theoretical prediction, 171.13 ($p < 0.01$). Similarly, the average retailer profit in LR is 596.37 versus the 536.44 prediction ($p < 0.01$). Conversely, in NR, retailers achieve an average profit of 818.50, which is actually lower than predicted, 866.67 ($p < 0.01$).

The last row of Table 4 shows the observed average supply chain profit. Although retailers achieve higher profits when exposed to bankruptcy risk (HR and LR), the total profits are still lower than the normative predictions (1015.38 vs. 1079.00 in HR and 1265.71 vs. 1349.60 in LR, both $p < 0.01$), due to insufficient order quantities. Last, we do not observe significantly different supply chain profits in NR, versus the normative prediction.

Figure 5 presents average net profits and predictions, which allows us to further understand the impact of deviations and to directly compare across treatments. Beginning with supplier net profit, NR achieves higher supplier profit than HR and LR (both $p < 0.05$), which runs counter to the normative predictions. Also, there is no significant difference between supplier profit in HR and LR. Turning to retailer net profit, we observe an opposite trend compared to the normative predictions. In particular, retailer net profit in HR is higher than that in LR ($p < 0.05$), which is then higher than NR ($p < 0.01$). In other words, HR > LR > NR whereas theory predicts the reverse,



Note: Normative predictions are represented by horizontal red dashed lines. Significance of regressions with random effects comparing observed with normative predictions are given by *** $p < 0.01$.

$HR < LR < NR$. Recall that we did not have a formal behavioral hypothesis around profits, but can now provide the following pertinent result:

RESULT 4. *In contrast to the normative predictions, suppliers earn a significantly higher profit when retailers are not exposed to bankruptcy risk ($NR > LR$ and $NR > HR$), and retailers earn a significantly higher profit when they are exposed to more bankruptcy risk ($HR > LR > NR$). Further, supply chain profits are significantly below the normative predictions when retailers are exposed to bankruptcy risk (HR and LR).*

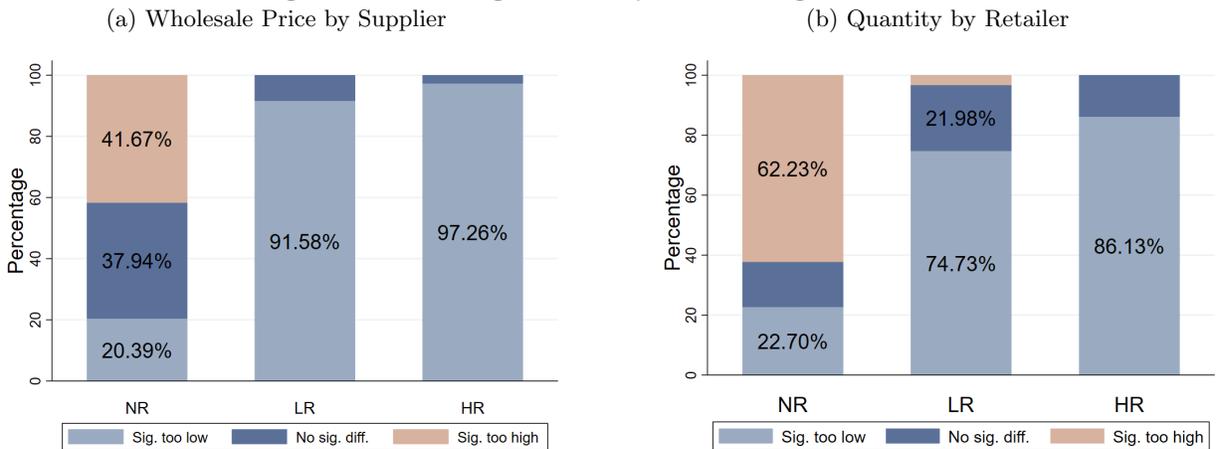
To understand the behavioral mechanism behind these outcomes, we propose a plausible behavioral model that can capture the observed decision deviations outcomes in Section 6.

5.3. Heterogeneity Analysis

In this subsection we conduct a heterogeneity analysis to examine whether our results are driven by a subset of participants. Figure 6 classifies participants into three categories based on whether

they consistently set lower (the bottom, light-blue portions in the figure) or higher (the top, light-orange portions) values than optimal, or did not significantly deviate in one direction (the middle, dark-blue portions). Thus, we compare participants' actual decisions with (conditionally) optimal decisions in all rounds,³ with a Wilcoxon signed-rank test (because of the limited sample size⁴) and a significance level of 5%. Beginning with the wholesale price decision in Figure 6a, the majority of suppliers consistently set the wholesale price lower in LR (91.58%) and HR (97.26%), consistent with the average observed results in Table 2. However, suppliers in NR exhibit mixed behavior, 41.67% of them choose to over-price, 20.39% prefer lower prices, and 37.94% stand in between. Turning to the quantity decision of retailers in Figure 6b, results are mostly consistent with the previous averages. Specifically, 62.23% of retailers in NR overstock, while 74.73% in LR and 86.13% in HR understock (and thus a lower bankruptcy risk).

Figure 6 Percentages of Participants Deviating in Decisions



Note: Classification of a participant is based on Wilcoxon signed-rank tests between observed decisions and (conditional) optimal decisions (at 5% level). Percentages below 20% are omitted because of limited space.

Taking a closer look at the three subgroups of suppliers in NR, the average wholesale prices of the “under-pricing” group, the “in between” group, and the “over-pricing” group are 18.29, 20.15, and 22.28, respectively. The absolute deviation from the optimal wholesale price, 20, is similar for both the “under-” and the “over-pricing” groups, which can account for the overall average price being just slightly higher than the normative prediction. Although deviations are less consistent in this case, we note that the “over-pricing” group is still the most common type (41.67%) driving the average decision. Overall, this analysis indicates that our aggregate results are not driven by a small group of outliers or individuals.

³ For the wholesale price, we compare a supplier's actual wholesale price and the theoretically optimal price (e.g., 25.87 in HR). For the quantity, we compare a retailer's actual quantity and the optimal quantity conditioning on the observed wholesale price in each round.

⁴ The sample size for each test is 20 before excluding rejected data.

