E-Companion of

When to Broadcast? Inventory Disclosure Policies for Online Sales of Limited Inventory

Zibo LIU
School of Management
Fudan University
Yangpu District, Shanghai 200433, P. R. China
liuzibo@fudan.edu.cn

Shi CHEN
Foster School of Business
University of Washington
Seattle, WA 98195
shichen@uw.edu

Kamran MOINZADEH
Foster School of Business
University of Washington
Seattle, WA 98195
kamran@uw.edu

Yong TAN
Foster School of Business
University of Washington
Seattle, WA 98195
ytan@uw.edu

Appendix A: Customer Survey

We conducted a customer survey via Amazon's MTurk to examine customers' behavior and preference when they are considering buying products from flash sales. We first present our survey contents and survey results, and then we do some brief analysis on the results.

Survey Consent:

This survey will ask you about your behavior when you are facing online sales with limited inventory (online flash sales). Typically, these sales offer a discounted product in limited time and with limited inventory and are called "flash sales". An example of flash sales is Amazon's lightning deals, which can be accessed by choosing "lightning deals" from "deal type" on the webpage https://www.amazon.com/gp/goldbox. You should know the definition of sales with limited inventory (flash sales) before you take this survey. You'll need to answer what you will do when you face different settings of online sales with limited inventory in this survey. The survey will take you roughly 20 minutes. If you agree to take the survey, please proceed.

No. of subjects taking part in the survey: 403

Questions:

About customer characteristics: (some questions used to get a general idea of the customers we are investigating.)

1. How often do you shop online?
   (a) Almost not shop online
   (b) Less than once a quarter
   (c) More than once a quarter but less than once a month
   (d) More than once a month but less than once a week
   (e) More than once a week

   Results: a: 1.24%; b: 2.23%; c: 14.39%; d: 49.13%; e: 33.00%.

2. How much do you know about flash sales?
   (a) Haven't heard about it before
   (b) Know what it is but did not participate in it
   (c) Checking flash sales website sometime and bought a few products
   (d) Checking flash sales website and buy products on it frequently

   Results: a: 0.75%; b: 18.86%; c: 70.47%; d: 9.93%.

3. If you try to buy something on flash sales website, what products would you check first?
   (a) Products with little time left on sales
   (b) Products with only a few units left
   (c) Products that are selling the fastest
   (d) Products that are of great deals

   Results: a: 11.17%; b: 7.69%; c: 10.67%; d: 70.47%.

4. If you know that inventory information (e.g., percentage of inventory claimed, or number of units left) about the product may be disclosed sometime during the sale, but it is concealed at the time when you check the sale, would you consider the inventory information more helpful to you compared to sales of the same product on a retailer when inventory information is never disclosed?

(a) Yes

(b) No

Results: a: 68.73%; b: 31.27%.

In following questions, you need to use scores 1-7 to show your willingness to buy the product. Typically, your willingness to buy at different scores can be represented as:

1: **Not interested** in the product at all

3: **A little bit interested** in the product (tend to not buy it, but with good quality signal will buy it)

5: **Moderately interested** in (you tend to buy it)

7: **Very interested** in (almost surely to buy it).

Suppose you are facing a flash sale for a product that lasts 2 hours, and your original willingness to buy without any other information (such as inventory and time left) is 4 (i.e., you're indifferent between buying and not buying). Answer the following questions: (answers to these questions are used in our brief analysis)

Scenario 1: The retailer does not disclose any information of inventory left during the sale:

1. Will you guess the claimed inventory when you come to the sale? (Note: if 25% inventory is claimed, it means that 75% of starting inventory is still available.)

   (a) Yes

   (b) No

   Results: a: 36.97%; b: 62.78%.

2. If you answer "Yes" in the previous question, please enter your estimated claimed inventory if you check the sale at different times. Note: if 25% inventory is claimed, then it means that 75% of starting inventory is still available. Sale lasts 2 hours, and 115 minutes after the beginning of the sale means that you have 5 minutes left before the sale ends.

   Results (taking average):

Minutes **after the start of the sale**	5	30	60	90	115
Estimated claimed inventory (%)	13.03	27.68	45.57	64.37	81.75

3. Based on what you have learned about flash sales, do flash sales have limited inventory?

   (a) Yes

   (b) No

4. Will you come back later to check the sale if you decide not to buy the product at first glance?

   (a) Yes, I will come back to check the sale later

   (b) No, I won't check the sale anymore.

   Results: a: 61.54%; b: 38.46%.

5. If you answer "Yes" in the previous question, would your willingness to buy change when you come back later during the sale? **Remember that your original willingness to buy without any information is 4**. Please enter your willingness to buy when you come back. Note: the first row of the table is minutes until the end of the sale, not minutes after the start of the sale. Your willingness to buy at different scores can be represented as:

1: **Not interested** in the product at all

3: **A little bit interested** in the product (tend to not buy it, but with good quality signal will buy it)

5: **Moderately interested** in (you tend to buy it)

7: **Very interested** in (almost surely to buy it).

Results (taking average):

Minutes **until the end of the sale (remaining time)**	90	60	30	5
Willingness to buy	3.47	3.95	4.58	4.95

Scenario 2: When the retailer always shows inventory information during the sale:

6. Please enter your willingness to buy the product using scores 1-7 in the following table. The table shows different situations when you check the sale. **Remember that your original willingness to buy without any information is 4**, and your willingness to buy can either go up or down. Note: if 25% inventory is claimed, then it means that 75% of starting inventory is still available. Sale lasts 2 hours, and 115 minutes after the beginning of the sale means that you have 5 minutes left before the sale end. Your willingness to buy at different scores can be represented as:

   1: **Not interested** in the product at all

   3: **A little bit interested** in the product (tend to not buy it, but with good quality signal will buy it)

   5: **Moderately interested** in (you tend to buy it)

   7: **Very interested** in (almost surely to buy it).

   Results (taking average):

% of inventory claimed / Minutes after the start of the sale	0%	25%	50%	75%	95%
5	2.75	3.35	4.06	4.75	5.17
30	2.65	3.16	3.77	4.39	4.96
60	2.54	2.94	3.54	4.05	4.58
90	2.45	2.87	3.35	3.89	4.42
115	2.35	2.76	3.24	3.78	4.26

7. Suppose that you first check the sale 30 minutes after the sale begins. **How long would you wait** before you come back to check the sale status if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.

   (a) 5 minutes

   (b) 10 minutes

   (c) 20 minutes

   (d) 30 minutes

   (e) 60 minutes (i.e., until 30 minutes before the sale ends)

   (f) 70 minutes (i.e., until 20 minutes before the sale ends)

   (g) 80 minutes (i.e., until 10 minutes before the sale ends)

   (h) 85 minutes (i.e., until 5 minutes before the sale ends)

(i) I won't come back to check the sale any more

Results: a: 3.72%; b: 12.41%; c: 10.67%; d: 16.13%; e: 23.82%; f: 3.72%; g: 4.22%; h: 2.98%; i: 21.84%.

8. Suppose that you first check the sale 60 minutes after the sale begins. **How long would you wait** before you come back to check the sale status if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.

   (a) 5 minutes

   (b) 10 minutes

   (c) 20 minutes

   (d) 30 minutes (i.e., until 30 minutes before the sale ends)

   (e) 40 minutes (i.e., until 20 minutes before the sale ends)

   (f) 50 minutes (i.e., until 10 minutes before the sale ends)

   (g) 55 minutes (i.e., until 5 minutes before the sale ends)

   (h) I won't come back to check the sale any more

   Results: a: 4.71%; b: 10.42%; c: 9.68%; d: 24.57%; e: 8.44%; f: 8.93%; g: 4.22%; h: 28.78%.

9. Suppose that you first check the sale 60 minutes after the sale begins. **How long would you wait** before you come back to check the sale status if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.

   (a) 5 minutes

   (b) 10 minutes

   (c) 20 minutes (i.e., until 10 minutes before the sale ends)

   (d) 25 minutes (i.e., until 5 minutes before the sale ends)

   (e) I won't come back to check the sale any more

   Results: a: 10.45%; b: 13.68%; c: 22.89%; d: 13.68%; e: 39.05%.

10. Please do not answer this question and leave it blank:

    (a) Won't buy it

    (b) Will buy it later

    (c) Will come back later to decide

    (d) Will buy it now

11. When you check the sale, what can be the reason for buying immediately or delaying the purchase? (You can choose multiple options for this question)

    (a) I will delay purchasing as I want to wait to see if the product is selling fast, which indicates that it is a popular product of good quality

    (b) I will delay purchasing to get more information about the product through other channels, discuss with friends and family to see if the product is worth it

    (c) If inventory is low, I will buy immediately so that I do not lose the chance to buy and own the product, as the product may be sold out if I wait

(d) Other reasons (please specify):____.

Results: a: 32.51%; b: 62.53%; c: 32.75%; d: 4.22%. Text of other reasons:

(i) If it's something I need and I know it's at a great price, I'll buy it immediately. It's rare that I research it and come back. If I don't need it enough to snap it up now, I'm probably not going to buy it at all.

(ii) Sometimes I need a little time to decide if I really want to spend the money and/or need the product

(iii) Price

(iv) As a check on myself from impulse buying

(v) If I need or want the item

(vi) I'll be very interested if the bargain is very good. If the product isn't something I need then I won't buy it.

(vii) I make my decisions pretty quickly. I don't focus on countdown clocks or number of items remaining. That has no effect on my choice. If I decide not to buy it then, I wont come back later so I don't delay it at all.

(viii) I will buy immediately if it is something I have been watching and know is being sold at a good price.

(ix) I will not check the inventory and purchase if I wanted it.

(x) Think through the purchasing process and ask if I really need this product (either now or at some later time)

(xi) I see it as a good deal and something I need or want and decide to buy. I don't care about the actions of others.

(xii) I will buy the product if I want it, not based on the amount sold in a vertain time, or if the inventory is really low

(xiii) Is it something I really need?

(xiv) It makes no difference

(xv) It's something I had already wanted/needed

(xvi) I'm fucking poor?

(xvii) Spend time consuming if this is the kind of purchase I really want to make at this time

Scenario 3: The retailer may disclose inventory information sometime during sale. Suppose, however, when you check the sale, inventory information is not yet disclosed. Note that the firm's disclosure time may be based on inventory claimed and time left:

12. Will you guess the claimed inventory when you come to the sale?

(a) Yes

(b) No

Results: a: 21.59%; b: 77.67%.

13. If you answer "Yes" in the previous question, please enter your estimated claimed inventory if you check the sale at a different time. Note: if 25% inventory is claimed, then it means that 75% of starting inventory is still available. Sale last 2 hours, and 115 minutes after the beginning of the sale means that you have 5 minutes left before the sale ends.

