

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, New York 14853;

^bShenzhen Finance Institute, School of Management and Economics, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), Shenzhen 518172, China; ^cNaveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080

*Corresponding author

Contact: adavis@cornell.edu,  <https://orcid.org/0000-0002-1689-2299> (AMD); huangrihuan@cuhk.edu.cn,

 <https://orcid.org/0000-0002-6790-6814> (RH); kyleb.hyndman@utdallas.edu,  <https://orcid.org/0000-0003-3666-8734> (KH)

Received: January 10, 2023

Revised: May 16, 2024; January 20, 2025

Accepted: April 18, 2025

Published Online in Articles in Advance:
May 16, 2025

<https://doi.org/10.1287/msom.2022.0653>

Copyright: © 2025 INFORMS

Abstract. *Problem definition:* 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. *Methodology/Results:* 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 with 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 baseline predictions. To account for these observed decisions, we show that a behavioral model of reference dependence and fairness organizes the data well. *Managerial implications:* Our work demonstrates that the presence of bankruptcy risk for a retailer significantly alters supply chain decisions in systematic ways, which has direct consequences on profits.

Funding: This work was supported by the National Natural Science Foundation of China [72402198].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/msom.2022.0653>.

Keywords: behavioral operations • OM-finance interface • supply chain management

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, up to 80% of trade is financed by credit or credit insurance (WTO 2019). 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 (Tang et al. 2018, Tunca and Zhu 2018).

One popular supply chain financing tool is trade credit, where a buyer (hereafter retailer) 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 retailer can either pay within 10 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 retailers, allowing them to purchase a desired quantity with relatively low financing costs compared with bank financing. Indeed, it is an important financing source for retailers that are unable to receive bank loans altogether (Biais and Gollier 1997, Giannetti et al. 2011). Trade credit also offers many further benefits, such as building relationships between firms (Wilson and Summers 2002), increasing

the likelihood of trade (Breza and Liberman 2017), inducing relationship-specific investments by suppliers (Dass et al. 2015), serving as a competitive tool against other suppliers (Lee et al. 2018), and conveying valuable information to external investors (Aktas et al. 2012). Therefore, it is not surprising that retailers, even those without capital constraints, may purchase from smaller suppliers through trade credit (Murfin and Njoroge 2015). Furthermore, suppliers, as the financing provider, can take advantage of trade credit. By setting appropriate contract terms, suppliers can indirectly affect the retailer's financing decisions or potentially induce higher ordering quantities from retailers.

Although trade credit brings benefits to the supply chain, it also comes with downsides. If retailers default because of 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. Therefore, by providing financing to retailers, such suppliers partially share bankruptcy risk with them (although suppliers can be better informed about the risk of providing financing to retailers, compared with financing providers outside of the supply chain such as a bank, when a retailer's financial status is private information (Biais and Gollier 1997)). However, suppliers may also be able to take advantage of retailers' exposure to bankruptcy risk and extract more profits instead.

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 baseline theory (that assumes risk-neutral expected profit-maximizing decision-makers) predicts that a retailer facing higher bankruptcy risk should set a higher quantity (than if without bankruptcy risk), making the supplier charge a higher wholesale price to extract more profits (Kouvelis and Zhao 2012). Although such baseline theory assumes profit-maximizing decision-makers, it is unclear whether behavioral biases can drive decisions away from these predictions. For example, decisions may be influenced by reference points (Ho et al. 2010) or fairness concerns (Fehr and Schmidt 2000). Moreover, if behavioral biases do play a role in decision-making, then the magnitude of the biases and how they 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 outcomes. We consider a two-tier supply chain consisting of a supplier and a capital-constrained retailer. 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 its 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, then 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 laboratory experiments with human participants to test the baseline 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? To the extent that decisions and outcomes deviate from the baseline predictions, we are also interested in identifying what behavioral bias(es) can account for such outcomes.

Our experimental results indicate that capital-constrained retailers significantly understock relative to the baseline benchmark (which reduces their *actual* bankruptcy risk). Suppliers set significantly lower wholesale prices in an attempt to induce higher quantities and reduce the profit gap between the two parties. Conversely, when retailers do not face any meaningful capital constraints (i.e., no bankruptcy risk), they set quantities that are slightly higher than the baseline predictions, and suppliers set wholesale prices close to the optimal predictions. In short, the presence of bankruptcy risk leads both retailers and suppliers to deviate from the baseline predictions in systematic ways relative to a setting without bankruptcy risk.

As a consequence of these decisions, 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 baseline predictions. For instance, the baseline 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. To test the robustness of these results, we conduct an additional experimental treatment where we alter the *supplier's* initial capital and find that similar deviations persist in both retailer and supplier decisions.

In an effort to account for these deviations under bankruptcy risk, we find that a reference-dependent retailer, whose stochastic reference point is the realized

profit from ordering mean demand, fits the data well. Given that suppliers do not directly face bankruptcy risk and are predicted to earn disproportionately higher profits when retailers are capital constrained, we find that a fairness-minded supplier best fits the observed wholesale price decisions. To further increase our confidence in these behavioral explanations, we show that these behavioral tendencies continue to organize the data well in our robustness treatment with low supplier capital.

Our work contributes to both practitioners and researchers. Regarding the former, we show that financially risky retailers earn a higher profit under a trade credit contract. Although this may further explain the prevalence of trade credit from a behavioral angle, our results also imply that these higher retailer profits come at the expense of suppliers, who should actually prefer trading with retailers who are not overly capital constrained. These insights are especially noteworthy for managers in that they demonstrate when and how supply chain parties deviate from the baseline predictions, which has direct implications on profits. 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 observe that supply decisions differ significantly when retailers face bankruptcy risk, the latter of which is a common feature for many companies. Furthermore, we show that behavioral factors, notably reference dependence and fairness, capture the data well. 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). We hope that our work will serve as an important building block for such studies.

2. Literature Review

Trade credit has been studied from a theoretical perspective for decades in the literature. Existing theoretical studies in finance and economics focus on various aspects of trade credit, such as financing advantages and price discrimination (Petersen and Rajan 1997). Our work is related mostly to theories on the risk-sharing role of trade credit (see, e.g., Yang and Birge 2018). Specifically, Brander and Lewis (1986) analytically introduced the “limited liability effect;” firms with limited liability can take advantage of debt to be more aggressive in a Cournot competition setting, which is consistent with the baseline model in our paper. In contrast to Brander and Lewis (1986), Showalter (1995) showed that in Bertrand competition, a firm’s debt choice depends on whether uncertainty originates from demand or cost.

Related theoretical work in the operations management literature has focused mainly on operational

decisions under trade credit contracts (see Seifert et al. 2013) for a comprehensive review). Dada and Hu (2008) found that a capital-constrained newsvendor will borrow funds from a bank and order less than would be ideal if the borrowing cost is not too high. They also showed that the channel can be coordinated by a nonlinear loan schedule. Lai et al. (2009) studied 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 financial constraints are absent. Kouvelis and Zhao (2012) studied a supply chain of a supplier and a capital-constrained retailer, with endogenous contract terms, and compared bank-financing-only to supplier-financing-only scenarios. They showed that the limited liability effect also exists in such a setting; that is, the retailer sets a higher quantity when facing higher bankruptcy risk, which leads the supplier to set a higher wholesale price. They further studied a two-tier supply chain with capital constraints in a more complex setting, such as how bankruptcy costs (fixed and variable) affect contract design when only bank financing is available (Kouvelis and Zhao 2016) and how exogenous default risk (dependent on credit rating) affects decisions when both financing tools are present (Kouvelis and Zhao 2018). We consider a simplified setting of the supplier-financing-only scenario in Kouvelis and Zhao (2012), and our experimental results contrast with theoretical predictions, such as significant understocking and underpricing when the retailer has bankruptcy risk.