6. Behavioral Model: Reference Dependence

So far we have shown that observed decisions significantly deviate from the normative theory. To understand what can account for these deviations, we now investigate a reference-dependence model. Reference dependence has been observed in various operational contexts (e.g., Ho et al. 2010, Baron et al. 2015, Tereyağoğlu et al. 2018). We first introduce a general framework of reference dependence which follows Uppari and Hasija (2019) and subsequently discuss potential reference points. We then fit these alternative reference points to the data, along with the normative theory, to determine the best-performing model. We show that considering only reference-dependent retailers, but not suppliers, based on a simple reference point, is sufficient to capture both parties' observed decisions in our experiments. Thus, our behavioral model is parsimonious and predictive.

Following the notation of Uppari and Hasija (2019), a reference-dependent decision maker compares its profit to a reference point \mathcal{P} and gains positive (negative) utility if its profit is higher (lower) than \mathcal{P} . Based on whether \mathcal{P} is dependent on demand d or any decision, reference points can be categorized by a) a fixed reference point (FRP) that is independent of any value, b) a prospect reference point (PRP) that depends on decisions, and c) a stochastic reference point (SRP) that depends on both realized demand and decisions. Define $\mathcal{D}_> = \{d \in [0, a] : \pi(d) > \mathcal{P}\}$ as the domain of demand that leads to a higher realized profit than the reference point (the gains domain). Similarly, let $\mathcal{D}_< = \{d \in [0, a] : \pi(d) < \mathcal{P}\}$ be the losses domain. A general reference-dependent utility function is

$$\mathbb{E}[\pi] + \left[\eta \int_{x \in \mathcal{X}_>} (\pi(x) - \mathcal{P}(x)) dF(x) - \lambda \int_{x \in \mathcal{X}_<} (\mathcal{P}(x) - \pi(x)) dF(x) \right]. \quad (8)$$

The utility in Equation (8) consists of expected profit (also called the consumption utility) and the gain-loss utility. The parameters η and λ are the psychological weights of gains and losses, respectively. We assume η and λ to be non-negative unless a negative parameter is meaningful.

We first investigate the quantity decision of the retailer and then examine the wholesale price decision of the supplier. Based on the observed quantity deviation, it is possible that the supplier assumes the retailer is reference-dependent rather than rational. Therefore, we need to first understand potential reference points for the retailer. For an introduced reference point, we only discuss the main ideas in this section; we present the detailed utility function in Appendix A.

To evaluate a reference point, we fit the observed decisions using maximum-likelihood estimation (MLE). In the estimation, we use truncated normal distributions for both the quantity and wholesale price decisions, with (conditionally) optimal values as the means and estimate the standard deviations. For the quantity, the distribution is truncated at lower bound 0 and upper bound 100 of demand. For the wholesale price, the normal distribution is truncated at unit production cost

$c = 10$ and unit selling price $p = 30$. The PDF of a truncated normal distribution $\varphi(x; \mu, \sigma, a, b)$ is defined by

$$\varphi(x; \mu, \sigma, a, b) = \frac{\phi\left(\frac{x-\mu}{\sigma}\right)}{\sigma \left(\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)\right)}. \quad (9)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF of the standard normal distribution, respectively.

We split the data into a training set (67%) and a testing set (33%).⁵ In each treatment, recall that there are 30 participants playing one role (either supplier or retailer). We assign 20 of 30 to the training set and the remaining 10 to the testing set. We first estimate each model using the training set. We then calculate the log-likelihood (LL) of the estimated parameters on the testing set and compare across models. Let I be retailers in the testing test (indexed by i) and J be suppliers in the testing test (indexed by j), respectively. For retailer i or supplier j , some rounds are excluded because of rejection. Therefore, we use T_i and T_j to denote the set of rounds that are included in the estimation, with t being the index. The total LL functions for retailers and suppliers are shown below.

$$LL_r(\eta_r, \lambda_r, \sigma_r) = \prod_{i \in I} \prod_{t \in T_i} \varphi(q_{it}; \tilde{q}(w_{it}, \eta_r, \lambda_r), \sigma_r, 0, 100), \quad (10)$$

$$LL_s(\eta_s, \lambda_s, \sigma_w) = \prod_{j \in J} \prod_{t \in T_j} \varphi(w_{jt}; \tilde{w}(\eta_s, \lambda_s), \sigma_w, 10, 30). \quad (11)$$

We use subscript r for all retailer-related estimated variables in Equation (10) and s for those related to suppliers in Equation (11). In addition, $\tilde{q}(w_{it}, \eta_r, \lambda_r)$ in Equation (10) is the optimal quantity of a reference-dependence model conditioning on w_{it} , the wholesale price observed by retailer i in round t . Similarly, $\tilde{w}(\eta_s, \lambda_s)$ in Equation (11) is the optimal wholesale price of a reference-dependence model.

6.1. Quantity

For the retailer, his initial capital r is a reasonable candidate for the reference point. During the experiments, the retailer was allowed to reject a supplier's offer and keep his initial capital. Therefore, it is likely that the retailer treats r as a target income. Formally, $\mathcal{P}^r = r$ is a FRP which is independent of quantity and demand. As the retailer is essentially making a newsvendor decision under a financial constraint, we also consider expected profit $\mathcal{P}^r = \mathbb{E}[\pi_r(q)]$ as a PRP suggested by Uppari and Hasija (2019). Although in our decision support tool, only realized profit with respect to demand was given instead of expected profit, it was still possible to infer the expected profit.

Table 5 shows the treatment-specific and total LL's from the MLE for each reference point and the normative model. Overall, the reference-dependence model, regardless of the reference point,

⁵ We believe such an out-of-sample validation and comparison is more robust and rigorous than an in-sample approach.

outperforms the normative model in terms of goodness of fit in each treatment. The total LL's of the two reference-dependence models are -2359.88 (initial capital) and -2360.73 (expected profit), compared with -2436.80 for the normative model. Performance of each reference point slightly varies from treatment to treatment. Beginning with NR, initial capital fits better than the other two models. Turning to LR, expected profit performs the best, but initial capital is just slightly worse. Both reference points are much better than the normative model. Finally, both reference points achieve similar LL's in HR, both of which are more favorable than that of the normative model. Given that the performance of the two reference points is relatively close in terms of both total likelihood and treatment-specific likelihoods, we consider both in the following analysis.