Results (taking average):

Minutes **after the start of the sale**	5	30	60	90	115
Estimated claimed inventory (%)	14.81	31.09	49.47	67.26	85.17

14. Please enter your willingness to buy the product use scores 1-7 in the following table. The table shows different situations when you first check out the sale. **Remember that your original willingness to buy without any information is 4**, and your willingness to buy can go up or down when you come to the sale and find that the inventory information is not disclosed yet. Sale lasts 2 hours, and 115 minutes after the beginning of the sale means that you have 5 minutes left before the sale ends. Your willingness to buy at different scores can be represented as:

    1: **Not interested** in the product at all

    3: **A little bit interested** in the product (tend to not buy it, but with good quality signal will buy it)

    5: **Moderately interested** in (you tend to buy it)

    7: **Very interested** in (almost surely to buy it).

    Results (taking average):

Minutes **after the start of the sale**	5	30	60	90	115
Willingness to buy	3.54	3.73	3.90	4.00	4.20

15. Based on what you know about flash sales, do flash sales have limited time?
    (a) Yes
    (b) No

16. Suppose that you first check the sale status 30 minutes after the sale begins. **How long would you wait** before you come back to check the sale again if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.
    (a) 5 minutes
    (b) 10 minutes
    (c) 20 minutes
    (d) 30 minutes
    (e) 60 minutes (i.e., until 30 minutes before the sale ends)
    (f) 70 minutes (i.e., until 20 minutes before the sale ends)
    (g) 80 minutes (i.e., until 10 minutes before the sale ends)
    (h) 85 minutes (i.e., until 5 minutes before the sale ends)
    (i) I won't come back to check the sale any more

    Results: a: 4.71%; b: 9.68%; c: 13.15%; d: 20.60%; e: 18.11%; f: 4.47%; g: 4.22%; h: 2.23%; i: 22.83%.

17. Suppose that you first check the sale 60 minutes after the sale begins. **How long would you wait** before you come back to check the sale status if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.
    (a) 5 minutes
    (b) 10 minutes
    (c) 20 minutes
    (d) 30 minutes (i.e., until 30 minutes before the sale ends)
    (e) 40 minutes (i.e., until 20 minutes before the sale ends)

(f) 50 minutes (i.e., until 10 minutes before the sale ends)
(g) 55 minutes (i.e., until 5 minutes before the sale ends)
(h) I won't come back to check the sale any more

Results: a: 4.47%; b: 9.68%; c: 12.90%; d: 22.58%; e: 9.68%; f: 7.69%; g: 3.97%; h: 28.78%.

18. Suppose that you first check the sale 60 minutes after the sale begins. **How long would you wait** before you come back to check the sale status if you decide not to buy the product at the first glance? Note: you can choose not coming back anymore.

(a) 5 minutes
(b) 10 minutes
(c) 20 minutes (i.e., until 10 minutes before the sale ends)
(d) 25 minutes (i.e., until 5 minutes before the sale ends)
(e) I won't come back to check the sale any more

Results: a: 7.94%; b: 17.62%; c: 22.83%; d: 11.91%; e: 39.70%.

19. When you check the sale, what can be the reason for buying immediately or delaying the purchase? (You can choose multiple options for this question)

(a) I want to wait until inventory information is available, so I can better assess how fast the product is sold, which is an indicator of the popularity and quality of the product
(b) I want to wait to get more information about the product through other channels, discuss with friends and family to see if the product is worth it
(c) If inventory is low, I will buy immediately so that I do not lose the chance to buy and own the product, as the product may be sold out if I wait
(d) Other reasons (please specify):____.

Results: a: 32.00%; b: 64.76%; c: 29.78%; d: 3.72%. Text of other reasons:

(i) I don't care about how much inventory there is. If it's a great price and something I need, I'll get it now. I rarely research something and come back. Almost never.
(ii) If I want it, I'm going to buy it immediately.
(iii) Again, price is everything
(iv) Again, as a check on myself to come back to make sure it's really want I want to avoid an impulse buy
(v) If I need or want the item.
(vi) The purchase has to be something worthwhile. It has to be a very good deal.
(vii) I make my decisions pretty quickly. I don't focus on countdown clocks or number of items remaining. That has no effect on my choice. If I decide not to buy it then, I wont come back later so I don't delay it at all.
(viii) I would buy immediately if it is a product I have been wanting and know that this is a good price
(ix) I will think about if I truly can afford it
(x) Ask whether I really need the product (either now or at some time in the future)
(xi) I won't come back unless I previously decided to purchase.
(xii) It makes no difference

(xiii) Needing the product

(xiv) Again, poor

(xv) Consider my budget

Note: Questions 3, 10, and 15 are used to test if the subjects are paying enough attention to the survey so that the answers can be used.

Some brief analysis:

Questions 5 and 14 show that the subjects' willingness to buy increases with the time elapsed. It is an evidence of the scarcity effect and in line with our assumption. Question 6, which comprehensively collects subjects willingness to pay under different situations when inventory information is available, indicates that customers' willingness to pay increases with percentage of inventory claimed and decreases with minutes after the start of the sale. This is true in every row and column of the table. This strongly supports our assumption on how the demand adjustment factor changes with claimed inventory and time elapsed.

Moreover, Question 2 and 13 shows that customers expect claimed inventory to be a linear function of time elapsed, which provides us the idea on how customers estimate current inventory when the inventory information is not available. Also, to investigate that if customers have different expectations under Scenario 1 and Scenario 3 when the inventory is not available, we compare the results of Questions 2 and 13. We have 101 valid observations of Question 2 and 47 valid observations of Question 13, which is enough for us to perform a group difference t-test to see if the differences of the results of Questions 2 and 13 are significant. The group difference t-test can be calculated as:

$$\frac{\mu_{Q13} - \mu_{Q2}}{\sqrt{\frac{Var_{Q13}+Var_{Q2}}{n_{Q13}+n_{Q2}}}}$$

In the above expression, μ is group mean, Var is group variance, and n is number of observations in the group. $Q2$ and $Q13$ indicate two groups we want to compare. The detailed t-test result is as follows:

Minutes **after the start of the sale**	5	30	60	90	115
Estimated claimed inventory of Q2 (%)	13.03	27.68	45.57	64.37	81.75
Estimated claimed inventory of Q13 (%)	14.81	31.09	49.47	67.26	85.17
Difference	1.78	3.41	3.90	2.89	3.42
Group difference t-statistic:	1.01	1.52	1.37	0.98	1.20

As none of the absolute values of group difference t-statistics exceed 1.96, the difference between two groups is insignificant at a 5% significant level. This result supports our reasoning on why we treat the situation of "never disclose" and where inventory information may be available in the future the same.

Appendix B: Mathematical Proofs.

Proof of Lemma 1. As Δ_x is a quadratic function of $q_x(Ks)$, we can solve $\Delta_x > 0$ to get:

$$q_x(Ks) > \frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)$$

As we have a reasonable assumption that customers with $q_{x,0} \geq q_{TQ}$ will buy the product immediately, we now focus on those customers with $q_{x,0} < q_{TQ}$. In this case, $q_x(Ks) = q_{TQ} - (1-\rho(s))(q_{TQ} - q_{x,0})$. Together with the inequality above, we can finally get the condition on $q_{x,0}$ for a customer to buy the product immediately:

$$q_{x,0} \geq q_{TQ} - \frac{q_{TQ} - \frac{1}{b}(c + \delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}))}{1-\rho(s)}$$

□

Proof of Lemma 2. As the probability of a future arrival for a customer of segment x during the remaining time of the sale is uniformly distributed, we have the probability of stock-out to be $\frac{1}{1-\tau}\int_{\frac{(1-s)\tau}{\eta s}\wedge(1-\tau)}^{(1-\tau)} 1 d\tau$, which can be translated to:

$$p(s,\tau|\eta) = \begin{cases} 1 - \frac{(1-s)\tau}{\eta s(1-\tau)} & \text{if } s > \frac{\tau}{\eta(1-\tau)+\tau} \\ 0 & \text{if } s \leq \frac{\tau}{\eta(1-\tau)+\tau} \end{cases}$$

If $s \leq \frac{\tau}{\overline{\eta}(1-\tau)+\tau}$, we have $p(s,\tau) = 0$. If $\frac{\tau}{\overline{\eta}(1-\tau)+\tau} < s \leq \frac{\tau}{\underline{\eta}(1-\tau)+\tau}$, then:

$$p(s,\tau) = \frac{1}{\overline{\eta}-\underline{\eta}}\int_{\frac{(1-s)\tau}{s(1-\tau)}}^{\overline{\eta}} \left(1 - \frac{(1-s)\tau}{\eta s(1-\tau)}\right) d\eta = \frac{1}{\overline{\eta}-\underline{\eta}}\left(\overline{\eta} - \frac{(1-s)\tau}{s(1-\tau)} - \frac{(1-s)\tau}{s(1-\tau)}\ln\left(\frac{\overline{\eta}s(1-\tau)}{(1-s)\tau}\right)\right)$$

If $s > \frac{\tau}{\underline{\eta}(1-\tau)+\tau}$, then:

$$p(s,\tau) = \frac{1}{\overline{\eta}-\underline{\eta}}\int_{\underline{\eta}}^{\overline{\eta}} \left(1 - \frac{(1-s)\tau}{\eta s(1-\tau)}\right) d\eta = 1 - \frac{1}{\overline{\eta}-\underline{\eta}}\frac{(1-s)\tau}{s(1-\tau)}\ln\left(\frac{\overline{\eta}}{\underline{\eta}}\right)$$

So we finally have:

$$p(s,\tau) = \begin{cases} 0 & \text{If } s \leq \frac{\tau}{\overline{\eta}(1-\tau)+\tau} \\ \frac{1}{\overline{\eta}-\underline{\eta}}\left(\overline{\eta} - \frac{(1-s)\tau}{s(1-\tau)} - \frac{(1-s)\tau}{s(1-\tau)}\ln\left(\frac{\overline{\eta}s(1-\tau)}{(1-s)\tau}\right)\right) & \text{If } \frac{\tau}{\overline{\eta}(1-\tau)+\tau} < s \leq \frac{\tau}{\underline{\eta}(1-\tau)+\tau} \\ 1 - \frac{1}{\overline{\eta}-\underline{\eta}}\frac{(1-s)\tau}{s(1-\tau)}\ln\left(\frac{\overline{\eta}}{\underline{\eta}}\right) & \text{If } s > \frac{\tau}{\underline{\eta}(1-\tau)+\tau} \end{cases}$$