Like analytical work, there has been empirical research on trade credit. Although it is common for powerful suppliers to provide trade credit to small retailers, Murfin and Njoroge (2015) found that even retailers without capital constraints purchase via trade credit from smaller and weaker suppliers. Studying the impact of restricting trade credit, Barrot (2016) showed that it significantly reduces liquidation risk of financially constrained suppliers (lenders), whereas Breza and Liberman (2017) found that it reduces the likelihood of trade by 11%. Lee et al. (2018) studied how, when competition exists, trade credit affects firm performance and found that when trade credit offered by suppliers exceeds industry-average levels, retailers’ performance is negatively associated with the amount of trade credit. Taking advantage of a quasi-natural setting provided by a regulatory shock, Aral et al. (2022) studied how a distressed buyer’s sourcing strategy is affected by bankruptcy risk. They showed that capital-constrained buyers are forced to under-diversify compared with those without capital constraints. Astvansh and Jindal (2022) found 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 an opposite effect. They argued that the main reason is the disparate nature of dependence in the supply chain. Chen et al. (2023) empirically investigated the impact of trade credit in four French retail sectors and found that a decline in trade credit usage can result in an inventory reduction (0.67% reduction in inventory per 1% reduction in trade credit usage) as well as a decline in revenue and gross profit. By conducting laboratory experiments, our work complements these existing empirical studies by uncovering operational decisions at a more detailed level in a controlled environment and how they affect supply chain outcomes in the presence of trade credit.

Turning to behavioral studies in finance and operations management, to our knowledge, none of the existing studies in the operations management literature have investigated trade credit and bankruptcy risk from a behavioral perspective (Donohue et al. 2019), whereas there are two papers that are noteworthy. Oechssler and Schuhmacher (2004) experimentally tested both Brander and Lewis (1986) and Showalter (1995) and found that the observed results are consistent only with Showalter (1995). For Cournot competition, human subjects recognize only their own benefit from limited liability, but not their competitors', which leads to lower debt levels than predicted. Our experimental results suggest similar deviations in a two-tier supply chain setting without competition, and we further show that reference dependence and fairness can account for the deviations. Chen et al. (2013) investigated how payment schemes affect order quantity decisions. Specifically, they considered a scenario where the retailer pays only after realized demand, which is essentially purchasing purely through trade credit. They found that order decisions are closest to optimal in this scenario compared with two other cases, one where the retailer pays when ordering and another where the retailer is paid by the customer in advance (similar to buyer financing). There are two main differences between their setting and ours. First, they focused on inventory decisions, whereas we consider both pricing and inventory decisions in a two-tier supply chain. Second, they did not consider that the retailer may not be able to repay in full and thus file for bankruptcy, whereas we study how such bankruptcy

risk affects decisions and outcomes in a supply chain (and we find significant understocking behavior in the presence of bankruptcy risk).


3. Baseline Model

Consider a retailer ordering a product from a supplier 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, it may lack sufficient cash and require financing. Depending on w , q , and r , there are two scenarios for the retailer's financing outcome. If $wq \leq r$, then the retailer's capital is sufficient to cover its purchasing cost. The retailer pays the supplier wq up front. If $wq > r$, then 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. In practice, the supplier may charge additional interest for the delayed payment; we simplify the setting by assuming that the risk-free rate is zero and the supplier sets its interest rate at the risk-free rate, which is optimal for the supplier (Kouvelis and Zhao 2012). If the realized demand is sufficiently high, then the retailer can repay the supplier in full, $wq - r$. If the demand is low, then the retailer transfers all of its revenue from realized demand to the supplier and earns zero profit, thus going bankrupt. By bankruptcy, we assume that the retailer bears limited liability, following the existing literature (see, e.g., Chod 2017, Yang and Birge 2018). Figure 1 summarizes the sequence of events.

For the baseline model introduced in this section, we make some underlying assumptions related to the financial aspects. As such, the model is parsimonious but captures key tradeoffs. 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, the cost, any initial capital, and demand distribution are common knowledge).

Figure 1. Sequence of Decisions and Events



3.1. Retailer's Decision

Given a wholesale price w , the retailer determines q (and thus the choice of financing) to maximize its 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

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

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

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

Kouvelis and Zhao (2012) showed 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. In general, the financing region is larger when the retailer has less capital r and therefore is exposed to higher bankruptcy risk. The two regions also indicate that the retailer requires financing when the

wholesale price is intermediate (i.e., $w_u < w < w_l$). If the wholesale price is too low, then 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, then the retailer will order a quantity that is so small that no financing is needed.


Figure 2 shows how q^* varies with respect to w and r . In Figure 2(a), 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 it 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). This is consistent with the "limited liability effect" (Brander and Lewis 1986). Figure 2(b) shows how q^* varies with respect to retailer capital r given the corresponding optimal wholesale price. Again, the retailer will order more when it has less capital (i.e., more exposed to bankruptcy risk) because it 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) showed 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

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

Figure 2. (Color online) Optimal Quantity Under Supplier Financing



Notes. $p = 30, c = 10, r = 100$. Demand follows a uniform distribution between 0 and 100.

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

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

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

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 $\bar{q} \leq q_l$, then $w^* = p\bar{F}(\bar{q})$, and the retailer does not need financing.
- (Financing) If $\bar{q} > q_l$, then 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_l, q_u) . 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.

Turning to how w^* varies with respect to r , Figure 3 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 the retailer.

Although theory predicts that a retailer exposed to higher bankruptcy risk should order a higher quantity (relative to lower bankruptcy risk), it is unclear whether human decision-makers would behave in this way. For instance, certain studies have shown that professional participants exhibit myopic loss aversion (to even a greater extent than student participants) in financial settings (Haigh and List 2005, Eriksen and Kvaløy 2010). In our case, a human retailer with low initial capital may not be willing to respond to the presence of limited liability and bear the high bankruptcy risk that comes with it and would rather 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.


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 possible only when retailers overorder by extreme amounts). Because we are interested mainly 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 at $p = 30$ and the supplier production cost at $c = 10$. 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 for this is twofold. a) The baseline theory shows that the supplier's optimal decision is unaffected by r_s , and b) $r_s = 400$ makes supplier bankruptcy rarely happen (the average supplier bankruptcy rate in our data is 0.39%), which means that we can focus on the retailer's bankruptcy risk as a key treatment

Figure 3. (Color online) Optimal Wholesale Price with Respect to Retailer Capital



Notes. $p = 30, c = 10, r_s = 400$. Demand follows a uniform distribution between 0 and 100.

variable. Note that if a supplier goes bankrupt in the game, it 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 baseline predictions of all three treatments are shown in Table 1. Understanding that the retailer is more protected by limited liability when it 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 ($k = 0$), whereas it is $k = 14.38$ in LR 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 it significantly overstocks.

Regarding predicted profits, supplier expected profit increases from 733.33 in NR to 813.16 in LR to 907.87 in HR because the supplier is better able to extract profit from the retailer as it 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.