Table 5 Out-of-sample Model Comparison of Log-likelihood for Retailer Decisions

	Normative Model	FRP: Initial Capital r	PRP: Expected Profit $\mathbb{E}[\pi_r]$
NR	-876.39	-860.27	-866.44
LR	-762.38	-721.55	-717.66
HR	-798.03	-778.07	-776.60
Total	-2436.80	-2359.88	-2360.70

Note: Log-likelihoods (LL's) are calculated on the testing set (33% of the data) with parameters estimated on the training set (67% of the data).

Table 6 shows the estimated parameters and predicted decisions for the two reference points. As we have shown, they both outperform the normative model using an out-of-sample validation; here we use the full dataset to estimate parameters to obtain better insights from the population. Beginning with predicted quantity, both reference points generate predictions that are more accurate than compared to the normative model. Between them, initial capital leads to better predictions despite its slightly worse fit in some treatments. Specifically, the quantity predicted by initial capital as a reference point versus the average observed value is remarkably close across all treatments (observed values are in the table note): 38.76 versus 36.38 in NR, 35.27 versus 35.14 in LR, and 40.44 versus 40.55 in HR. This indicates that a simple FRP is sufficient to capture the observed quantities.

Turning to the estimated parameters in Table 6, note that the magnitude of estimated parameters partially depends on the reference point and varies across treatments. For example, different reference points can affect the size of the gains and losses domains in utility function (8) and therefore the weights η_r and λ_r . As a result, it is less applicable to compare parameters between treatments. However, we can compare $\hat{\eta}_r$ and $\hat{\lambda}_r$ within each treatment. Beginning with NR, although the estimated $\hat{\eta}_r$ and $\hat{\lambda}_r$ by expected profit are close to zero, initial capital (FRP) estimates a higher

Table 6 Estimation Results for Quantity Decisions of Reference-dependent Retailers

	NR		LR		HR	
	FRP: r	PRP: $\mathbb{E}[\pi_r]$	FRP: r	PRP: $\mathbb{E}[\pi_r]$	FRP: r	PRP: $\mathbb{E}[\pi_r]$
Predicted q	38.76	31.22	35.27	34.68	40.44	35.11
Gain: $\hat{\eta}_r$	3.362	0.009	2.815	1.294	0.789	1.050
Loss: $\hat{\lambda}_r$	1.651	0.019	5.412	1.722	7.137	2.105

Note: \hat{q} is the predicted quantity. The average observed quantity is 36.38 in NR, 35.14 in LR, and 40.55 in HR. Two reference points are fitted: i) initial capital r as a FRP, and ii) expected profit $\mathbb{E}[\pi_r]$ as a PRP.

$\hat{\eta}_r$ than $\hat{\lambda}_r$ for NR, indicating that retailers in NR care more about gains than losses. As a comparison, both reference points predict the opposite in LR and HR: $\hat{\lambda}_r$ being higher than $\hat{\eta}_r$, which suggests that retailers exposed to bankruptcy risk care more about losses than gains. This can also explain the directions of the observed quantity deviations. In NR, retailers assigned greater weight to the gains domain where the demand is high, so they overstocked to earn a higher profit in this domain. On the other hand, retailers in LR and HR weighted the losses domain more and chose to understock. This suggests that the presence of bankruptcy risk reverses retailers' preferences around gain and loss domains.

6.2. Wholesale Price

We have shown that the reference-dependence model can capture the retailer's decision well. For the supplier's decision, we begin by determining whether only embedding a reference-dependent quantity response in the supplier's problem is sufficient to predict the observed wholesale price deviations. Given that the observed quantities significantly deviate from the normative predictions, it is reasonable to assume that suppliers did not presume that retailers made quantity decisions in a rational way, especially in LR and HR. On top of the retailer response, we assume the supplier is still rational and maximizing her expected profit function, (i.e., Equations (5a) and (5b)).

We consider the same two reference points for quantity response discussed in Section 6.1: initial capital r as a FRP and expected profit $\mathbb{E}[\pi_r]$ as a PRP. We fit wholesale price decisions using the same approach as quantity decisions: we split the dataset into a training set (67%) and a testing set (33%). Parameters are estimated using the training set, and models are compared using the testing set. For each model, we estimate η_{s-r} and λ_{s-r} in the retailer's quantity response. Note that subscript $s-r$ is used because these are parameters "believed" by the supplier and may be different from the true parameters of the retailer.

Table 7 compares out-of-sample LL's of different quantity responses for each treatment. Overall, initial capital (FRP) outperforms expected profit (PRP), as well as the normative model, in all treatments. Regarding treatment-specific LL's, the normative model fits similarly well as initial capital in NR because the supplier did not deviate from the optimal wholesale price to a large

Table 7 Out-of-sample Model Comparison of Log-likelihood for Supplier Decisions

Quantity Response	Normative model	FRP: initial capital r	PRP Expected Profit $\mathbb{E}[\pi_r]$
NR	-476.29	-476.26	-479.42
LR	-464.54	-401.12	-401.72
HR	-531.26	-449.60	-527.32
Total	-1472.10	-1326.99	-1408.46

Note: Log-likelihoods (LL's) are calculated on the testing set (33% of the data) with parameters estimated on the training set (67% of the data). In the models fitted, suppliers are rational but assume that retailers determine quantity following the model in each column.

extent (20.70 observed vs. 20.00 predicted). As for the other two treatments, a reference-dependent quantity response is necessary to capture the supplier's decision. In LR, the performance of both reference points is similar and much better than the normative model, with initial capital being marginally better. In HR, initial capital achieves a more favorable LL than the other two models. Overall, initial capital (FRP) once again fits best in terms of cumulative LL, across treatments.

Table 8 Estimation Results for Wholesale Price Decisions of Suppliers

	NR		LR		HR	
	FRP: r	PRP: $\mathbb{E}[\pi_r]$	FRP: r	PRP: $\mathbb{E}[\pi_r]$	FRP: r	PRP: $\mathbb{E}[\pi_r]$
Predicted w	20.64	21.01	19.01	19.40	19.79	19.75
$\hat{\eta}_{s-r}$	0.349	0.440	3.854	0.004	0.408	2.985
$\hat{\lambda}_{s-r}$	0.033	0.088	7.859	0.780	8.488	7.386

Note: \tilde{w} is the predicted wholesale price. The average observed wholesale price is 20.70 in NR, 19.05 in LR, and 19.91 in HR. In the models fitted, suppliers are rational and retailers determine quantity following the model in each column.