□

Proof of Proposition 1. We first take a look at the situation where inventory information is not available. In this case, when the customer arrives at τ, she estimates the current claimed inventory to be $\widetilde{s}$, and only a customer with $q_{x,0} \geq \frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(\widetilde{s},\tau)}}{1+\sqrt{p(\widetilde{s},\tau)}}\right)$ will purchase the product. Similar to the analysis of the queues with limited waiting space, we assume that there is an infinite population of potential arriving customers. This implies that if the customer decides not to buy immediately upon arrival, then she joins the pool of customers and is not tracked (loss queue systems). Thus, let $H(\cdot)$ denotes the probability distribution function of $q_{x,0}$, the proportion of customers buying the product is:

$$1 - H\left(\frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(\widetilde{s},\tau)}}{1+\sqrt{p(\widetilde{s},\tau)}}\right)\right)$$

Together with the assumption that customers' prior belief of the product quality $q_{x,0}$ follows a uniform distribution $U(q_{min}, q_{max})$ and the average arrival rate to be $\Lambda(t)$, we derive the demand rate as:

$$\theta(t) = \Lambda(t)\left(1 - H\left(\frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(\widetilde{s},\tau)}}{1+\sqrt{p(\widetilde{s},\tau)}}\right)\right)\right) = \Lambda(t)\frac{q_{max} - \frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(\widetilde{s},\tau)}}{1+\sqrt{p(\widetilde{s},\tau)}}\right)}{q_{max} - q_{min}}$$

Similarly, when the inventory information is available, according to Lemma 1, only customers with $q_{x,0} \geq q_{TQ} - \frac{q_{TQ} - \frac{1}{b}\left(c+\delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)}{1-\rho(s)}$ will purchase the product. As customers not buying are not tracked, the proportion of customers purchasing the product should be:

$$1 - H\left(q_{TQ} - \frac{q_{TQ} - \frac{1}{b}\left(c+\delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)}{1-\rho(s)}\right)$$

Finally, considering the uniformly distributed $q_{x,0}$ and the average arrival rate of all customers $\Lambda(t)$, we have the demand rate as:

$$\begin{aligned}\lambda_{Ks}(T\tau) &= \Lambda(t)\left(1 - H\left(q_{TQ} - \frac{q_{TQ} - \frac{1}{b}\left(c+\delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)}{1-\rho(s)}\right)\right)\\ &= \Lambda(t)\frac{q_{max} - q_{TQ} + \frac{q_{TQ} - \frac{1}{b}\left(c+\delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)}{1-\rho(s)}}{q_{max} - q_{min}}\end{aligned}$$

□

Proof of Proposition 2. First, we show that the general form of solutions to (9)-(10) is as follows:

$$\begin{aligned}Q_0(t|\zeta,j) &= C_0 e^{-\int_\zeta^t \lambda_j(s)ds},\\ Q_i(t|\zeta,j) &= e^{-\int_\zeta^t \lambda_{j+i}(s)ds}\left[\int_\zeta^t e^{\int_\zeta^s \lambda_{j+i}(x)dx}\lambda_{j+i-1}(s)Q_{i-1}(s|\zeta,j)ds\right], i = 1,\cdots,K-j,\end{aligned}$$

where C_0 has a non-negative value such that the boundary conditions are satisfied. Note that,

$$\begin{aligned}\frac{dQ_0(t|\zeta,j)}{dt} &= -\lambda_j(t)\cdot C_0 e^{-\int_\zeta^t \lambda_j(s)ds} = -\lambda_j(t)Q_0(t|\zeta,j)\\ \frac{dQ_i(t|\zeta,j)}{dt} &= -\lambda_{j+i}(t)Q_i(t|\zeta,j) + e^{-\int_\zeta^t \lambda_{j+i}(s)ds}\left[e^{\int_\zeta^t \lambda_{j+i}(x)dx}\lambda_{j+i-1}(t)Q_{i-1}(t|\zeta,j)\right]\\ &= -\lambda_{j+i}(t)Q_i(t|\zeta,j) + \lambda_{j+i-1}(t)Q_{i-1}(t|\zeta,j).\end{aligned}$$

Furthermore, in order to satisfy the boundary conditions, we should have $C_0 = 1$ such that $Q_0(\zeta|\zeta,j) = 1$ and $Q_i(\zeta|\zeta,j) = 0$ for any $i \in \{1,\cdots,K-j\}$.

□

Proof of Theorem 1. We first show that the following statement is true: If it is optimal to disclose inventory at time t when j sales have been made until t, then it is always better to reveal than conceal inventory at any time $t' < t$ when j sales have been made until t'.

Suppose that at time τ instead of t, where $t - \tau = \Delta$ is an infinitesimal time interval, the j-th sale occurs and the retailer has not broadcasted its inventory. If the retailer decides to reveal its inventory immediately at τ, then its expected future sales would be

$$(1-\lambda_j(\tau)\Delta)\Gamma(T|t,j) + \lambda_j(\tau)\Delta[\Gamma(T|t,j+1)+1] + o(\Delta). \quad \text{(B.1)}$$

In contrast, if the retailer decides not to reveal its inventory immediately, its expected future sales is:

$$(1-\theta(\tau)\Delta)\Gamma(T|t,j) + \theta(\tau)\Delta[\Gamma(T|t,j+1)+1] + o(\Delta). \quad \text{(B.2)}$$

In (B.1) or (B.2), with probability $(1-\lambda_j(\tau)\Delta)$ or $(1-\theta(\tau)\Delta)$, respectively, there will be no additional sale made in $[\tau, t]$. Then, at time t, the retailer will start to broadcast its inventory as it is assumed to be optimal to broadcast inventory levels if there are j sales until t, and, as a result, the expected future sales thereafter will be $\Gamma(T|t,j)$. Otherwise, with probability $\lambda_j(\tau)\Delta$ or $\theta(\tau)\Delta$, respectively, there will be $j+1$ sales until t. In that case, the retailer will also start to broadcast its inventory at t, and the expected future sales thereafter will be $\Gamma(T|t,j+1)$. The reason is that, if the retailer finds itself better off revealing than concealing its inventory at time t with j sales, it must also find itself better off revealing than concealing its inventory at the same time t with more sales (e.g., $j+1$ sales) made, as, according to Assumption 2, $\beta(j+1,t) > \beta(j,t)$ and thus $\lambda_{j+1}(t) > \lambda_j(t)$.

Next, note that a necessary condition for revealing inventory information at time t with j sales is: $\lambda_j(t) \geq \theta(t)$, which implies $\beta(j,t) \geq 1$. Then, according to Assumption 2, we have $\beta(j,\tau) > \beta(j,t) \geq 1$ given that $\tau < t$, which implies $\lambda_j(\tau) > \theta(\tau)$ as well. In addition, it is clear that $\Gamma(T|t,j+1)+1 > \Gamma(T|t,j)$. That is, we have $\lambda_j(\tau) > \theta(\tau)$ and $\Gamma(T|t,j+1)+1 \geq \Gamma(T|t,j)$, which together establish the fact that (B.1) is greater than (B.2), so the retailer must be better off revealing its inventory at τ if there are j sales made, given that it is also optimal to do so at time $t = \tau + \Delta$. Using backward induction, the conclusion holds for any time $t' < t$.

Furthermore, since the above statement is true, there must exist a threshold t_j^* such that the retailer should broadcast inventory information if and only if j sales are made before t_j^*. The reason is as follows. Suppose that it will be better to reveal than conceal inventory information at certain time t when j sales are made, then the retailer should also reveal its inventory at any time $t' \leq t$ when the j-th sale occurs. Let t_j^* be the largest t such that the retailer prefers revealing to concealing its inventory, given that there are j sales made. As a result, if the j-th sale occurs before t_j^*, it is optimal to broadcast inventory; otherwise, it is optimal to conceal. Such a policy is exactly what the proposed time-dependent threshold policy describes.

Last, the monotonicity of t_j^*'s is a result of Assumption 1, which indicates that the retailer cannot "flip-flop" its inventory disclosure decisions. Specifically, suppose that there exists certain j such that $t_{j+1}^* < t_j^*$. As long as the j-th sale occurs before t_j^*, the retailer should have already revealed its inventory according to the proposed time-dependent policy. After that, even though the $(j+1)$-th sale may occur after t_{j+1}^*, it does not make a difference as the retailer will not reverse its previous disclosure decision due to Assumption 1. In that case, indeed we can set $t_{j+1}^* = t_j^*$ without affecting the outcome of implementing the proposed policy.

□

Proof of Proposition 3. Recall that $V^*(T|t,j)$ is the expected "sales-to-go" given that the retailer does not disclose its inventory at the current time t when there are j sales. However, the retailer may decide to disclose its inventory in the future.

The retailer may start to broadcast its inventory upon the $(j+1)$-th sale in the future. The probability is $q_1(y|t,j)$ given that the $(j+1)$-th sale occurs at $[y-dy, y]$. In such a case, the expected "sales-to-go" includes the one additional sale that will occur at $[y-dy, y]$ and the expected sales thereafter, $\Gamma(T|y,j+1)$. Note that $y \in [t, t_{j+1}^*]$ because, otherwise, the retailer will not optimally broadcast its inventory at $[y-dy, y]$ upon the occurrence of the $(j+1)$-th sale.

Similarly, the retailer may start to broadcast its inventory levels upon the $(j+i)$-th sale in the future, where $i \in \{2, \cdots, K-j\}$. The probability is $q_i(y|t,j)$ given that the $(j+i)$-th sale occurs at $[y-dy, y]$. The expected "sales-to-go" includes the i additional sales that will occur during $[t, y]$ and the expected sales thereafter, $\Gamma(T|y,j+i)$. It is clear that $y \le t^*_{j+i}$ because, otherwise, the retailer will not optimally disclose its inventory at $[y - dy, y]$ upon the occurrence of the $(j+i)$-th sale. We also note that $y > t^*_{j+i-1}$: suppose $y \le t^*_{j+i-1}$, then the $(j+i-1)$-th sale must have occurred before t^*_{j+i-1}, implying that the retailer must have already disclosed its inventory upon the occurrence of the $(j+i-1)$-th sale at the latest.

Last, the retailer may conceal its inventory until the end of the sale. The third term of (18) gives the expected "sales-to-go" given that the retailer conceals its inventory until the end of the sale period. Specifically, with probability $r_n(T|t,j)$, there will be n additional sales while the retailer keeps concealing the inventory levels during $[t,T]$. Note that, since $t^*_K = T$, n cannot exceed $K - j - 1$ as, otherwise, the retailer would have revealed its inventory during $[t,T]$.

Now we turn to calculate $q_i(y|t,j)$ and $r_i(y|t,j)$. In (19), with probability $r_n(t^*_{j+i-1}|t,j)$, there will be n additional sales and the retailer will not broadcast its inventory during $[t, t^*_{j+i-1}]$. Note that n cannot be equal to or greater than $i-1$ as, otherwise, the retailer would have broadcast its inventory. Next, with probability density $\theta(y)p\big(i-1-n, \Theta(t^*_{j+i-1}, y)\big)$, there will be $i-n$ more sales in $(t^*_{j+i-1}, y]$, and, in particular, the $(i-n)$-th additional sale since t^*_{j+i-1} (i.e., the i-th additional sale since t) will occur at $[y-dy, y]$.