Each treatment included 60 participants recruited from Cornell University, where cash was the only incentive offered. Each participant was randomly assigned a role and stayed in that role throughout the game, which included 20 rounds in total. 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. To make sure participants had a complete understanding of the game, they were required to answer several multiple-choice comprehension questions before proceeding. Based on the responses, we found little evidence that the alternative framing of bankruptcy affected participants' understanding of the game.

Our experiment was implemented through oTree (Chen et al. 2016) synchronously over Zoom with webcams (because of COVID-19; see Li et al. 2021). Before each session started, a researcher read the instructions aloud and answered any questions. Participants were then required to answer the aforementioned comprehension questions. 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. Please see Online Appendix E for sample instructions.

4.2. Behavioral Hypotheses

Although the baseline theory provides predictions for a profit-maximizing decision-maker, existing experimental evidence suggests that human participants can exhibit behavioral biases and deviate from rationality. In this subsection, we discuss two behavioral

Table 1. Experimental Parameter Settings and Predictions

	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	1,600.00 [500.00]	1,349.60 [549.60]	1,079.00 [579.00]

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

hypotheses related to the order quantity and wholesale price, 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 baseline 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, it may become a reference point for them, making them averse to earning less than r and, therefore, averse to bankruptcy. In other words, retailers will attempt to face an actual bankruptcy risk that is lower than the baseline theory predicts. This can be accomplished by setting an order quantity that is below the baseline 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 with a setting where the retailer is exposed to a relatively lower bankruptcy risk (e.g., LR; recall Figure 2(b)). However, the increased protection of limited liability in HR may not dominate this bias (which should be more salient when bankruptcy risk is high). As a result, retailers may have less incentive to order the high predicted quantity in HR and to a lesser extent in LR. Therefore, we offer the following hypothesis (Hypothesis 1).

Hypothesis 1. *Retailers in LR and HR will set lower quantities than the baseline theory predicts in order to achieve a lower bankruptcy risk. The relative deviation $(q - q^*)/q^*$ will be larger in absolute value in HR than in LR.*

The supplier provides financing to the retailer and, when retailers are capital constrained, may have fairness concerns as the retailer directly faces the bankruptcy risk and earns a much lower expected profit (Fehr and Schmidt 1999), as shown in Table 1. In addition, the supplier may also anticipate the retailer's understocking behavior. If so, the supplier will set a more generous wholesale price to induce a higher quantity. We also expect the relative deviation to be larger in HR than in LR, for two reasons. First, retailers in HR have much lower initial capital than suppliers, potentially causing stronger fairness concerns. Second, 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. Thus, we have Hypothesis 2.

Hypothesis 2. *Suppliers in LR and HR will set lower wholesale prices than the baseline theory predicts. The*

relative deviation $(w - w^)/w^*$ in absolute value will be larger in HR than in LR.*

Unlike decisions, hypothesizing about how profits may deviate from the theory is more challenging. If both Hypotheses 1 and 2 are supported, then we can anticipate that the supplier's profit will be lower than predicted. However, how the retailer's profit will change is unclear; although the retailer can be hurt by its suboptimal ordering, it may also benefit from the lower wholesale price set by the supplier. Therefore, we aim to find the answers around profits by conducting a laboratory experiment and conduct a more detailed investigation in Section 5.

5. Results

We first present observed decisions and any deviations and then show how these decisions translate into profits. Recall that a retailer was allowed to reject a supplier's offer if the retailer 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 (standard errors clustered at the cohort level) for hypothesis tests (Hyndman and Embrey 2019), with details given in Online Appendix C.

Before showing the results, we first check whether any significant learning took place. Overall, there is only mild learning for the wholesale price and none 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, $p = 0.097$) and LR (18.84 vs. 19.30, $p = 0.093$) but not in HR (19.80 vs. 20.00, $p = 0.38$). Although there are some differences, the magnitude is small; 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, all $p > 0.1$). Given that the learning is either slight or insignificant, we keep the full data for all analyses (rejected data are still excluded).

Regarding bankruptcy effects across rounds, we do not find significant evidence that retailers deviated more in quantity, or tended to reject offers more, after experiencing bankruptcy in any treatment. Similarly, we do not find significant evidence that suppliers set wholesale prices differently after their paired retailers went bankrupt in the previous round.

5.1. Decisions

Table 2 presents the average observed decisions (left side) and the theoretical predictions (right side). For

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]

Notes. Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Predicted quantities and bankruptcy thresholds are conditioned on observed wholesale prices. Unconditional baseline predictions, when applicable, are reported in square brackets. Significance of regressions with random effects (standard errors clustered at the cohort level) compared with observed vs. (conditional) baseline predictions is given by [†] $p < 0.01$ and [‡] $p < 0.05$.

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 vs. NR relative to the baseline predictions. Specifically, retailers significantly understock when they are exposed to bankruptcy risk, 40.55 observed vs. 55.76 predicted in HR ($p < 0.01$, namely 27.27% lower) and 35.14 observed vs. 46.46 predicted in LR ($p < 0.01$, 24.37% lower), but overstock when they are not exposed to bankruptcy risk in NR, 36.38 observed vs. 31.46 predicted (although not statistically significant, $p = 0.14$). 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 ($p < 0.01$). It is worth noting that the high initial capital in the NR treatment implies that retailers rely heavily on their own financing instead of trade credit, which makes it similar to the Scheme O treatment in Chen et al. (2013). The slight overstocking behavior we observe is consistent with findings in their study.

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., $(q - q^*(w))/q^*(w)$, $p = 0.92$). 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 baseline 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 nonzero bankruptcy risk.

Turning to wholesale prices, the average observed value is significantly lower than predicted in HR, 19.91 vs. 25.87 ($p < 0.01$, 23.04% lower), and in LR, 19.05 vs. 21.47 ($p < 0.01$, 11.27% lower). The relative deviation $(w - w^*)/w^*$ is significantly larger in HR than in LR ($p < 0.01$). Hypothesis 2 is fully supported. Turning to NR, the observed wholesale price is slightly higher than predicted, 20.70 vs. 20.00 ($p < 0.05$, 3.5% higher). In other words, similar to retailers, suppliers exhibit a different direction of deviation in response to retailers' degree of bankruptcy risk. Although suppliers in NR set prices slightly higher than the baseline prediction, note that the absolute difference is relatively small. Combining these observations around wholesale prices yields the following result.

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

Finally, Table 2 also indicates that suppliers do not offer wholesale prices that, in theory, lead to lower bankruptcy risk for retailers relative to the baseline 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 ($p < 0.01$), although we emphasize that the absolute differences are relatively small.

5.2. Profits

Next, we examine how the observed decisions translate into expected profits, first, comparing to the baseline predictions, and second, comparing across treatments. To this end, Table 3 shows the observed average profits on the left side and the baseline 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 because we observe underpricing and understocking in HR and LR. As for NR, although suppliers achieve slightly higher profits than predicted on average, 757.49 vs. 733.33, the difference is not significant ($p = 0.45$).

In contrast to supplier profit, we observe higher average retailer profit in HR and LR relative to the baseline 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 vs. 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 3 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 baseline predictions (1,015.38 vs. 1,079.00 in HR and 1,265.71 vs. 1,349.60 in LR, both $p < 0.01$), because of insufficient order quantities. Last, we do not observe significantly different supply chain profits in NR versus the baseline prediction ($p = 0.31$).