Table 8 displays predicted wholesale prices and fitted parameters for both reference-dependence models. To fully utilize the data, we conduct the estimations in this table using the full dataset. Compared with expected profit, initial capital as a reference point again generates the best, and fairly accurate, predictions. Specifically, the wholesale price predicted by the initial capital versus average observed value (in table note) is 20.64 versus 20.70 in NR, 19.01 versus 19.05 in LR, and 19.79 versus 19.91 in HR. Although we may not directly compare the estimated parameters between treatments or models, one can still observe that the estimated $\hat{\eta}_{s-r}$ and $\hat{\lambda}_{s-r}$ share a pattern similar to those in Table 6. In particular, $\hat{\eta}_{s-r}$ is larger than $\hat{\lambda}_{s-r}$ in NR, but the opposite is the case in LR and HR. Thus, although the supplier may not accurately anticipate the degree of reference dependence of the retailer, she can largely infer the retailer's preference for gains and losses (which switches in the presence of bankruptcy risk, LR and HR versus NR).

Overall, the estimation results of both quantity and wholesale price decisions show that assuming a reference-dependent retailer can organize the data well. Although there are various candidates for reference points, our results suggest that the retailer's initial capital r is both simple and sufficient to predict the outcomes.

6.3. Discussion

While we show that reference dependence can explain the observed deviations, we also acknowledge that there are other possible behavioral factors. Here we provide a brief discussion of these factors, mainly focusing on the retailer side, as we only assume the retailer is reference dependent in our behavioral model.

First, one may argue that the understocking behavior in LR and HR is due to probability weighting errors (Gonzalez and Wu 1999), where the retailer over-estimates the bankruptcy probability and thus reacts by ordering less. However, recall that the retailer's optimal quantity increases in the bankruptcy risk because of limited liability. As a result, the retailer should prefer a higher quantity if he believes the bankruptcy probability is higher, which is inconsistent with our observed deviation. Therefore, probability weighting cannot be the driver of the understocking we observe in HR and LR.

Another possible bias for the retailer is disappointment aversion (Camerer and Ho 1994), where the retailer will feel disappointed upon bankruptcy. In other words, when calculating expected utility, the retailer has different weights on outcomes depending on whether he is bankrupt or not. Disappointment aversion can account for the understocking in LR and HR, but cannot explain the overstocking in NR. Because we aim to find a simple behavioral model to explain all of the observed deviations, we opt for reference dependence which can predict deviations in both directions.

Turning to the supplier, it is possible that the under-pricing in LR and HR is driven by fairness concerns, as the retailer is predicted to earn much less than the supplier. While fairness can indeed predict a lower wholesale price, we have shown that embedding reference-dependent quantity response in the supplier's problem is sufficient to account for the deviation. Therefore, considering fairness adds unnecessary complexity to the behavioral model without improving the predictive ability significantly. More importantly, recall that in NR, we observed overstocking by retailers and wholesale prices that were slightly higher than the normative prediction, which would be directionally inconsistent with fairness.

7. Robustness Check: Automated-Retailer Treatment

In Section 6.2, we showed that one does not need to model the behavioral bias of the supplier in order to capture their decisions. Assuming that the retailer is reference-dependent can well predict the observed wholesale prices in all treatments. In other words, we implicitly assume that suppliers

are rational and are merely anticipating the quantity decisions of reference-dependent retailers. In this section, we empirically test this assumption by conducting three additional treatments where we automate the retailer decision. If supplier decisions become significantly closer to the original normative predictions, particularly in LR and HR (where wholesale prices deviated significantly from the normative theory), then this behavioral assumption is reasonable. We name the three new treatments NR-A, LR-A, and HR-A, where A represents an automated retailer.

7.1. Experimental Design and Implementation

Different from the main experiments, participants only play as a supplier with an automated retailer, who sets an order quantity following the normative theory. Therefore, there is no group interaction and it takes less time to finish a round. We use a within-subject design here so that we can observe a participant's decisions in all three treatments.⁶ Specifically, participants went through three treatments in a random order. At the beginning of each treatment, they were notified of the change in the retailer's initial capital. To minimize the ordering effect, we randomly assigned an equal number of participants to each possible sequence of three treatments (six combinations of sequence in total). The treatment sequence was unknown to them before the game started.

Participants were told that the retailer they were playing with was automated and would determine its quantity in an expected-profit maximizing way. We provided the same decision support to participants as in the main experiment, a realized profit figure with respect to all possible demand levels. Because there is experimental evidence that fairness concerns can exist even when human participants interact with automated players (Kalkançı et al. 2014), realized retailer profit from the decision support and results were not displayed. Other design elements remained the same as in the main experiment, including the participant pool and the protocol.

Regarding sample size, we assign five participants to each sequence of treatments, with 30 new participants in total. Each treatment has 15 rounds (45 rounds in total). At the beginning of each treatment, they are notified of the retailer's initial capital for the next 15 rounds. We reduced the number of rounds for each treatment because of the limited session length, and we did not observe significant learning in the data. Experiments were again implemented synchronously through Zoom, and the average earnings were \$24.20.

7.2. Results

Average wholesale prices of the automated-retailer treatments, along with those of the original treatments and normative predictions, are shown in Table 9. To better compare decisions from the supplier's perspective, in the second row we show average optimal quantities conditioning on observed wholesale prices.

⁶ This also allows us to avoid issues associated with running out of participants, given a limited participant pool.