For the calculation of $r_i(y|t,j)$ in (20) and (21), first, as $r_n(t^*_{j+k}|t,j)$ denotes the probability that n sales come and the retailer does not disclose the inventory information, the $(j+1)$-th sale cannot come before t^*_{j+1}, so we have (20). Next, in (21), with probability $r_m(t^*_{j+k-1}|t^*_{j+1}, j)$, there will be $m \le n$ additional sales and the retailer will not disclose its inventory in $[t^*_{j+1}, t^*_{j+k-1}]$. Note that m cannot be equal to or greater than $k-1$ as, otherwise, the retailer would have disclosed its inventory before t^*_{j+k-1}. Moreover, with probability $p\big(n-m, \Theta(t^*_{j+k-1}, t^*_{j+k})\big)$, there will be $n-m$ more sales in $(t^*_{j+k-1}, t^*_{j+k}]$. Put these probabilities together, (21) gives the probability that there will be n additional sales and the retailer will not disclose its inventory in $[t^*_{j+1}, t^*_{j+k}]$. As indicated, n cannot exceed $k-1$; otherwise, the retailer would have disclosed its inventory before t^*_{j+k}.

Now we show the first two steps, $j = K-1$ and $j = K-2$, which helps to better understand the calculating procedure in Proposition 3. First, when $j = K$, we can set $t^*_K = T$ without loss of generality, so we start from $j = K-1$. The goal is to find $t^*_{K-1} \in [0, t^*_K] = [0, T]$.

As $q_1(y|t, K-1)$ denotes the probability density function of the following event: given that there have been $K-1$ sales made until t, the disclosure of inventory will be at $[y-dy, y]$ when the first future demand occurs. Then, $V^*(T|t, K-1)$ can be written as

$$V^*(T|t, K-1) = \int_t^{t^*_K = T} [1 + \Gamma(T|y, K)] q_1(y|t, K-1) dy. \tag{B.3}$$

Note that $q_1(y|t, K-1) = \theta(y)p(0, \Theta(t,y))$ and $\Gamma(T|y,K) = 0$. Thus, (B.3) becomes

$$V^*(T|t, K-1) = \int_t^{t^*_K = T} \theta(y)p(0, \Theta(t,y)) dy = P(1, \Theta(t,T)).$$

This makes sense as, with probability $P(1, \Theta(t,T)) = \sum_{x=1}^{\infty} p(x, \Theta(t,T))$, there will be one more sale during $[t,T]$ given that there is only one item left at t.

Hence, t^*_{K-1} will be the largest value of t such that $\Gamma(T|t, K-1) \ge V^*(T|t, K-1) = P(1, \Theta(t,T))$; that is, the expected "sales-to-go" if the retailer discloses its inventory at t is greater than or equal to the expected "sales-to-go" if the retailer does not disclose its inventory at that time. If $\Gamma(T|t, K-1) < V^*(T|t, K-1) = P(1, \Theta(t,T))$ for all t, we define $t^*_{K-1} = 0$ and all our following results still hold.

Next, suppose we have obtained t^*_{K-1} and consider $j = K-2$. The goal is to find $t^*_{K-2} \in [0, t^*_{K-1}]$. Then, $V^*(T|t, K-2)$ can be written as

$$V^*(T|t, K-2) = \tag{B.4}$$
$$\int_t^{t^*_{K-1}} [1 + \Gamma(T|y, K-1)] q_1(y|t, K-2) dy + \int_{t^*_{K-1}}^T [2 + \Gamma(T|y, K)] q_2(y|t, K-2) dy + r_1(T|t, K-2).$$

In (B.4), the first term is the expected future sales, providing that the following scenario happens: The first future demand (i.e., the $(K-2+1)$-th sale) will occur at $[y-dy, y]$ and the retailer will start to disclose its inventory from then on, which requires $y \le t^*_{K-1}$. In that case, the expected sales during $[t, T]$ will be $1 + \Gamma(T|y, K-1)$. The second term can be explained in a similar way. But, since the retailer will disclose its inventory at time $[y-dy, y]$ upon the occurrence of the second future demand (i.e., the $(K-2+2)$-th sale), we must have $y > t^*_{K-1}$ because, otherwise, the retailer would have disclosed its inventory upon the occurrence of the first future demand (i.e., the $(K-2+1)$-th sale). The third term gives the expected sales from t to T given that the retailer will not disclose its inventory until T. In that case, demand during $[t, T]$ cannot be greater than or equal to two items as, otherwise, the retailer would have disclosed its inventory (note that $t^*_K = T$).

Furthermore, for any $t \le t^*_{K-1}$,

$$q_1(y|t, K-2) = r_0(t|t, K-2)\theta(y)p(0, \Theta(t, y)) = \theta(y)p(0, \Theta(t, y));$$
$$q_2(y|t, K-2) = r_0(t^*_{K-1}|t, K-2)\theta(y)p(1, \Theta(t^*_{K-1}, y)).$$

Here, we provide an explanation for the expression of $q_2(y|t, K-2)$. In order to have the retailer disclose its inventory upon the occurrence of the second future demand at $[y-dy, y]$, two conditions are required: (1) There will be no demand during $[t, t^*_{K-1}]$ because, otherwise, the retailer would have disclosed its inventory upon the occurrence of the first future demand; the probability of such a condition is $r_0(t^*_{K-1}|t, K-2)$. (2) There will be demand for two items during $(t^*_{K-1}, y]$, and, in particular, the second future demand will occur at $[y-dy, y]$; the probability density of such a condition is $\theta(y)p(1, \Theta(t^*_{K-1}, y))$.

Therefore, calculation of (B.4) reduces to a problem of calculating $r_0(t^*_{K-1}|t, K-2)$ and $r_1(T|t, K-2) = r_1(t^*_K|t, K-2)$, which actually have closed-form expressions. Now,

$$r_0(t^*_{K-1}|t, K-2) = p(0, \Theta(t, t^*_{K-1}))r_0(t^*_{K-1}|t^*_{K-1}, K-2) = p(0, \Theta(t, t^*_{K-1}));$$
$$r_1(t^*_K|t, K-2) = p(0, \Theta(t, t^*_{K-1}))r_1(t^*_K|t^*_{K-1}, K-2) = p(0, \Theta(t, t^*_{K-1}))p(1, \Theta(t^*_{K-1}, t^*_K)).$$

Finally, once the calculation of (B.4) is done, t^*_{K-2} will be the largest value of t such that $\Gamma(T|t, K-2) \ge V^*(T|t, K-2)$. Now, we have completed the calculation for $j = K-2$.

We can also find the expressions of $q_i(y|t, j)$ and $r_i(y|t, j)$ when $j = K-2$ or $j = K-1$ following (19)-(21). For instance, for $j = K-2$, (19) reduces to $q_1(y|t, K-2) = \theta(y)p(0, \Theta(t, y))$ and $q_2(y|t, K-2) = r_0(t^*_{K-1}|t, K-2)\theta(y)p\big(1, \Theta(t^*_{K-1}, y)\big)$. Moreover, (20) and (21) together lead to $r_0(t^*_{K-1}|t, K-2) = p(0, \Theta(t, t^*_{K-1}))r_0(t^*_{K-1}|t^*_{K-1}, K-2) = p(0, \Theta(t, t^*_{K-1}))$ when $k = 1$ and $n = 0$, and $r_1(t^*_K|t, K-2) = p(0, \Theta(t, t^*_{K-1}))r_1(t^*_K|t^*_{K-1}, K-2) = p(0, \Theta(t, t^*_{K-1}))p(1, \Theta(t^*_{K-1}, t^*_K))$ when $k = 2$ and $n = 1$. These results are consistent with what we have derived for the case of $j = K-2$.

One can verify that, when $j = K-1$, (18) reduces to $V^*(T|t, K-1) = \int_t^{t^*_K} [1 + \Gamma(T|y, K)] q_1(y|t, K-1) dy$, which is indeed (B.3). When $j = K-2$, (18) reduces to $\int_t^{t^*_{K-1}} [1 + \Gamma(T|y, K-1)] q_1(y|t, K-2) dy + \int_{t^*_{K-1}}^{t^*_K} [2 + \Gamma(T|y, K)] q_2(y|t, K-2) dy + r_1(T|t, K-2)$, which is (B.4).

□

Appendix C: Numerical Study Extensions

Impact of b. Figure C.1 shows all relative improvements when b varies. As we discussed the behaviors of R_2 and R_4 in the main body of the paper, we only analyze R_1 and R_3 here. Note that R_1 increases and then decreases with b. When b is small, the observational learning (herding) effect has little impact on the customers' purchase decisions, which explains why "always disclose" is inferior to "never disclose" in such cases. As b increases, customers' purchasing decisions are more influenced by herding, which benefits the "always disclose" policy. However, when b is too large, more customers have a high enough utility of buying the product immediately, and the product will be almost sold-out at the end of the sale regardless of the inventory disclosure policy. As the result, when b is large, increasing b narrows the gap between two simple policies.

Figure C.1 R_1 **and** R_3 **as** b **varies.**

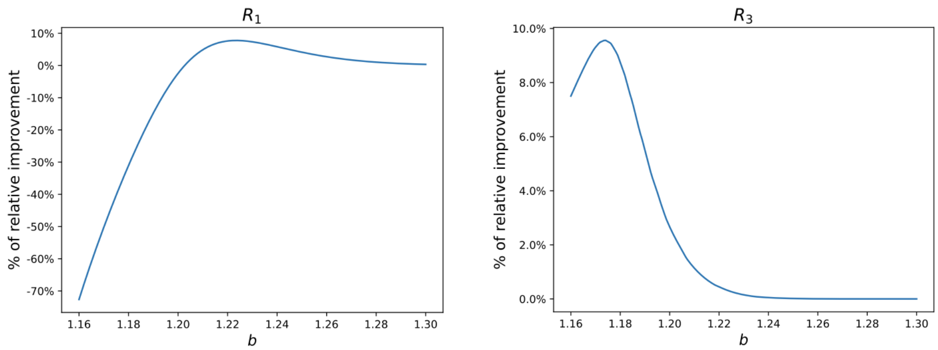

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

Next, we examine the additional improvement of switching from the optimal fixed threshold policy to the optimal time-dependent threshold policy. As can be seen, R_3 has a unimodal behavior with respect to b. To better analyze the behavior of R_3, Figure C.2 shows the optimal thresholds under both the fixed and time-dependent threshold policies as b varies. The vertical axis of the left graph represents the optimal fixed threshold, while the vertical axis of the right graph reveals the optimal time thresholds for different levels of inventory sold. Our observations are as below.