Figure 4 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 baseline predictions. Also, there is

no significant difference between supplier profit in HR and LR ($p = 0.64$). Turning to retailer net profit, we observe an opposite trend compared with the baseline 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, $HR < LR < NR$. Recall that we did not have a formal behavioral hypothesis around profits but can now provide the following pertinent result.

Result 3. *In contrast to the baseline 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 ($HR > LR > NR$). Furthermore, supply chain profits are significantly below the baseline 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

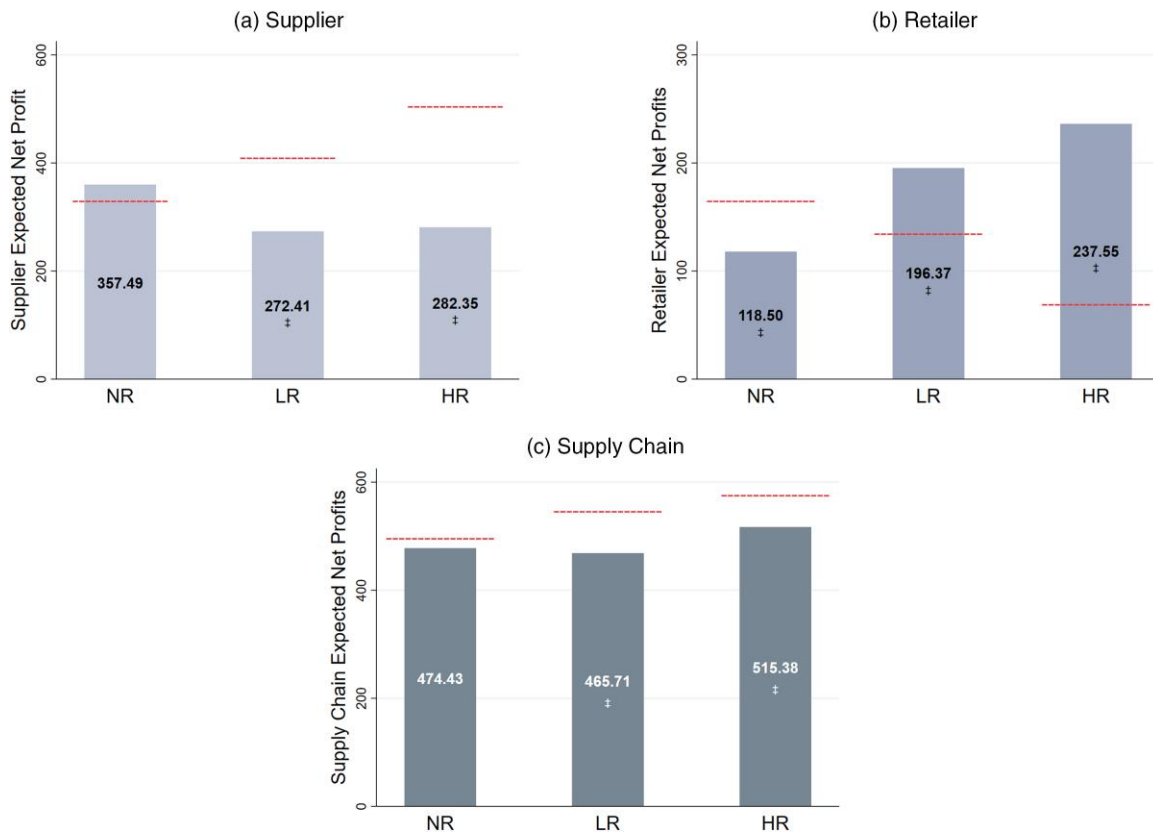
In this subsection, we conduct a heterogeneity analysis to examine whether our results are driven by a subset of participants (see Online Appendix D for distributions of observed decisions by treatment). Figure 5 classifies participants into three categories based on whether they consistently set lower (the bottom, light-green portions in the figure) or higher (the top, blue-green portions) values than optimal or did not significantly deviate in one direction (the middle, dark-green portions). Thus, we compare participants' actual decisions with (conditionally) optimal decisions in all rounds with a Wilcoxon signed-rank test, a significance level of 5%, and a sample size of 20 for each test before excluding rejections. Specifically, for the wholesale price, we compare a supplier's actual wholesale price and the optimal price according to the baseline theory (e.g., 25.87 in HR). For the quantity, we compare a retailer's actual quantity and the

Table 3. 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	1,574.43 (16.72)	1,265.71 [†] (13.75)	1,015.38 [†] (14.07)	1,600.00	1,349.60	1,079.00

Notes. Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Significance of regressions with random effects (standard errors clustered at the cohort level) compared with observed vs. baseline predictions is given by [†] $p < 0.01$.

Figure 4. (Color online) Average Observed Net Profit




Notes. Baseline predictions are represented by horizontal dashed lines. Significance of regressions with random effects (standard errors clustered at the cohort level) compared with observed vs. baseline predictions is given by † $p < 0.01$.

profit-maximizing quantity conditioning on the observed wholesale price in each round.

Beginning with the wholesale price decision in Figure 5(a), 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 overprice, 20.39% prefer lower prices, and 37.94% stand in between. Turning to the quantity decision of retailers in Figure 5(b), the results are mostly consistent with the previous averages. Specifically, 62.23% of retailers in NR overstock, whereas

Figure 5. (Color online) Percentages of Participants Deviating in Decisions



Notes. 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.

74.73% in LR and 86.13% in HR understock (and thus face a lower bankruptcy risk).

Taking a closer look at the three subgroups of suppliers in NR, the average wholesale prices of the “underpricing” group, the “in between” group, and the “overpricing” group are 18.29, 20.15, and 22.28, respectively. The absolute deviation from the optimal wholesale price, 20, is similar for both the “underpricing” and the “overpricing” groups, which can account for the overall average price being just slightly higher than the baseline prediction. Although deviations are less consistent in this case, we note that the “overpricing” 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.

6. Behavioral Models

So far, we have shown that observed decisions significantly deviate from the baseline theory. To understand what can account for these deviations, we now investigate behavioral models for both parties and test their performance by fitting them to the data.

6.1. Quantity: Reference Dependence

We consider a reference-dependence model for the quantity decision, which has been observed in various operational contexts (see, e.g., Ho et al. 2010, Baron et al. 2015, Tereyağoglu et al. 2018). We first introduce a general framework of reference dependence that follows Uppari and Hasija (2019) and subsequently discuss potential reference points.

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 a decision but is independent of demand, and c) a stochastic reference point (SRP) that depends on demand but is independent of a decision. 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 for the retailer is

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

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 nonnegative.

For the retailer, its initial capital r is a reasonable candidate for the reference point. Recall that, during the experiments, the retailer was allowed to reject a supplier's offer and keep its initial capital. Therefore, it is plausible that the retailer treats r as a target profit. Formally, $\mathcal{P} = r$ is an FRP that is independent of quantity and demand. Because the retailer is essentially making a newsvendor decision under a financial constraint, we also consider expected profit $\mathcal{P} = \mathbb{E}[\pi_r(q)]$ as a PRP and realized profit when the retailer orders the mean demand $\mathcal{P} = \pi_r(50, d)$ as a SRP, as suggested by Uppari and Hasija (2019). For the PRP, 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. For the SRP, the mean demand is a typical anchor point in the newsvendor model, and also, the retailer could see all of the realized profit outcomes if the mean demand was tested as a quantity. See Online Appendix A for theoretical details of each model.