Table 9 Average Decisions of Automated Retailer Treatments

	Automated Retailer			Original Treatment			Normative Prediction		
	NR-A	LR-A	HR-A	NR	LR	HR	NR	LR	HR
Wholesale Price	20.47 (0.32)	21.21 (0.30)	22.46 (0.71)	20.70 (0.34)	19.05 [‡] (0.29)	19.91 [‡] (0.47)	20.00	21.47	25.87 [‡]
Optimal Quantity	31.88 (1.09)	38.14 (1.22)	49.16 (1.61)	31.15 (1.18)	46.23 [‡] (0.85)	55.37 [‡] (1.13)	33.33	38.73	42.49 [‡]

Note: Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Average optimal quantity of original treatments is collapsed within suppliers, which are slightly different from those in Table 2. Significance of regressions with random effects compared with decisions of automated-retailer treatments, and the normative predictions, is given by [†] $p < 0.05$ and [‡] $p < 0.01$.

Beginning with NR, wholesale prices did not significantly change moving from a human retailer to an automated retailer or compared with the normative prediction (20.47 in NR-A vs. 20.70 in NR and 20 predicted). This is not surprising, as suppliers had made relatively rational decisions when facing human retailers. Moving to the two treatments where retailers are exposed to bankruptcy risk (LR-A and HR-A), suppliers set wholesale prices much closer to optimal when retailers are rational. Specifically, in LR-A there is no significant difference between the observed average wholesale price and the normative prediction (21.21 observed vs. 21.47 predicted). However, the average price in LR-A is significantly higher than in LR (21.21 vs. 19.05, $p < 0.01$). Finally, although the average wholesale price in HR-A is lower than the normative prediction (22.46 vs. 25.87 predicted, $p < 0.01$), it is significantly higher (closer to optimal) than in HR (22.46 vs. 19.91, $p < 0.01$). One possible explanation is that the predicted optimal price of 25.87 is relatively extreme in HR, which makes suppliers less willing to choose. Overall, compared with the original treatments, wholesale prices were indeed set higher in LR-A and HR-A, but not in NR-A, which is consistent with the estimation results in Section 6.2. Summarizing the findings, we have the following result.

RESULT 5. *When facing automated retailers who are exposed to bankruptcy risk (LR and HR), suppliers set wholesale prices significantly higher than when interacting with human retailers. No significant difference is observed between human and automated retailers in NR.*

8. Conclusion

In practice, a firm's financial position is usually intertwined with its operational decisions (Goentzel and Rice 2015), which suggests the importance of studying them jointly. In this paper, we experimentally investigate how bankruptcy risk affects decision-making and performance under a trade credit contract. Specifically, we consider a two-tier supply chain consisting of a supplier and a capital-constrained retailer. The retailer can purchase from the supplier and pay with his capital up front, or place a higher quantity through trade credit and repay the supplier after demand is realized. If the actual demand is too low for the retailer to repay the supplier in full, the retailer

will transfer all of his cash, including initial capital and any revenue from demand, to the supplier and file for bankruptcy. We consider different levels of exposure to bankruptcy risk by varying the retailer's initial capital and conduct controlled lab experiments with human participants to observe decisions and performance. In the main experiment, we consider three levels of bankruptcy risk: no risk (NR), low risk (LR), and high risk (HR).

The normative theory suggests that a retailer exposed to higher bankruptcy risk is more protected by limited liability and should set a higher quantity compared with the case without capital constraints. Knowing this, the supplier will set a higher wholesale price to extract more profit from the retailer. However, our experimental results show that retailers who are exposed to bankruptcy risk (HR and LR) significantly understock to minimize their actual bankruptcy risk, and that suppliers set lower wholesale prices to induce higher quantities. In contrast, retailers who should, in theory, face no bankruptcy risk (NR) choose to overstock, bringing themselves higher bankruptcy risk, and suppliers set slightly higher wholesale prices. As a consequence of these deviations, a key insight is that observed retailer net profit (excluding initial capital) actually *increases* in bankruptcy risk, whereas the normative theory predicts the opposite. Similarly, another insight is that supplier profits are actually higher when interacting with a retailer that does not face bankruptcy risk, which also runs counter to the normative theory.

To account for these anomalies, we develop a simple reference-dependence model and show that the retailer's initial capital as a simple fixed reference point is sufficient to fit both retailers' and suppliers' decisions in the presence of bankruptcy risk. In particular, the reference-dependence model can predict the observed understocking from retailers in LR and HR. Embedding a reference-dependent retailer in the supplier's problem can also fit suppliers' low wholesale prices in LR and HR. To empirically validate how suppliers anticipate retailers' quantity response, we conduct automated-retailer treatments for each of the original treatments. These new experimental results suggest that wholesale prices remain similar in NR but are significantly higher in LR and HR, supporting our behavioral model.

We believe our paper not only provides important managerial implications but also advances the research literature in key ways. Beginning with the former, suppliers should be aware that retailers who face bankruptcy risk tend to understock, in order to minimize the likelihood of going bankrupt. Although offering a more generous wholesale price can induce a higher quantity, it ultimately results in a supplier profit that is still lower than the normative predictions. Due to this, suppliers actually earn a higher profit when partnering with a retailer who is not capital constrained (which runs counter to theory), since retailers do not understock in such an environment. Regarding retailers, although it is reasonable that retailers understock to avoid bankruptcy, our data suggest that they may also fail to realize the potential benefits of limited liability. Turning to our contribution to the

literature, to our knowledge, this is the first behavioral investigation on bankruptcy risk in supply chains. Our experimental results show that behavioral factors play a significant role in decision-making when bankruptcy risk is present. Further, note that the observed deviations are surprisingly large despite our simplified setting. Indeed, we do not consider other components in supply chain finance such as interest rates, bank financing, factoring, and insurance, which could exacerbate these effects. Now that we have established initial results in supply chains with bankruptcy risk, we believe more research considering such complexities is critical to providing a comprehensive understanding of behavioral factors in this field.

Finally, there are limitations in our study that we consider as opportunities for future research. First, we consider a price-only trade credit contract without an interest rate. In practice, suppliers may also provide a wholesale price and an interest rate, allowing retailers to pay early with a discount or pay later with interest (Yang and Birge 2018). Such a contract, also known as an early-payment-discount contract, requires suppliers to optimize an interest rate and a wholesale price simultaneously. Experimental research dedicated to studying both price and interest rates would be interesting, especially given that Kouvelis and Zhao (2012) show that a rational supplier should always set her interest rate as a risk-free rate. Second, while supplier financing has become increasingly important, bank financing as a traditional financing tool still prevails in practice. Comparing both financing tools in terms of supply chain performance can generate useful managerial implications. Finally, sometimes suppliers can also be capital-constrained and in favor of offering trade credit due to competition Lee and Rhee (2011). The supplier alone, or both supplier and retailer, being capital-constrained will be an interesting direction for future research.