Figure C.2 **The optimal thresholds of the fixed (left) and time-dependent (right) threshold policies as** b **varies.**

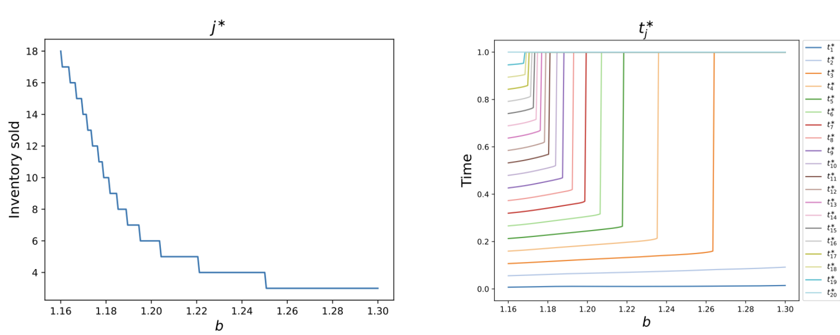

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

- The optimal fixed threshold decreases as b increases. This is intuitive, as a small b implies that the observational learning effect is insignificant. As b becomes large, the observational learning effect becomes the dominating factor, and the platform is motivated to disclose inventory information earlier (i.e., lower the threshold for disclosure) to benefit from the customers' herding behaviors.
- The optimal time thresholds associated with different levels of inventory sold, t_j^*'s, are increasing in b, which corresponds to the fact that j^* decreases with b. Also, for every t_j^*, there is a "kink" point where it starts to increase dramatically to 1. If we draw a curve going through these kink points, the curve would be similar to that of j^*, showing the underlying relationship between time-dependent threshold policy and fixed threshold policy.

Having discussed properties of the optimal thresholds, we now focus on the behavior of R_3. Note that the fixed threshold policy can be viewed as a special case of the time-dependent threshold policy. That is, suppose j^* is the optimal fixed threshold such that the platform will disclose inventory information once the amount of product sold reaches j^*. Indeed, such a policy is equivalent to a time-dependent threshold policy with $t_j = 0$ for $j < j^*$ and $t_j = T$ for $j \geq j^*$.

Our explanation for the location of the peak of R_3 comes from the observation of the "stability" of j^* when b varies. If j^* increases or decreases rapidly when b varies, then it means that the j^* at this region is not "stable" enough, and it would be much better if we can find a policy which can tell us in what circumstances would, for example, $j^* - 1$, or j^*, or $j^* + 1$, be the best threshold. This is exactly where the time-dependent threshold policy can help. In fact, when j^* is not stable, we can also observe densely distributed kink points near the value of b for the time-dependent threshold policy, which indicates that the optimal time-dependent thresholds also varies dramatically. Thus, we can conclude that the most conspicuous improvement of time-dependent threshold policy occurs when j^* changes the fastest with b. This explanation can also help us illustrate the behavior of R_3 when other parameters varies in the following analysis.

Impact of ρ_{max}. Figure C.3 depicts the measures of improvements when ρ_{max} changes. We first focus on R_1, which increases with ρ_{max}. The increasing pattern is due to the fact that for larger values of ρ_{max}, the learning ability of customers is higher, and thus the "always disclose" policy benefits more from the higher learned product quality. Also, ρ_{max} does not have any impact on the demand rate in the "never disclose" case.

Figure C.3 R_1 **and** R_3 **as** ρ_{max} **varies.**

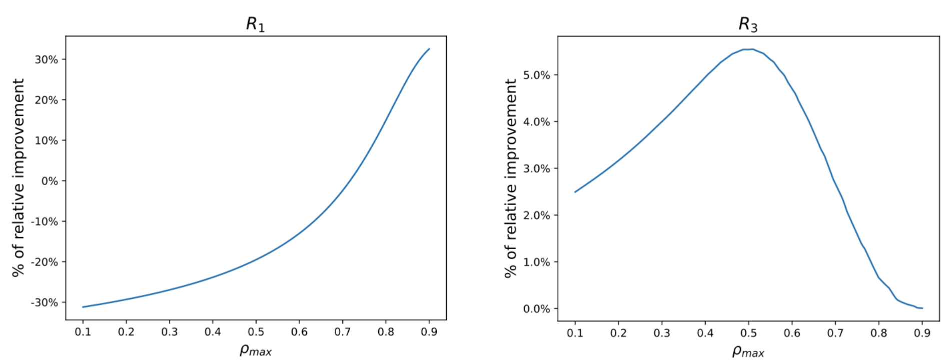

Note. $b = 1.2$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

To investigate the behavior of R_3, we again resort to the optimal fixed threshold and time-dependent thresholds. Our analysis shows that j^* decreases with ρ_{max} and t_j^*'s increases with ρ_{max}. R_3 peaks at near 0.5, which is where j^* and t_j^*'s varies most rapidly, which is consistent with our previous logic.

Impact of q_{min}. Figure C.4 illustrates the impact of q_{min} on the four measures of improvements. We find that R_1 is increasing in q_{min} while R_2, R_3, and R_4 are decreasing in q_{min}. Increasing q_{min} reduces the heterogeneity of the customers' prior beliefs of product quality and undermines the importance of observational learning. It also increases the demand rates under all policies and shrinks the gaps of performance between different policies. Hence, the disadvantage of "always disclose" policy compared with "never disclose" policy diminishes, and the fixed and time-dependent threshold policies can improve less when q_{min} increases.

Figure C.4 Measures of improvements as q_{min} varies.

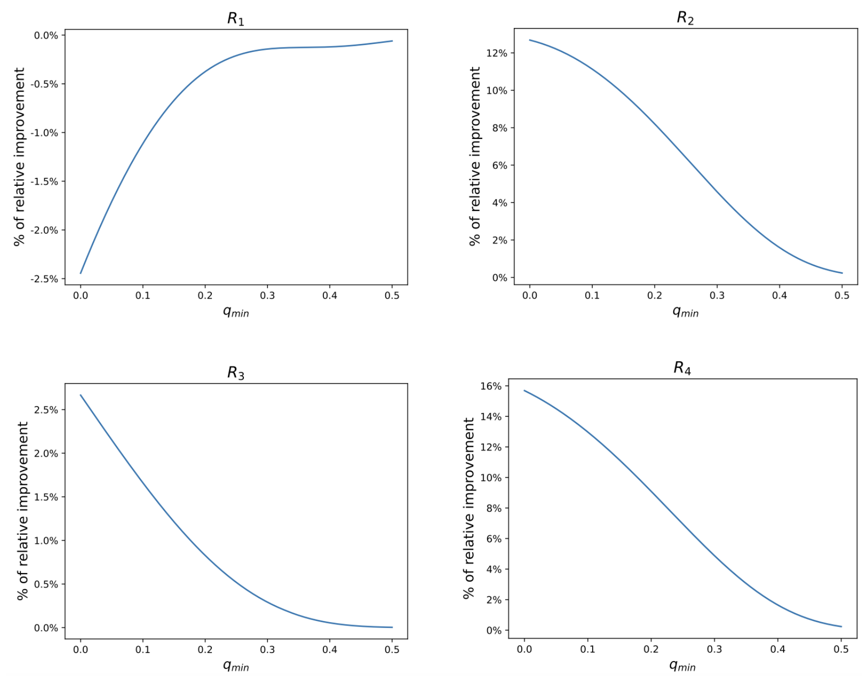

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

Impact of $E\eta$. Figure C.5 displays the improvements of different policies when $E\eta$ varies. First, let's consider R_1, the improvement of "always disclose" compared to "never disclose." When $E\eta$ is small, as our previous analysis shows, the scarcity effect of the "never disclose" case is very weak compared to that of the "always disclose" case. As $E\eta$ increases, the scarcity effect of the "never disclose" case rises and even surpasses that of the "always disclose" case when the selling speed is not high. Thus, the relative performance of the "always disclose" policy deteriorates and R_1 decreases with $E\eta$.

Nest, we investigate R_3. Figure C.5 shows that R_3 is first increasing and then decreasing in $E\eta$. Moreover, the optimal time thresholds associated with different levels of inventory sold are decreasing in $E\eta$. The decreasing trend of the time thresholds implies that disclosure of inventory information becomes less favorable, which is consistent with our reasoning above that a higher $E\eta$ induces a relatively more serious negative consequence of slow sales, thus

Figure C.5 R_1 **and** R_3 **as** $E\eta$ **varies.**

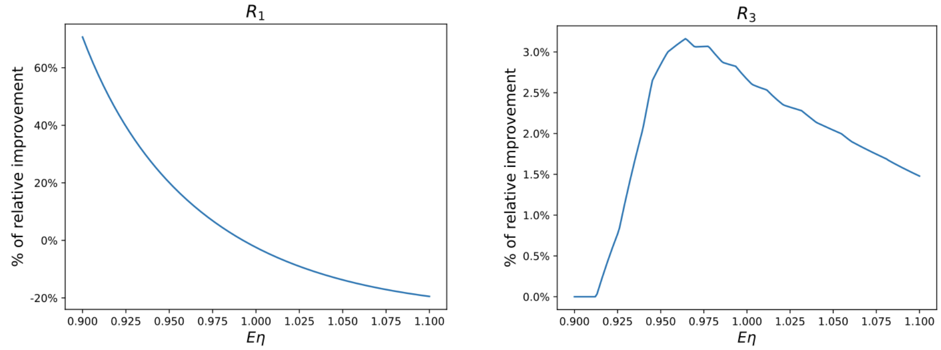

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\overline{\eta} - \underline{\eta} = 0.2$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

discouraging the platform from disclosing its inventory. The peak of R_3 is located at the place where j^* increases the fastest and the "kink" points of t_j^*'s are densely distributed.

Impact of δ_ε**.** Figure C.6 depicts the improvements of different policies when δ_ε changes. First, we observe that R_1 first increases then decreases with δ_ε. When δ_ε is large, the demand rate of the "always disclose" case is quite low, but the expected probability of stock-out is not affected by the low demand rate for the "never to disclose" case, so it is better to not reveal inventory information. As δ_ε shrinks, the demand rate of both simple policies rises. For the "always disclose" case, an additional benefit is that the scarcity effect becomes stronger and results in a higher R_1. Finally, as δ_ε continue to decrease, the product uncertainty becomes small and the demand rates of both simple policies soar. As a result, the sale is sold out irrespectively of the policy used, so R_1 becomes smaller again.