To evaluate a reference point, we fit the observed decisions using maximum-likelihood estimation (MLE). In the estimation, we use a truncated normal distribution for the quantity decisions, with quantities that maximize Equation (8), given the observed wholesale prices, as the means and estimate the standard deviations. The distribution for the quantity is truncated at lower bound 0 and upper bound 100 of demand. 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, respectively, of the standard normal distribution.

We use the full sample for estimation and generate standard errors by bootstrapping. We then calculate the log-likelihood (LL) and compare across models using the Bayesian information criterion (BIC). Let I be retailers in the full sample (indexed by i). For retailer i , some rounds are excluded because of rejection. Therefore, we use T_i to denote the set of rounds that are included in the estimation, with t being the index. The total likelihood function for retailers is shown below:

$$\text{Likelihood}_r(\eta, \lambda, \sigma_r) = \prod_{i \in I} \prod_{t \in T_i} \varphi(q_{it}; \tilde{q}(w_{it}, \eta, \lambda), \sigma_r, 0, 100). \quad (10)$$

Note that $\tilde{q}(w_{it}, \eta, \lambda)$ 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 .

Table 4 shows the log-likelihood values of treatment specific for each reference point and the baseline model along with the Bayesian information criterion measures. Overall, the reference-dependence model, regardless of the reference point, outperforms the baseline model in terms of goodness of fit in each treatment. Comparing the BIC values that take the number of model parameters into account leads to the same conclusion; that is, all reference-dependence models achieve a lower BIC value than the baseline model. Among the three reference points, mean demand (SRP) performs the best in all treatments, especially in LR and HR, which are of particular interest in our study. In NR, reference dependence only slightly improves the fit over the baseline model, potentially because retailers in NR deviate only slightly from the baseline theory. The favorable performance of the mean demand (SRP) variant can be further supported by the average mean absolute percentage error (MAPE) between predicted and observed quantities, shown in the square brackets in Table 4. In short, the SRP achieves the lowest prediction error on average in all treatments. We note that this performance ranking (LL, BIC, and MAPE) remains unchanged if η and λ are constrained to be the same across treatments in the estimation (see Online Appendix B.2 for detailed results).

Given that the SRP performs best among the potential reference points, we report the gains and losses parameters for this setting in Table 5. To understand why the gains parameters are consistently greater than the losses parameters in all treatments, we first start with LR and HR. Note that the gains and losses regions with an SRP depend on whether the actual quantity is below or above the mean. If a quantity is below the mean, which is common in both the LR data (459 out of 546 decisions) and HR data (404 out of 548 decisions), then the gains region is actually when the realized demand is low, making it so that ordering below the mean is preferred. Therefore, a greater weight on the

Table 5. Estimated Parameters of SRP (Mean Demand)

	NR	LR	HR	Combined
Gains parameter: $\hat{\eta}$	3.55 (0.99)	3.42 (0.21)	3.67 (1.28)	3.46 (0.13)
Losses parameter: $\hat{\lambda}$	1.33 (1.56)	1.35 (0.14)	0.71 (0.49)	1.15 (0.16)

Notes. Parameters estimated using the full sample, with rejected data excluded. Standard errors in parentheses are derived from bootstrapping 1,000 times.

gains region leads to a lower predicted quantity, which fits the understocking behavior in these two treatments. Turning to NR, although ordering below the mean is also largely the case in NR (440 out of 564 decisions), there is much more heterogeneity in terms of understocking and overstocking behavior (as shown in Figure 5(b)), which may also explain why all of the MAPEs in NR are higher than those in LR and HR, for all models, in Table 4.

6.2. Wholesale Price: Fairness

Turning to the behavioral drivers for the supplier, who does not directly face bankruptcy risk, our experimental results are informative. Table 3 shows that profit distributions in LR and HR are more equitable than the theoretical predictions, that is, supplier profits being lower and retailer profits being higher, indicating that fairness can be a potential driver of wholesale price decision (Fehr and Schmidt 2000). Given that the supplier's initial capital can be higher or lower than the retailer's, we consider both advantageous and disadvantageous fairness concerns; that is, the supplier suffers from earning more or earning less than the retailer does. The utility function of a fairness-minded supplier is given by Equation (11).

$$u_s = \mathbb{E}[\pi_s(w)] - \theta_a(\mathbb{E}[\pi_s(w)] - \mathbb{E}[\pi_r(q(w))])^+ - \theta_d(\mathbb{E}[\pi_r(q(w))] - \mathbb{E}[\pi_s(w)])^+, \quad (11)$$

where θ_a and θ_d represent the supplier's psychological weight for advantageous fairness and disadvantageous

Table 4. Model Comparison of Log-likelihood and Prediction Error for Retailer Decisions

	Baseline model	FRP: initial capital r	PRP: expected profit $\mathbb{E}[\pi_r(q)]$	SRP: mean demand $\pi_r(50, d)$
NR	-2,460.33 [86.70%]	-2,423.18 [101.86%]	-2,416.98 [87.86%]	-2,378.08 [72.24%]
LR	-2,340.65 [70.64%]	-2,259.34 [45.87%]	-2,255.66 [45.39%]	-2,017.61 [34.60%]
HR	-2,448.45 [87.00%]	-2,299.33 [49.47%]	-2,293.37 [52.12%]	-2,133.16 [41.59%]
Total LL	-7,249.43	-6,981.85	-6,966.01	-6,528.85
BIC	14,521.10	14,030.42	13,998.74	13,124.42

Notes. Sample size N is 564 in NR, 546 in LR, and 548 in HR. Log-likelihoods (LL) are calculated using the full sample, with rejected data excluded. Average mean absolute percentage errors (MAPE) are presented in square brackets, calculated as $\frac{1}{N} \sum_{i \in I} \sum_{t \in T_i} \left| \frac{q_{it} - \hat{q}(w_{it}, \eta, \lambda)}{q_{it}} \right|$. BIC is Bayesian information criterion.

fairness, respectively. In Equation (11), we assume that the retailer's quantity response follows the baseline theory, consistent with the decision support tool provided to suppliers.

We continue to use MLE to fit the observed wholesale prices to determine the performance of the fairness model. Similar to our estimation for quantity decisions, we consider a truncated normal distribution as in Equation (9) for the wholesale price decisions, with wholesale prices that maximize Equation (11) as the means, and estimate the standard deviations. The normal distribution for the wholesale price is truncated at unit production cost $c = 10$ and unit selling price $p = 30$. Let J be suppliers in the full sample with rejected rounds excluded (indexed by j). With T_j denoting the set of rounds included in the estimation and t being the index, the total likelihood function for suppliers is shown below:

$$\text{Likelihood}_s(\theta_d, \theta_a, \sigma_w) = \prod_{j \in J} \prod_{t \in T_j} \varphi(w_{jt}; \tilde{w}(\theta_d, \theta_a), \sigma_w, 10, 30), \quad (12)$$

where $\tilde{w}(\theta_d, \theta_a)$ is the optimal wholesale price of the fairness model.