References

- Aral KD, Giambona E, Wang Y (2021) Buyer's bankruptcy risk, sourcing strategy, and firm value: Evidence from the supplier protection act. *Management Science* Articles in Advance, <https://doi.org/10.1287/mnsc.2021.4270>.
- Astvansh V, Jindal N (2022) Differential effects of received trade credit and provided trade credit on firm value. *Production and Operations Management* 31(2):781–798.
- Baron O, Hu M, Najafi-Asadolahi S, Qian Q (2015) Newsvendor selling to loss-averse consumers with stochastic reference points. *Manufacturing & Service Operations Management* 17(4):456–469.
- Camerer CF, Ho TH (1994) Violations of the betweenness axiom and nonlinearity in probability. *Journal of risk and uncertainty* 8(2):167–196.
- Chen DL, Schonger M, Wickens C (2016) otree – an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9:88–97.
- Chod J (2016) Inventory, risk shifting, and trade credit. *Management Science* 63(10):3207–3225.

- Chod J, Lyandres E, Yang SA (2019) Trade credit and supplier competition. *Journal of Financial Economics* 131(2):484–505.
- Cui TH, Raju JS, Zhang ZJ (2007) Fairness and channel coordination. *Management Science* 53(8):1303–1314.
- Cunat V (2007) Trade credit: suppliers as debt collectors and insurance providers. *The Review of Financial Studies* 20(2):491–527.
- Dada M, Hu Q (2008) Financing newsvendor inventory. *Operations Research Letters* 36(5):569–573.
- Davis A, Katok E, Santamaría N (2014) Push, pull, or both? a behavioral study of how the allocation of inventory risk affects channel efficiency. *Management Science* 60(11):2666–2683.
- Donohue K, Katok E, Leider S, eds. (2019) *The Handbook of Behavioral Operations* (Wiley).
- Giannetti M, Burkart M, Ellingsen T (2011) What you sell is what you lend? explaining trade credit contracts. *The Review of Financial Studies* 24(4):1261–1298.
- Goentzel J, Rice JB (2015) Guest voices: Managing supply chains is intertwined with financial management. <https://www.wsj.com/articles/guest-voices-managing-supply-chains-is-intertwined-with-financial-management-1437662343>, accessed August 10, 2019.
- Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cognitive psychology* 38(1):129–166.
- Gupta D, Wang L (2009) A stochastic inventory model with trade credit. *Manufacturing & Service Operations Management* 11(1):4–18.
- Gurnani H, Ramachandran K, Ray S, Xia Y (2014) Ordering behavior under supply risk: An experimental investigation. *Manufacturing & Service Operations Management* 16(1):61–75.
- Haley CW, Higgins RC (1973) Inventory policy and trade credit financing. *Management science* 20(4-part-i):464–471.
- Ho TH, Lim N, Cui TH (2010) Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science* 56(11):1891–1910.
- Hyndman K, Embrey M (2019) Econometrics for experiments. Donohue K, Katok E, Leider S, eds., *The Handbook of Behavioral Operations*, chapter 2, 35–88 (Wiley).
- Kahneman D, Knetsch JL, Thaler RH (1991) Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic perspectives* 5(1):193–206.
- Kalkanci B, Chen KY, Erhun F (2014) Complexity as a contract design factor: A human-to-human experimental study. *Production and Operations Management* 23(2):269–284.
- Kalkanci B, Chen KY, Erhun F (2014) Complexity as a contract design factor: A human-to-human experimental study. *Production and Operations Management* 23(2):269–284.
- Kouvelis P, Zhao W (2012) Financing the newsvendor: supplier vs. bank, and the structure of optimal trade credit contracts. *Operations Research* 60(3):566–580.

- Kouvelis P, Zhao W (2018) Who should finance the supply chain? impact of credit ratings on supply chain decisions. *Manufacturing & Service Operations Management* 20(1):19–35.
- Lai G, Debo LG, Sycara K (2009) Sharing inventory risk in supply chain: The implication of financial constraint. *Omega* 37(4):811–825.
- Lee CH, Rhee BD (2011) Trade credit for supply chain coordination. *European Journal of Operational Research* 214(1):136–146.
- Lee HH, Zhou J, Wang J (2018) Trade credit financing under competition and its impact on firm performance in supply chains. *Manufacturing & Service Operations Management* 20(1):36–52.
- Luo W, Shang KH (2019) Managing inventory for firms with trade credit and deficit penalty. *Operations Research* 67(2):468–478.
- Murfin J, Njoroge K (2015) The implicit costs of trade credit borrowing by large firms. *The Review of Financial Studies* 28(1):112–145.
- Murphy R (2018) How supply chain finance is offering companies a new cash source. <https://www.forbes.com/sites/forbesbusinessdevelopmentcouncil/2018/04/09/how-supply-chain-finance-is-offering-companies-a-new-cash-source/#daac23422f16>, accessed August 6, 2019.
- Ning J (2022) Strategic trade credit in a supply chain with buyer competition. *Manufacturing & Service Operations Management* Articles in Advance, <https://doi.org/10.1287/msom.2022.1103>.
- Peura H, Yang SA, Lai G (2017) Trade credit in competition: A horizontal benefit. *Manufacturing & Service Operations Management* 19(2):263–289.
- Seifert D, Seifert RW, Protopappa-Sieke M (2013) A review of trade credit literature: Opportunities for research in operations. *European Journal of Operational Research* 231(2):245–256.
- Tang CS, Yang SA, Wu J (2018) Sourcing from suppliers with financial constraints and performance risk. *Manufacturing & Service Operations Management* 20(1):70–84.
- Tereyağoğlu N, Fader PS, Veeraraghavan S (2018) Multiattribute loss aversion and reference dependence: Evidence from the performing arts industry. *Management Science* 64(1):421–436.
- Tunca TI, Zhu W (2018) Buyer intermediation in supplier finance. *Management Science* 64(12):5631–5650.
- Uppari BS, Hasija S (2019) Modeling newsvendor behavior: A prospect theory approach. *Manufacturing & Service Operations Management* 21(3):481–500.
- WTO (2019) Trade finance and the compliance challenge. https://www.wto.org/english/res_e/booksp_e/tradefinance19_e.pdf, accessed December 25, 2020.
- Yang SA, Bakshi N, Chen CJ (2021) Trade credit insurance: Operational value and contract choice. *Management Science* 67(2):875–891.
- Yang SA, Birge JR (2018) Trade credit, risk sharing, and inventory financing portfolios. *Management Science* 64(8):3667–3689.