Figure C.6 R_1 **and** R_3 **as** δ_ε **varies.**

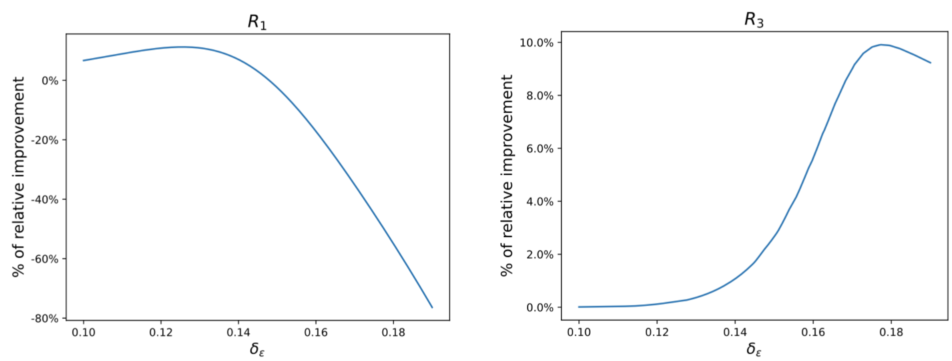

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

Next, note that R_3 exhibits a unimodal trend as δ_ε increases. The reasoning behind this observation is the same as previous ones. Actually, as δ_ε increases, j^* increases while t_j^*'s decrease, implying that an increase in δ_ε discourages the platform to broadcast inventory information under both threshold policies. Finally, the trend of R_4 is due to a combined effect of R_2 and R_3.

Impact of $\Delta\eta$**.** Figure C.7 depicts the influence of the range of customer's expected selling speed moderator, $\Delta\eta = \overline{\eta} - \underline{\eta}$, on the various measures of improvements and associated thresholds. As we conduct the following numerical

Figure C.7 Measures of improvements as $\Delta\eta$ varies.

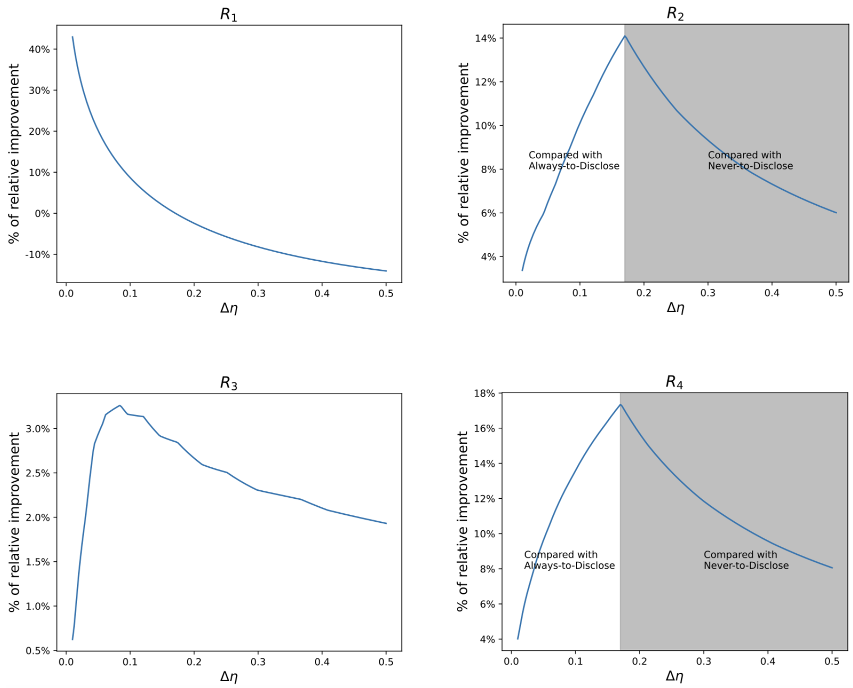

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $E\eta = 1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

experiment, we fix the average expected selling speed moderator $E\eta = (\overline{\eta} + \underline{\eta})/2$ at 1. The result shows that R_1 is decreasing in $\Delta\eta$, while R_2, R_3, and R_4 are all increasing then decreasing in $\Delta\eta$. When $\Delta\eta$ increases, it seems that the "never disclose" policy benefits more from it, as without any information of inventory, a higher $\Delta\eta$ creates better chance of the product being sold-out, thus stimulating the demand. On the other hand, the "always disclose" policy is not affected by $\Delta\eta$ as much as "never disclose" policy, as customers have information of inventory to help them estimate the probability of stock-out. So for the two simple policies, the "always disclose" policy is a better policy when $\Delta\eta$ is small and vice versa. The peak of R_2 is where two simple policies have similar performance. As for the unimodal pattern of R_3, the reasoning resembles the analysis of R_3's above. Lastly, R_4 integrates the impact in R_2 and R_3.

Analysis on the impact of difference factors on R_3. To better understand the situations where R_3 is large and where R_3 is 0, we perform the following analysis on our core testbed. The result is in Table C.1.

To explore the case when R_3 is large, we find all cases where $R_3 > 75\%$ of all positive R_3's. In these cases, their R_3 is relative large. We calculate the mean of different factors for these cases. The result shows that the average b is small to medium, average ρ_{max} is small, average $E\eta$ is medium, and the average δ_ε is large. This corresponds to the cases where the observational learning effect is small, the scarcity effect is medium, and the product uncertainty is high. In these cases, the scarcity effect dominates the observational learning effect, and together with a high product uncertainty, the randomness in the process of the flash sale is large. In these situations, the fixed threshold policy becomes less flexible with only one fixed threshold, and the relative improvement of the time-dependent threshold policy is large.

Figure C.8 Optimal price c^*'s as δ_ε varies.

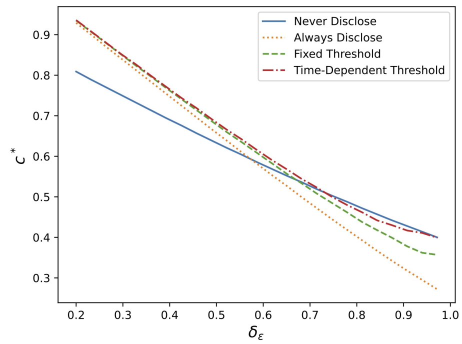

Note. $b = 1.5$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

For the cases where $R_3 = 0$, the mean of different factors shows that the average b is small to medium, average ρ_{max} is medium to large, average $E\eta$ is small, and the average δ_ε is small. This corresponds to the cases where the observational learning effect is medium to large, the scarcity effect is small, and the product uncertainty is low. In these cases, the observational learning effect dominates the scarcity effect, and together with a low product uncertainty, the randomness in the process of the flash sale is small. In these situations, the requirement on the flexibility of the model is relatively lower, and the fixed threshold policy can achieve the optimal expected sales, which is exactly the same with the time-dependent threshold policy.

	b	ρ_{max}	$E\eta$	δ_ε
Cases where $R_3 > 75\%$ of all positive R_3's	1.234	0.647	1.007	0.189
Cases where $R3 = 0$	1.200	0.787	0.955	0.121

Table C.1 Mean of different factors when the value of R_3 is in different levels

Optimal price c^*'s behavior when pricing decision is considered. The parameters most related with price c are b and δ_ε. As the retailer is optimizing over c and the utility of the customers is $bq_x(Ks) - c + \varepsilon_x$, we only need to fix one of b and δ_ε and let the other one vary, as fixing the other one would lead to exactly the same results. We choose to fix b and let δ_ε vary. For each value of δ_ε, we find the optimal price c^* numerically. To be specific, we search the optimal price in the interval $[0, b]$ (to ensure that the expected utility of buying the product immediately is positive for some customers so that the retailer has a positive revenue), and we use a step size of 0.001. For every c, we calculate the expected revenue, and we finally find the optimal c^* which corresponds to the highest revenue. Figure C.8 illustrates how optimal price c^*'s under different inventory disclosure policies change with δ_ε. We find that the optimal price decreases with δ_ε. The reason is that customers' uncertainty over product quality increases with δ_ε and results in lower expected sales when price is fixed, thus encouraging retailers to lower price. Also, The optimal price of the time-dependent threshold policy is always higher than that of the fixed threshold policy, which supports our results in Section 5.3. Another observation here is that the optimal prices of two threshold policies are higher than that of the "never disclose" policy when δ_ε is small, while the opposite is true when δ_ε is large. The high optimal price of two threshold policies when δ_ε is low can be attributed to the fact that their performances are much better than that of the "never disclose" policy so that they can raise price. The reason of the phenomenon when δ_ε is high is that the demand rates of threshold policies are affected more by the increasing δ_ε. For two threshold policies, the probability of stockout $p(s,\tau)$ depends on the real-time s, while for the "never disclose" policy the estimated probability of stockout $p(\widetilde{s},\tau)$ is not affected by the real s. Therefore, when δ_ε is high, the uncertainty is high and the impact of $p(s,\tau)$ is quite large. As a result, if price is too high, the demand rates of two threshold policies will drop significantly. Hence, their optimal prices at this case would become lower than that of the "never disclose" policy, whose demand rates are affected less due to the estimated and more stable $p(\widetilde{s},\tau)$. We can also observe the same pattern in Figure 5 in Section 5.3, where the revenue of the "never disclose" policy is the least sensitive with respect to price c. Figure C.9 shows the relative improvements of the fixed threshold policy and the time-dependent threshold policy when pricing decision is made. The results is similar to the main result regarding δ_ε in Section 5.2. This means that after taking the pricing decision of the retailer into consideration, the two threshold policies can still make significant improvements under proper conditions.

Figure C.9 Measures of Improvements as δ_ε varies when pricing decision is made.

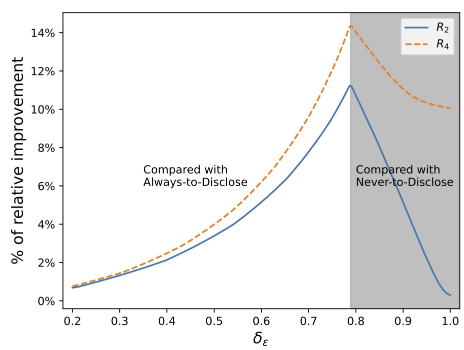

Note. $b = 1.5$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 200$, $a_0 = 0$, $K = 20$, and $T = 1$.

Analysis of the extended optimal policy. The extended optimal policy when the retailer can flip-flop the inventory disclosure decision is quite straightforward. As the seller can easily flip-flop the inventory disclosure decision, at any point in time t, suppose the current sales is j, then the seller can compare the demand rate when inventory is hidden and disclosed, and decide to disclose or not. In other words, under the extended optimal policy, the demand rate is:

$$\psi_j(t) = \max\{\theta(t), \lambda_j(t)\}$$

Then, similar to the cases of fixed threshold policy and the time-dependent threshold policy, we can derive the following differential equation:

$$W(T|t,j) = (1-\psi_j(t)\Delta t)W(T|t+\Delta t, j) + \psi_j(t)\Delta(t)(W(T|t+\Delta t, j+1)+1) + o(\Delta t)$$

where $W(T|t,j)$ denotes the expected "sales to go" during $[t,T]$ under the true optimal policy, given that there have been j sales at time t. We can then recursively solve for all $W(T|t,j)$'s as follows:

$$W(T|t,K) = 0$$
$$W(T|t,j) = e^{-\int_t^T \psi_j(s)ds} \int_t^T e^{\int_\tau^T \psi_j(s)ds} \psi_j(\tau)(W(T|\tau, j+1)+1)d\tau, j = 0,1,...,K-1$$

The expected sales under the extended optimal policy is actually $W(T|0,0)$. Thus, we can use numerical study to find out its improvement over other policies. We denote the improvement from the time-dependent threshold policy to the extended optimal policy as R_5. Table C.2 shows the result on our core testbed.