Table 6 compares treatment-specific and total LL values from the MLE for two fairness models plus the baseline model. For the fairness model, we consider a full model with both types of fairness concerns along with a nested fairness model where only one-sided fairness concern is included. In NR, only disadvantageous fairness is considered because the supplier has less initial capital than the retailer, whereas only advantageous fairness is considered in LR and HR because the supplier takes much less risk than the retailer in these two scenarios. The estimation results in Table 6 suggest that a fairness model can fit the wholesale price decisions much better than the baseline model, especially in LR and HR, where a larger magnitude of deviation is observed. Furthermore, treatment-specific LLs in Table 6 suggest that one-

sided fairness is sufficient to achieve the same goodness of fit, leading to a lower BIC value because of fewer parameters. Turning to the MAPEs in the square brackets, first note that the prediction errors are generally lower for the wholesale price compared with those for the quantity (seen in Table 4). The reason is that both parties' profits are more sensitive to the wholesale price, compared with the quantity, and thus there are generally smaller errors in the wholesale price. Regarding the MAPE values across models in Table 6, beginning with NR, neither fairness model generates lower prediction errors than the baseline model, which is expected because the improvement of LL brought by the fairness component is relatively small. This is also consistent with fact that suppliers in NR only slightly deviate from the baseline theory (20.70 observed vs. 20.00 predicted). Regarding LR and HR, the one-sided fairness model performs equal to the full-fairness model while generating much lower prediction errors compared with the baseline theory, 10.17% vs. 15.89% in LR and 12.88% vs. 33.53% in HR.

Table 7 displays the fitted parameters for the one-sided fairness model. The estimated disadvantageous fairness concern in NR is relatively small, 0.16, which is consistent with the slight wholesale price deviation in NR. One possible explanation is that although suppliers in NR have lower initial capital than retailers, they are not the ones directly facing bankruptcy risk. On the other hand, the estimated advantageous fairness parameters in LR and HR are 0.36 and 0.46, respectively. Compared with NR, the higher degree of fairness concerns in LR and HR explains the significant deviation in wholesale price in the two treatments. In LR, although both parties have the same initial wealth, the supplier takes much less risk than the retailer and is in an advantageous position, thus setting a lower wholesale price. As for HR, this effect is even more pronounced; a higher fairness concern suggests a larger deviation in wholesale prices, which we observed in HR.

Table 6. Model Comparison of Log-likelihood and Prediction Error for Supplier Decisions

	Baseline model	Full fairness model	One-sided fairness model
NR	−1,396.76 [10.79%]	−1,382.57 [10.79%]	−1,382.57 [10.79%]
LR	−1,437.62 [15.89%]	−1,238.40 [10.17%]	−1,238.40 [10.17%]
HR	−1,615.63 [33.53%]	−1,387.49 [12.88%]	−1,387.49 [12.88%]
Total LL	−4,450.01	−4,008.46	−4,008.46
BIC	8,922.26	8,083.64	8,061.40

Notes. Sample size N is 564 in NR, 546 in LR, and 548 in HR. Log-likelihoods (LL) are calculated using the full sample, with rejected data excluded. Average mean absolute percentage errors (MAPE) are presented in square brackets, calculated as $\frac{1}{N} \sum_{j \in J} \sum_{t \in T_j} \left| \frac{w_{jt} - \tilde{w}(\theta_d, \theta_a)}{w_{jt}} \right|$. BIC is Bayesian information criterion. "One-sided fairness" model means disadvantageous fairness only in NR and advantageous fairness only in LR and HR.

Table 7. Estimated Parameters of One-Sided Fairness

	NR	LR	HR
Advantageous fairness: $\hat{\theta}_a$	—	0.36 (0.02)	0.46 (<0.01)
Disadvantageous fairness: $\hat{\theta}_d$	0.16 (0.13)	—	—

Notes. Parameters estimated using the full sample, with rejected data excluded. Standard errors in parentheses are derived from bootstrapping 1,000 times.

6.3. Discussion

Although reference dependence and fairness can explain the observed deviations, it is possible that there are other possible behavioral factors. Here, we provide a brief discussion of them.

One may posit that the understocking behavior in LR and HR is due to probability weighting errors in which decision makers overweigh small probabilities and underweigh large probabilities. In our context, this would mean that retailers overestimate the bankruptcy probability and thus react by ordering less. Common modeling approaches include those of Tversky and Kahneman (1992) and Gonzalez and Wu (1999). Yet when we apply their frameworks to our setting, the model predicts overstocking behavior, which is inconsistent with our data. An alternative approach is to simply assume that the retailer has a weight on k that affects its perception of the bankruptcy probability. We show that it can fit the deviations we observe but does not outperform the SRP. Further discussion around probability weighting and estimation results are provided in Online Appendix B.4.

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 it is bankrupt or not. Disappointment aversion can account for the understocking in LR and HR but cannot explain the overstocking in NR. A similar bias is guilt aversion (Battigalli and Dufwenberg 2007); that is, the retailer has negative utility if its choice of quantity results in bankruptcy, which adversely affects the supplier's payoff.

Because the retailer may have more or less wealth than the supplier, depending on the treatment, it is possible that the retailer is also fairness minded. We find that fairness can indeed directionally capture the observed quantity deviations but is less successful at organizing the data than our model of reference dependence. A possible reason is that rather than comparing its own payoff with the supplier's payoff, the retailer focuses more on its own quantity decision and bankruptcy risk and/or profits (e.g., reference points).

Given the favorable performance of the reference dependence model for retailer-determined quantity decisions, it may potentially work for supplier-determined wholesale prices as well. We constructed the supplier's initial capital as an FRP and expected profits of both parties as two PRPs. All of them fit slightly better than the baseline model but are much worse than the fairness model (see Online Appendix B.5). The reason is that the supplier's realized profit is independent of demand unless the retailer is bankrupt. Therefore, the gain-loss utility will be constant in most cases, making the reference point less effective.

One may also consider that the supplier is rationally responding to the reference-dependent retailer. We fit the model of a profit-maximizing supplier and a reference-dependent retailer, with the SRP as the reference point to the supplier's data, by estimating the retailer's degree of reference dependence anticipated by the supplier. Results show that it fits slightly worse than the one-sided fairness model despite it requiring more parameters (see Online Appendix B.6). Therefore, the one-sided fairness model appears to be the better explanation for the supplier's decision, which is further supported in the next section, where we test a different level of r_s .

Although there are some alternative explanations for our data, none are perfect, nor have they been stress-tested. To this end, in the next section we conduct an alternative treatment in which we reduce the supplier's available capital. This allows us to investigate another theoretical prediction, namely that the supplier's capital does not affect the optimal wholesale price. Beyond this, it will also provide insight on the underlying mechanisms and, as we will show, cast doubt on retailer guilt aversion as a viable explanation for our data as well as speak to supplier and retailer fairness.