Appendix A: Theoretical Details of the Reference-dependence Model

A.1. Fixed Reference Point: Initial Capital

When the retailer's reference point is his initial capital r , the retailer is in the gains domain when demand is larger than $\frac{wq}{p}$, and is in the losses domain otherwise. Depending on whether the retailer need financing or not, his utility function is shown as follows.

- If $wq - r > 0$:

$$\mathbb{E}[u_r^{FRP}(q)] = \mathbb{E}[\pi_r(q)] + \eta \int_{\frac{wq}{p}}^q [p \min(x, q) - wq] dF(x) - \lambda \int_0^{\frac{wq}{p}} [r - (px - wq + r)^+] dF(x)$$

- Otherwise:

$$\mathbb{E}[u_r^{FRP}(q)] = \mathbb{E}[\pi_r(q)] + \eta \int_{\frac{wq}{p}}^q (p \min(x, q) - wq) dF(x) - \lambda \int_0^{\frac{wq}{p}} (-px + wq) dF(x)$$

A.2. Prospect Reference Point: Expected Profit

When the retailer's reference point is his expected profit given a quantity q , the reference point $\mathcal{X}(q)$ is dependent on q .

- If $wq - r > 0$:

$$\begin{aligned} \mathbb{E}[u_r^{PRP}(q)] = & \mathbb{E}[\pi_r(q)] + \eta \int_{q - \frac{q^2}{2a} + \frac{(r-wq)^2}{2ap^2}}^q [p \min(x, q) - wq + r - \mathbb{E}[\pi_r(q)]] dF(x) \\ & - \lambda \int_0^{q - \frac{q^2}{2a} + \frac{(r-wq)^2}{2ap^2}} [\mathbb{E}[\pi_r(q)] - (px - wq + r)^+] dF(x) \end{aligned}$$

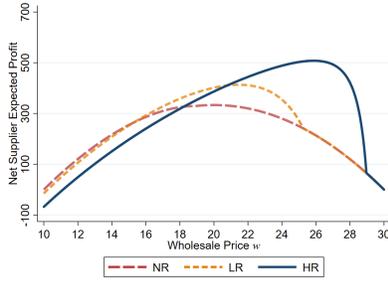
- Otherwise:

$$\begin{aligned} \mathbb{E}[u_r^{PRP}(q)] = & \mathbb{E}[\pi_r(q)] + \eta \int_{q - \frac{q^2}{2a}}^q [p \min(x, q) - wq + r - \mathbb{E}[\pi_r(q)]] dF(x) \\ & - \lambda \int_0^{q - \frac{q^2}{2a}} [\mathbb{E}[\pi_r(q)] - px + wq - r] dF(x) \end{aligned}$$

Appendix B: Additional Figures

Figure B.1 Net Expected Profit with Respect to Wholesale Price

(a) Supplier



(b) Retailer

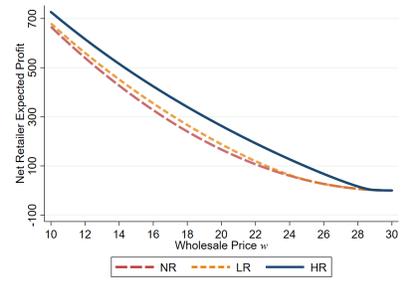
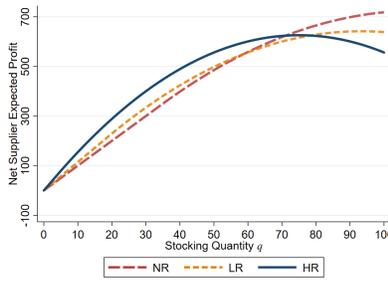


Figure B.2 Net Expected Profit with Respect to Order Quantity

(a) Supplier



(b) Retailer

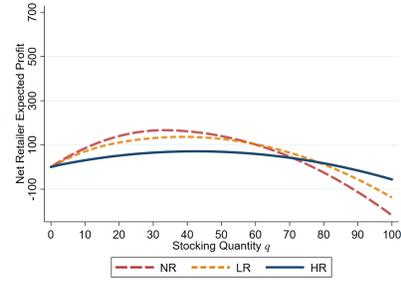
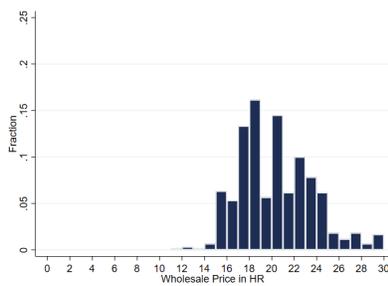
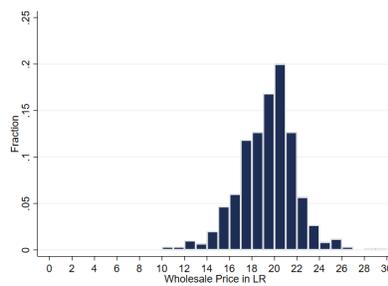