	Mean	Standard Deviation	25% quantile	50% quantile	75% quantile	Maximum
Cases with $R_2 > 0$	5.11%	6.35%	0.60%	2.95%	7.16%	45.21%
Cases with $R_3 > 0$	2.37%	2.65%	0.36%	1.57%	3.48%	24.17%
Cases with $R_4 > 0$	7.29%	6.97%	2.37%	5.57%	9.94%	50.97%
Cases with $R_5 > 0$	2.09%	2.19%	0.27%	1.35%	3.31%	11.99%

Table C.2 Summary statistics of the relative improvements measured by R_2, R_3, R_4, and R_5

Among the total 84,375 cases, the extended optimal policy can achieve better expected sales ($R_5 > 0$) in 69,414 cases, which is a number similar to that of R_3. However, we can observe that the mean and all percentiles of R_5 is smaller than those of R_3. This means that the further improvement from the time-dependent threshold policy to the extended optimal policy is smaller than, though comparable to, that from the fixed threshold policy to the time-dependent threshold policy.

Impact of a_0. Figure C.10 illustrates the impact of a_0, which actually indicates the monotonicity of customers' average arrival rate. As a_0 greatly affects the demand rate when no inventory information is available, we set $\Theta(0,T) = 0.75 \times K$ instead of $\Lambda_0 = 200$ to make the comparison of policies meaningful. R_1 is decreasing with a_0, showing that never-to-dicslose policy is more favorable when a_0 is larger. This can be explained as follows. When a_0 is negative and its absolute value is large, it corresponds to the case where customers rush to browse the products at very early time and the arrival rate gradually drops with time. In this case, the difference between different demand rates is large at early sale time but small later. Thus, retailers focus more on early stages and fear less that bad performance in early stage will harm the later stage of the sale, making the always-to-disclose policy better than never-to-disclose policy. As a_0 increases, the importance of late stage increases, thus making retailers more cautious when making the inventory disclosure decision. When a_0 is positive and large, retailers prefer to be conservative and use the never-to-disclose policy. This also leads to the unimodal behavior of R_2. R_3 reaches its peak when a_0 is negative, and the reasoning resembles the analysis of R_3's above. Finally, R_4 is an integrated effect of R_2 and R_3.

Impact of Risk Aversion. In the presence of risk aversion, the customer incurs disutility that increases with her perceived uncertainty. In our extended model, we use the variance of posterior belief, $\sigma_x^2(Ks)$, to represent the perceived uncertainty, and the utility of a customer with risk aversion can be written as:

$$U_x(s) = b(q_x(Ks) - \beta\sigma_x^2(Ks)) - c + \varepsilon_x \quad \text{(C.1)}$$

Figure C.10   Measures of improvements as a_0 varies.

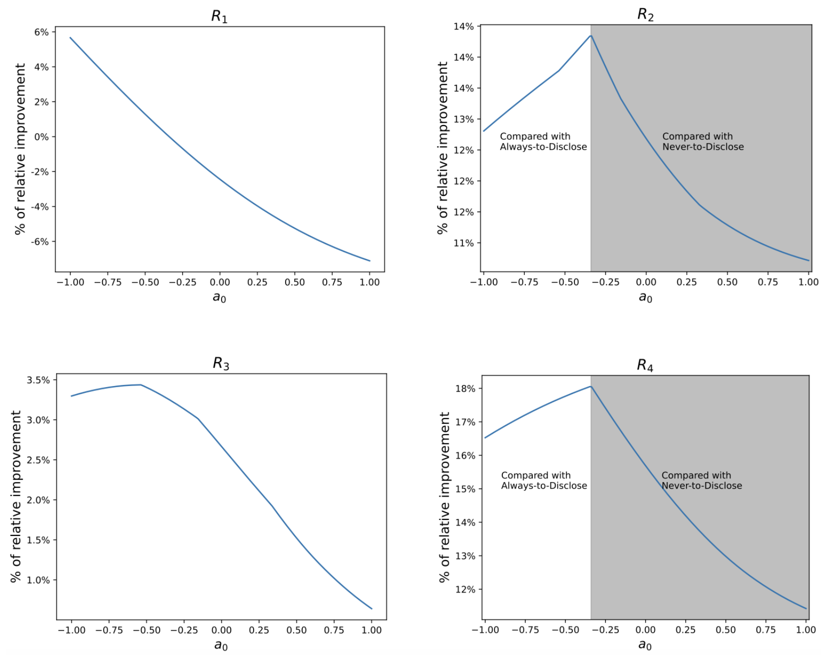

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Theta(0,T) = 0.75 \times K$, $K = 20$, and $T = 1$.

Here β represents the extent of risk aversion. $\sigma_x^2(Ks)$ is updated together with the mean of posterior belief $q_x(Ks)$ under our Bayesian updating framework, and we can calculate it as:

$$\sigma_x^2(Ks) = \frac{1}{\frac{1}{\sigma_0^2} + \frac{Ks}{\sigma_{TQ}^2}} = (1-\rho(s))\sigma_0^2 \tag{C.2}$$

Remember that $\rho(s)$ increases with s, we can find that the variance on product quality decreases with s, which means that the perceived risk of buying the product decreases with the number of units sold. Notice that, in the case of risk aversion, even a customer with a high prior belief on the product quality (i.e., $q_{x,0} > q_{TQ}$) may still not buy the product due to its high quality uncertainty, and her belief evolves with s and t in a similar way to those with a low prior belief. However, if the customer holds an extremely high prior belief (i.e., $q_{x,0} - \beta\sigma_0^2 > q_{TQ}$), we assume that even quality uncertainty would not stop her from buying the product immediately, which is reasonable. Plugging (C.2) in (C.1), and following the derivation process in Section 3, we can get the demand rates $\theta(t)$ and $\lambda_i(t)$:

$$\theta(t) = \Lambda(t)\frac{q_{max} - \beta\sigma_0^2 - \frac{1}{b}\left(c + \delta_\varepsilon \frac{1-\sqrt{p(\widetilde{s},\tau)}}{1+\sqrt{p(\widetilde{s},\tau)}}\right)}{q_{max} - q_{min}}.$$

$$\lambda_i(t) = \Lambda(t)\frac{q_{max} - \beta\sigma_0^2 - q_{TQ} + \frac{q_{TQ} - \frac{1}{b}\left(c+\delta_\varepsilon \frac{1-\sqrt{p(s,\tau)}}{1+\sqrt{p(s,\tau)}}\right)}{1-\rho(s)}}{q_{max} - q_{min}}.$$

Comparing the above expression of demand rates with those in the risk-neutral case, we can find that we actually substitute $q_{max} - \beta\sigma_0^2$ for q_{max}. Thus, the properties of the demand rate functions does not change. So under Assumptions 1 and 2, the analytical results in Section 4 still hold, and the time-dependent threshold policy is still the optimal policy.

Figure C.11 **Measures of improvements as β varies.**

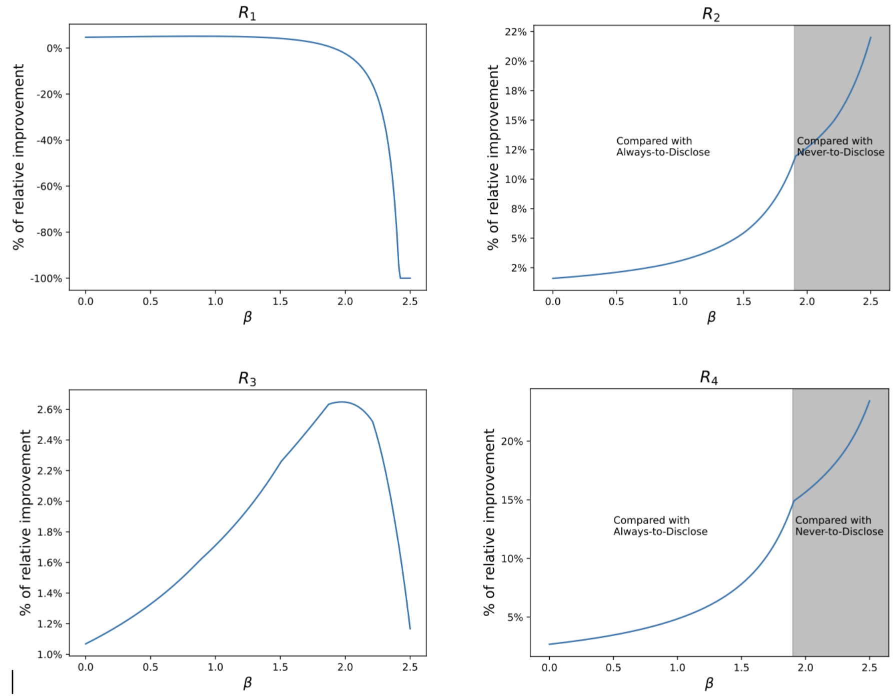

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Theta(0,T) = 0.75 \times K$, $\sigma_0^2 = 0.1$, $a_0 = 0$, $K = 20$, and $T = 1$.

We then leverage numerical studies to find how exactly risk aversion affect the performance of different policies. We first investigate how the degree of risk aversion, β, affects the performance of various policies. We normalize σ_0^2 to 0.1, and we assume that $\Theta(0,T) = 0.75 \times K$ to make situations under different β's comparable, as β has a great impact on demand rates if $\Lambda(t)$ is kept constant. Figure C.11 illustrates the impact of β. As β increases, customers are less likely to buy the product due to increasing concerns on quality uncertainty. Also, more customers with high prior belief hesitate to buy the product at first sight due to the increasing impact of quality uncertainty. At the same time, herding effect becomes more important, as more units sold leads to not only a higher perceived quality, but also a smaller variance. As a result, when β increases, the utility of all customers decrease, and fewer customers with high prior belief would buy immediately. Under the always-to-disclose policy, units sold would increase more slowly with time, which contributes to the decreasing trend of R_1. R_2 is strictly increasing, which can be explained as follows. As β increases, more customers with high prior belief would hesitate to purchase immediately and turn to rely on the herding effect to make decisions, and a flexible threshold makes the fixed threshold policy taking much more advantage from the herding effect compared with the two simple policies. The behavior of R_3 can be explained similar to the analysis of R_3's above. Finally, R_4 is an integrated effect of R_2 and R_3.