7. Robustness Check: Alternative Supplier Capital

For this new treatment, we lower the supplier's initial capital r_s from 400 to 100 and fix the retailer's initial capital r at 100, that is, the same as HR, where the retailer incurs the highest bankruptcy risk. The reasons for this are threefold. First, we observe the most significant deviations in HR (27.27% understocking on average), indicating a strong behavioral effect. Second, with $r_s = 100$ and $r = 100$, the supplier now faces non-trivial bankruptcy risk. Facing higher risk, the supplier is more likely to deviate from the theory. In other words, we designed this treatment in a way such that a behavioral effect is more likely to be observed. Following the theory in Section 3.2, the supplier also bears limited liability and will earn a profit of zero if it goes bankrupt; that is, the transfer

payment from the retailer together with the supplier's initial capital r_s are insufficient to cover the supplier's production cost. To maintain consistency with the original experiment, we did not show the supplier's bankruptcy probability to participants. In addition, in the case that the supplier's capital is insufficient to cover its production cost, we assume that the supplier will receive an interest-free bank loan to tease out the impact of any interest rate and focus on r_s , which we acknowledge as a necessary limitation of our design (see Section 8 for further discussion). Finally, because the reduced supplier capital puts the supplier in a more precarious situation, bankruptcy by the retailer is more likely to trigger bankruptcy by the supplier. A guilt-averse retailer would like to avoid this and should therefore underorder relative to the original HR treatment. We refer to this new treatment as "S-HR." Other parameters and experimental protocols follow from the main experiment. The theoretical predictions in S-HR are the same as HR, except for the supplier's expected profit and supply chain channel profit, because of the change of r_s , but the expected net profits remain unchanged. We recruited an additional 48 participants from the same subject pool for this new S-HR treatment.

Excluding the rejected data (11.46% in S-HR), Table 8 depicts the average decisions and bankruptcy thresholds of both parties, in S-HR and HR, compared with the theoretical predictions. The supplier bankruptcy threshold is the minimum demand such that the supplier will not go bankrupt, calculated as $(cq - r_s - r)^+ / p$. In S-HR, we observe directionally similar deviations as in HR; both wholesale prices and quantities are set lower than predicted. The average wholesale price in S-HR is 20.92, lower than the prediction of 25.87 ($p < 0.01$, 19.13% lower). Similarly, the average quantity of 36.38 is significantly lower than the conditionally

optimal prediction of 51.80 ($p < 0.01$, 29.77% lower). The lower quantities result in significantly lower bankruptcy risk for both retailers and suppliers. For retailers in S-HR, the predicted bankruptcy threshold on average is 30.51, whereas the actual threshold is much lower, 20.86 ($p < 0.01$). For suppliers, the actual bankruptcy threshold is 5.79 vs. 10.58 predicted ($p < 0.01$). Fitting these decisions to the behavioral models discussed in Section 6, the SRP for the retailer and one-sided fairness for the supplier continue to outperform the other models (Online Appendix B.3 provides a detailed discussion around how fairness is a key driver of behavior for suppliers but not retailers in S-HR). Estimation results are provided in Online Appendix B.1. Finally, comparing decisions of the two treatments directly to each other, we do not find statistically significant differences for wholesale price or the quantity, suggesting that guilt aversion is not the mechanism for the understocking behavior ($p = 0.30$ for wholesale price and $p = 0.14$ for quantity). Thus we have the following result.

Result 4. *When retailers bear high bankruptcy risk, whether suppliers are exposed to substantial bankruptcy risk or not leads to similar deviations in decisions, relative to the baseline theory: lower order quantities by retailers and lower wholesale prices by suppliers.*

We also investigate whether the similar decisions in S-HR and HR result in similar profits. Table 9 presents average net profit of both parties and the channel in S-HR and HR compared with predicted values. Note that net profit predictions are identical for both treatments because the only difference is the supplier's initial capital. Compared with the predictions, again we observe similar deviations in S-HR: significantly higher retailer net profit with significantly lower supplier net profit and channel net profit. Compared with

Table 8. Decision Comparison between S-HR and HR

	Observed		Predicted	
	S-HR	HR	S-HR	HR
Wholesale price	20.92 [‡] (0.77)	19.91 [‡] (0.47)	25.87	25.87
Quantity	36.38 [‡] (2.13)	40.55 [‡] (1.77)	51.80 (1.76) [42.49]	55.76 (0.67) [42.49]
Retailer bankruptcy threshold	20.86 [‡] (1.30)	22.66 [‡] (1.11)	30.51 (0.87) [33.31]	32.71 (0.32) [33.31]
Supplier bankruptcy threshold	5.79 [‡] (0.49)	1.22 [‡] (0.17)	10.58 (0.78) [7.50]	2.15 (0.24) [0.00]

Notes. Standard errors, across participants, are reported in parentheses. Rejected data are excluded. Predicted quantities and bankruptcy thresholds are conditioned on observed wholesale prices. Unconditional baseline predictions, when applicable, are reported in square brackets. Significance of regressions with random effects (standard errors clustered at the cohort level) compared with observed vs. (conditional) baseline predictions is given by [‡] $p < 0.01$.

Table 9. Net Profit Comparison between S-HR and HR

	Observed		Predicted	
	S-HR	HR	S-HR	HR
Supplier net profit	284.43 [‡] (13.99)	282.35 [‡] (9.29)	507.87	507.87
Retailer net profit	204.76 [‡] (16.72)	237.55 [‡] (9.54)	71.13	71.13
Supply chain net profit	480.56 [‡] (14.34)	515.38 [‡] (14.07)	579.00	579.00

Notes. Standard errors, across subjects, are reported in parentheses. Rejected data are excluded. Significance of regressions with random effects (standard errors clustered at the cohort level) compared with observed vs. baseline predictions is given by [‡] $p < 0.01$.

HR, net profits of both parties and the channel are not significantly different in S-HR ($p = 0.94$ for supplier, $p = 0.26$ for retailer, and $p = 0.17$ for channel). Putting aside these insignificant differences, the lower retailer and channel net profit in S-HR are due to two reasons. First, although wholesale price decisions in S-HR are not significantly different from those in HR, the average price in S-HR is slightly higher (Table 8), which hurts the retailer. Second, higher wholesale prices reduce retailers' quantities, which decreases channel profit.

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 supply chain 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 its 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 its 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 laboratory 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 baseline 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 (i.e., the limited liability effect). 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 lower their actual bankruptcy risk and that suppliers set lower wholesale prices to achieve a fairer profit distribution. 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 baseline 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 baseline theory. We further show that the supplier's risk has little impact on both parties' decisions in an additional treatment (S-HR) adapted from HR, consistent with what the baseline theory predicts.

To account for these anomalies, we develop a behavioral model consisting of a reference-dependent retailer and a fairness-minded supplier and show that it fits both parties' decisions reasonably well. In particular, we find that a stochastic reference point (realized profit if ordering the mean demand) performs the best in terms of capturing retailers' decisions in all treatments. For suppliers, advantageous fairness concerns are sufficient to capture their underpricing behavior in the presence of bankruptcy risk. We further show that the performance of the model is robust when both the supplier and the retailer bear substantial bankruptcy risk in S-HR.

We believe that 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 reduce the retailer's risk and possibly induce a higher quantity, it ultimately results in a supplier profit that is lower than the baseline predictions. Because of this, suppliers actually earn a higher profit when partnering with a retailer who is not capital constrained (which runs counter to theory) because retailers do not understock in such an environment. Regarding retailers, although it is reasonable that retailers understock to avoid bankruptcy, our data suggest that doing so forgoes the protection of limited liability and leads to a lower expected profit. In summary, our study provides managers with details around how supply chain decisions are made when alternative levels of bankruptcy risk are present, which in turn, sheds light on profit outcomes (thus helping with forecasting, planning, and other operational activities). 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 in a two-tier supply chain. Furthermore, 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 that 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, compared with a standard newsvendor case, the retailer's decision in LR and HR may be affected by purchasing through trade credit and bankruptcy risk. The two forces are interdependent in our setting, but separately quantifying the impact of each would help us better understand how people react to trade credit. Second, 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) showed that a rational supplier should always set its interest rate as a risk-free rate. Finally, we assume that the supplier receives an interest-free bank loan if its cash is insufficient to cover the production cost. However, the supplier's decision may be affected by the bank interest rate, which requires further investigation.