Figure B.3 Distribution of Observed Wholesale Prices

(a) HR



(b) LR



(c) NR

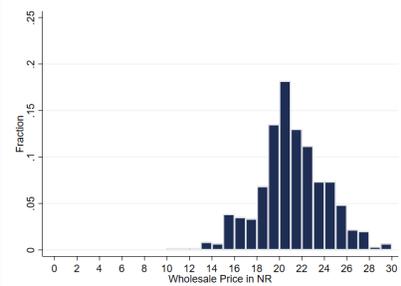
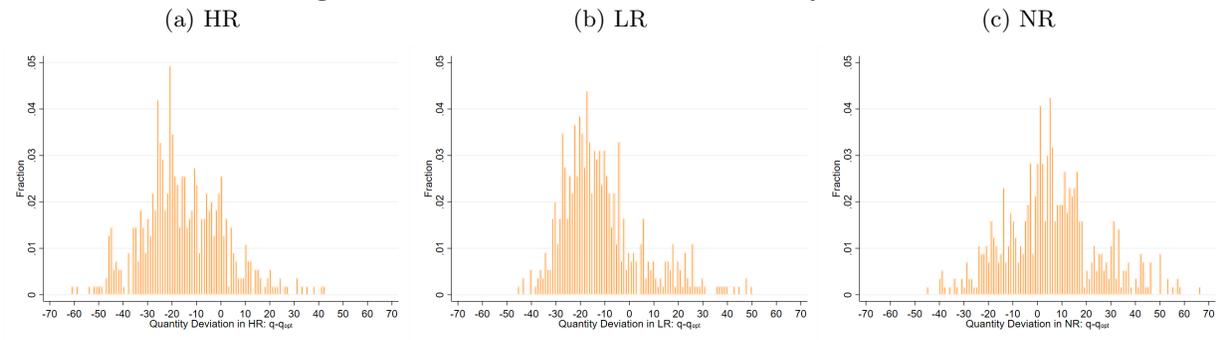


Figure B.4 Distribution of Observed Quantity Deviations



Appendix C: Experimental Details

C.1. Instructions for HR

You are about to participate in a decision making experiment. If you follow these instructions carefully and make good decisions, you will earn money that will be paid to you by PayPal at the end of the session. Your earnings will depend on your decisions, the decisions of other participants, and chance. Please keep your camera on and do not unmute yourself throughout the game. If you have any questions, or you accidentally close your browser, please ask the experimenter by chat.

Game Overview

You will be randomly assigned to one of two roles for the duration of the session: a retailer or a supplier. The retailer independently purchases units of a product from the supplier at a wholesale price per unit, and sells units to customers for \$30 per unit (all \$ are laboratory dollars). Customer demand is randomly and independently determined in each round, from 0 to 100, with each integer in that range equally likely. The retailer cannot make extra money from any remaining units.

The supplier starts with \$400 cash and produces units at a cost of \$10 per unit. The retailer starts with \$100 cash to purchase units from the supplier. If the retailer has sufficient cash, it pays the supplier upfront. If the cash is insufficient, the retailer pays the supplier after earning revenue by satisfying the (random) customer demand. If the retailer cannot pay the supplier in full after satisfying its demand, the retailer transfers all its remaining cash to the supplier and earns a profit of \$0.

Timeline in Each Round

You will play in 20 rounds. Each round has 2 stages. Decisions at each stage are specified as follows.

a) If you are the supplier:

- Stage 1: Set a wholesale price.
- Stage 2: Wait for the retailer to decide its stocking quantity.

b) If you are the retailer:

- Stage 1: Wait for the supplier to decide a wholesale price.
- Stage 2: Set a stocking quantity. If you are not satisfied with the wholesale price set by the supplier, you can reject the wholesale price offer by clicking on the “Reject The Wholesale Price” button without setting a stocking quantity. In this case, both parties keep their initial cash.

Profit Calculations

If the retailer can pay the supplier in full:

- Supplier Profit = (Wholesale Price – \$10) × Quantity + \$400
- Retailer Profit = \$30 × Units Sold + \$100 – Wholesale Price × Quantity

If the retailer cannot pay the supplier in full:

- Supplier Profit = \$30 × Units Sold + \$100 + \$400 – \$10 × Quantity
- Retailer Profit = \$0

“Units Sold” is the lower number of demand and quantity. That is, you cannot sell more than the quantity you ordered.

Decision Support

In each stage, there will be a testing section and a decision-making section, such that you can test your decision before submission. Screenshots of the two stages are shown below. Slide the scroll bar and you will see actual profits of both parties when demand is 0, 2, 4, ..., 98, 100, respectively. You will also see the likelihood that the retailer will not be able to pay the supplier in full. Note that for the supplier, the profits and the likelihood of the retailer not paying in full in this testing function are based on the assumption that the retailer chooses the stocking quantity to maximize its average profit. The true numbers will be different if the retailer makes a different choice.

Results

After both stages, demand will be revealed. Then you will see all of the information of that round in the results page, including your profit and your partner's profit.

This concludes one round. In total there will be 20 rounds. At the beginning of each round, you will be randomly re-matched with another participant. Note that customer demand in one round is completely independent from customer demand in any other round.

Before the game starts, you will be assigned to a breakout room. Note also that you will not be matched with any other participant in the same room.

Example

These numbers are simply used to illustrate the sequence of decisions and should not be construed as "good" or "bad" decisions. The supplier has initial cash \$400. The retailer has initial cash \$100.

Decisions & Realized Demand:

- The supplier sets a wholesale price of \$18 (note that decimals are permitted for the wholesale price).
- The retailer sets a stocking quantity of 60. Full payment to the supplier is $\$18 \times 60 = \$1,080$.

Outcome 1: Realized demand is 68. However, since the retailer sets a stocking quantity of 60, the units sold is 60. The retailer's total cash ($\$30 \times 60 + \$100 = \$1,900$) is sufficient.

- Supplier Profit = $(\$18 - \$10) \times 60 + \$400 = \880
- Retailer Profit = $\$30 \times 60 + \$100 - \$18 \times 60 = \820

Outcome 2: Realized demand is 12, which is lower than the stocking quantity set by the retailer. The units sold is 12. The retailer's total cash ($\$30 \times 12 + \$100 = \$460$) is insufficient.

- Supplier Profit = $\$30 \times 12 + \$100 + \$400 - \$10 \times 60 = \$260$
- Retailer Profit = \$0

Payment

At the end of the session, one round will be randomly picked to calculate your actual earnings from the game. Each round is equally likely to be picked. Therefore, it is in your best interest to maximize your profit in each round. The one-round profit will be converted to US dollars at the rate of 50.0 laboratory dollars for 1 US dollar. These profits will be added to your \$7.00 show-up fee, displayed on your screen, and paid to you by PayPal after the session. Please fill out the survey at the end of the session to provide your payment information.

Acknowledgments