Acknowledgments

The authors acknowledge Sung Kwan Lee for helpful discussions and seminar participants at Xi'an Jiaotong-Liverpool University and Hong Kong University of Science and Technology. The authors thank Cornell University for its financial support. The authors also thank Feryal Erhun, the associate editor, and two reviewers for their comments.

References

- Aktas N, De Bodt E, Lobe F, Statnik JC (2012) The information content of trade credit. *J. Banking Finance* 36(5):1402–1413.
- Aral KD, Giambona E, Wang Y (2022) Buyer's bankruptcy risk, sourcing strategy, and firm value: Evidence from the supplier protection act. *Management Sci.* 68(11):7940–7957.
- Astvansh V, Jindal N (2022) Differential effects of received trade credit and provided trade credit on firm value. *Production Oper. 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 Oper. Management* 17(4):456–469.
- Barrot JN (2016) Trade credit and industry dynamics: Evidence from trucking firms. *J. Finance* 71(5):1975–2016.
- Battigalli P, Dufwenberg M (2007) Guilt in games. *Amer. Econom. Rev.* 97(2):170–176.
- Biais B, Gollier C (1997) Trade credit and credit rationing. *Rev. Financial Stud.* 10(4):903–937.
- Brander JA, Lewis TR (1986) Oligopoly and financial structure: The limited liability effect. *Amer. Econom. Rev.* 76(5):956–970.
- Breza E, Liberman A (2017) Financial contracting and organizational form: Evidence from the regulation of trade credit. *J. Finance* 72(1):291–324.
- Camerer CF, Ho TH (1994) Violations of the betweenness axiom and nonlinearity in probability. *J. Risk Uncertainty* 8(2):167–196.
- Chen CJ, Jain N, Yang SA (2023) The impact of trade credit provision on retail inventory: An empirical investigation using synthetic controls. *Management Sci.* 69(8):4591–4608.
- Chen L, Kök AG, Tong JD (2013) The effect of payment schemes on inventory decisions: The role of mental accounting. *Management Sci.* 59(2):436–451.
- Chen DL, Schonger M, Wickens C (2016) otree—An open-source platform for laboratory, online, and field experiments. *J. Behavioral Exp. Finance* 9:88–97.
- Chod J (2017) Inventory, risk shifting, and trade credit. *Management Sci.* 63(10):3207–3225.
- Dada M, Hu Q (2008) Financing newsvendor inventory. *Oper. Res. Lett.* 36(5):569–573.
- Dass N, Kale JR, Nanda V (2015) Trade credit, relationship-specific investment, and product market power. *Rev. Finance* 19(5):1867–1923.
- 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 Sci.* 60(11):2666–2683.
- Donohue K, Katok E, Leider S, eds. (2019) *The Handbook of Behavioral Operations* (Wiley, Hoboken, NJ).
- Eriksen KW, Kvaløy O (2010) Do financial advisors exhibit myopic loss aversion? *Financial Marketing Portf. Management* 24(2):159–170.
- Fehr E, Schmidt KM (1999) A theory of fairness, competition, and cooperation. *Quart. J. Econom.* 114(3):817–868.
- Fehr E, Schmidt KM (2000) Fairness, incentives, and contractual choices. *Eur. Econom. Rev.* 44(4–6):1057–1068.
- Giannetti M, Burkart M, Ellingsen T (2011) What you sell is what you lend? Explaining trade credit contracts. *Rev. Financial Stud.* 24(4):1261–1298.
- Goentzel J, Rice JB (2015) Guest voices: Managing supply chains is intertwined with financial management. Accessed August 10, 2019, <https://www.wsj.com/articles/guest-voices-managing-supply-chains-is-intertwined-with-financial-management-1437662343>.
- Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cogn. Psychol.* 38(1):129–166.
- Haigh MS, List JA (2005) Do professional traders exhibit myopic loss aversion? An experimental analysis. *J. Finance* 60(1):523–534.
- Ho TH, Lim N, Cui TH (2010) Reference dependence in multilocation newsvendor models: A structural analysis. *Management Sci.* 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 (Wiley, Hoboken, NJ), 35–88.
- Kouvelis P, Zhao W (2012) Financing the newsvendor: Supplier vs. bank, and the structure of optimal trade credit contracts. *Oper. Res.* 60(3):566–580.

- Kouvelis P, Zhao W (2016) Supply chain contract design under financial constraints and bankruptcy costs. *Management Sci.* 62(8):2341–2357.
- Kouvelis P, Zhao W (2018) Who should finance the supply chain? Impact of credit ratings on supply chain decisions. *Manufacturing Service Oper. 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 HH, Zhou J, Wang J (2018) Trade credit financing under competition and its impact on firm performance in supply chains. *Manufacturing Service Oper. Management* 20(1):36–52.
- Li J, Leider S, Beil D, Duenyas I (2021) Running online experiments using web-conferencing software. *J. Econom. Sci. Assoc.* 7(2):167–183.
- Murfin J, Njoroge K (2015) The implicit costs of trade credit borrowing by large firms. *Rev. Financial Stud.* 28(1):112–145.
- Murphy R (2018) How supply chain finance is offering companies a new cash source. Accessed August 6, 2019, <https://www.forbes.com/sites/forbesbusinessdevelopmentcouncil/2018/04/09/how-supply-chain-finance-is-offering-companies-a-new-cash-source/#daac23422f16>.
- Oechssler J, Schuhmacher F (2004) The limited liability effect in experimental duopoly markets. *Internat. J. Indust. Organ.* 22(2):163–184.
- Petersen MA, Rajan RG (1997) Trade credit: Theories and evidence. *Rev. Financial Stud.* 10(3):661–691.
- Seifert D, Seifert RW, Protopappa-Sieke M (2013) A review of trade credit literature: Opportunities for research in operations. *Eur. J. Oper. Res.* 231(2):245–256.
- Showalter DM (1995) Oligopoly and financial structure: Comment. *Amer. Econom. Rev.* 85(3):647–653.
- Tang CS, Yang SA, Wu J (2018) Sourcing from suppliers with financial constraints and performance risk. *Manufacturing Service Oper. Management* 20(1):70–84.
- Tereyağoglu N, Fader PS, Veeraraghavan S (2018) Multiattribute loss aversion and reference dependence: Evidence from the performing arts industry. *Management Sci.* 64(1):421–436.
- Tunca TI, Zhu W (2018) Buyer intermediation in supplier finance. *Management Sci.* 64(12):5631–5650.
- Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertainty* 5(4):297–323.
- Uppari BS, Hasija S (2019) Modeling newsvendor behavior: A prospect theory approach. *Manufacturing Service Oper. Management* 21(3):481–500.
- Wilson N, Summers B (2002) Trade credit terms offered by small firms: Survey evidence and empirical analysis. *Bus. Fin. Account.* 29(3–4):317–351.
- WTO (2019) Trade finance and the compliance challenge. Accessed December 25, 2020, https://www.wto.org/english/res_e/booksp_e/tradefinnace19_e.pdf.
- Yang SA, Birge JR (2018) Trade credit, risk sharing, and inventory financing portfolios. *Management Sci.* 64(8):3667–3689.