Next, we examine if our main results still hold when customers are risk averse. We fix $\beta = 1.5$, and Figures C.12-C.15 show the corresponding curves. We find that the curves do not change much, which means that the results in the risk-averse case resembles those in the risk-neutral case.

Appendix D: Algorithm Pseudocodes and Complexity Analysis

The pseudocode of the optimal fixed threshold policy algorithm in the paper is presented in Algorithm 1.

Figure C.12 Measures of improvements as b varies when customers are risk averse.

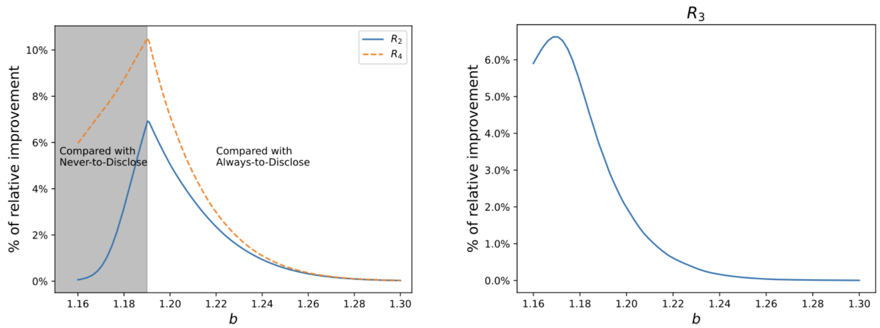

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 150$, $\beta = 1.5$, $\sigma_0^2 = 0.1$, $a_0 = 0$, $K = 20$, and $T = 1$.

Figure C.13 Measures of improvements as ρ_{max} varies when customers are risk averse.

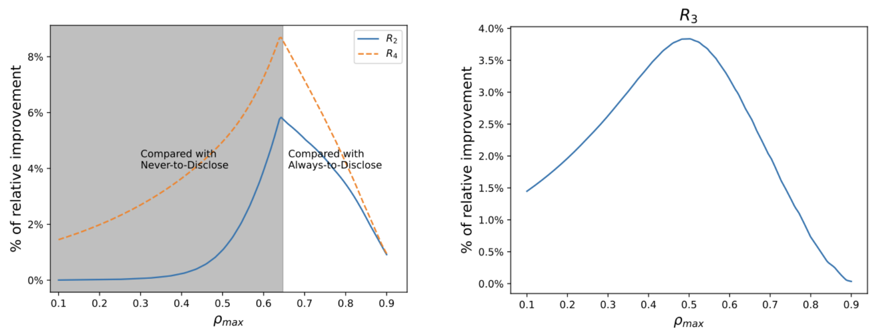

Note. $b = 1.2$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 150$, $\beta = 1.5$, $\sigma_0^2 = 0.1$, $a_0 = 0$, $K = 20$, and $T = 1$.

Figure C.14 Measures of improvements as $E\eta$ varies when customers are risk averse.

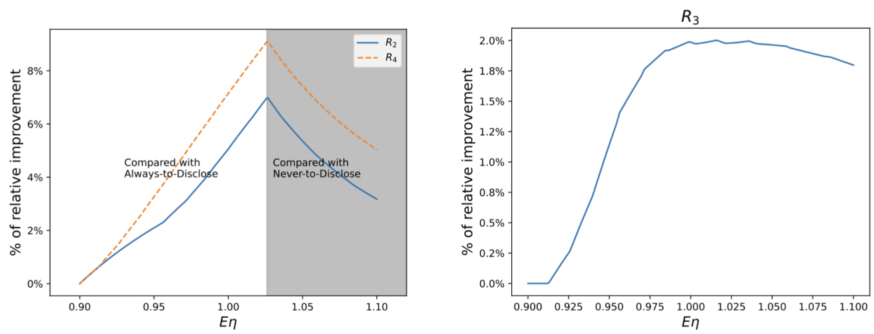

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\Delta\eta = 0.2$, $\Lambda_0 = 150$, $\beta = 1.5$, $\sigma_0^2 = 0.1$, $a_0 = 0$, $K = 20$, and $T = 1$.

Actually, the algorithm can be improved. We do not need to calculate every state probability Q. Instead, we can directly derive differential equations of Γ's with similar procedures. The differential equations are as follows:

$$\frac{d\Gamma(T|t,j)}{dt} = \lambda_j(t)\Gamma(T|t,j) - \lambda_j(t)(\Gamma(T|t,j+1)+1)$$

We can then derive Algorithm 2 based on the differential equations.

Figure C.15 **Measures of improvements as δ_ε varies when customers are risk averse.**

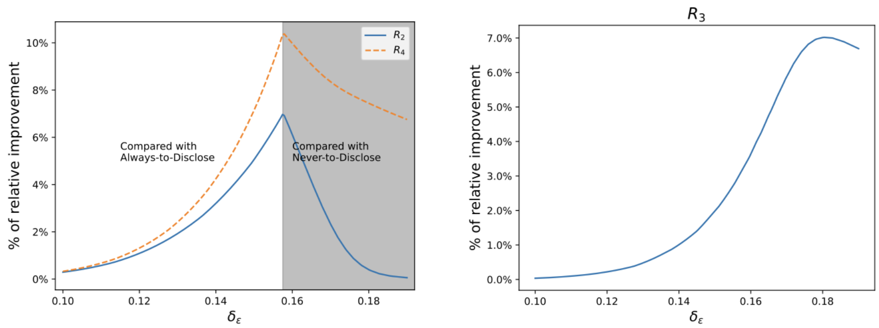

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $\Lambda_0 = 150$, $\beta = 1.5$, $\sigma_0^2 = 0.1$, $a_0 = 0$, $K = 20$, and $T = 1$.

Algorithm 1 Fixed Threshold Policy

Calculate $\Theta(0,t) = \int_0^t \theta(s)ds$

for $j = K-1, K-2, \ldots, 0$ **do**

  Calculate $\Lambda_j(0,t) = \int_0^t \lambda_j(s)ds$

  Calculate $Q_0(t|\zeta,j) = e^{-(\Lambda_j(0,t)-\Lambda_j(0,\zeta))}$

  for $i = 1, 2, \ldots, K-j$ **do**

    Calculate $Q_i(t|\zeta,j) = e^{-(\Lambda_{j+i}(0,t)-\Lambda_{j+i}(0,\zeta))} \left[\int_\zeta^t e^{(\Lambda_{j+i}(0,s)-\Lambda_{j+i}(0,\zeta))} \lambda_{j+i-1}(s) Q_{i-1}(s|\zeta,j)ds\right]$

  Calculate $\Gamma(T|\zeta,j) = \sum_{i=0}^{K-j} i \cdot Q_i(T|\zeta,j)$

Calculate $S^*(0) = \Gamma(T|0,0)$

for $j = 1, 2, ..., K-1$ **do**

  Calculate $S^*(j) = \int_0^T \Gamma(T|\zeta,j)\theta(\zeta)p(j-1,\Theta(0,\zeta))d\zeta + \sum_{x=0}^{j-1} x \cdot p(x,\Theta(0,T)) + j \cdot P(j,\Theta(0,T))$

  Calculate $\Delta S^*(j) = S^*(j) - S^*(j-1)$

  if $\Delta S^*(j) < 0$ **then**

    The optimal $j^* = j-1$

    return j^*

Calculate $S^*(K) = \sum_{x=0}^{K-1} x \cdot p(x,\Theta(0,T)) + K \cdot P(K,\Theta(0,T))$

Calculate $\Delta S^*(K) = S^*(K) - S^*(K-1)$

if $\Delta S^*(K) < 0$ **then**

  The optimal $j^* = K-1$

else

  The optimal $j^* = K$

return j^*

Algorithm 2 Fixed Threshold Policy (Efficient Version)

Calculate $\Theta(0,t)=\int_0^t \theta(s)ds$
$\Gamma(T|t,K)=0$
Calculate $S^*(K)=\sum_{x=0}^{K-1} x\cdot p(x,\Theta(0,T))+K\cdot P(K,\Theta(0,T))$
for $j=K-1,K-2,\ldots,0$ **do**
  Calculate $\Lambda_j(0,t)=\int_0^t \lambda_j(s)ds$
  Calculate $\Gamma(T|\zeta,j)=e^{-(\Lambda_j(0,T)-\Lambda_j(0,\zeta))}\left[\int_\zeta^T e^{(\Lambda_j(0,T)-\Lambda_j(0,s))}\lambda_j(s)(\Gamma(T|s,j+1)+1)ds\right]$
  if $j>0$ **then**
    Calculate $S^*(j)=\int_0^T \Gamma(T|\zeta,j)\theta(\zeta)p(j-1,\Theta(0,\zeta))d\zeta+\sum_{x=0}^{j-1} x\cdot p(x,\Theta(0,T))+j\cdot P(j,\Theta(0,T))$
  else
    Calculate $S^*(0)=\Gamma(T|0,0)$
  Calculate $\Delta S^*(j+1)=S^*(j+1)-S^*(j)$
  if $\Delta S^*(j+1)>0$ **then**
    The optimal $j^*=j+1$
    return j^*
The optimal $j^*=0$
return j^*

The complexity of Algorithm 2 can be analyzed as follows. First, to calculate $\Theta(0,t)=\int_0^t \theta(s)ds$, we need to simulate N_t time points, so the time complexity is $O(N_t)$. Next, we analyze what happens at every j. Calculating $\Lambda_j(0,t)=\int_0^t \lambda_j(s)ds$ is a cost of $O(N_t)$. To calculate $\Gamma(T|\zeta,j)$, we can do the numerical integrations incrementally as ζ gradually decreases from T to 0. So the cost is another $O(N_t)$. Then, when it comes to $S^*(j)$, we can first calculate $\sum_{x=0}^{j-1} x\cdot p(x,\Theta(0,T))$ and $P(j,\Theta(0,T))$ incrementally, which is a cost of $O(1)$ at each j. For every $S^*(j)$, the integration needs $O(N_t)$ time. So the time cost of all calculations at a j is $O(N_t)$. Considering that the search of j^* ends when $\Delta S^*(j)>0$ (we search backwards from K to 0), the expected number of j's to be searched is $O(K)$. So adding all steps up, the complexity of the efficient version of the algorithm is $O(N_tK)$, which indicates that the algorithm is very efficient. Note that even without the unimodality assumption of $S^*(j)$, we just need to calculate $S^*(j)$'s from K all the way back to 0, which means that the number of j's to be searched is still $O(K)$. Thus, the complexity of the efficient version of the algorithm is still $O(N_tK)$.

The pseudocode of the optimal time-dependent threshold policy algorithm in the paper is presented in Algorithm 3.

The efficiency of the time-dependent threshold policy can also be improved. Similar with the case of the Γ's in the fixed threshold policy, we can derive the differential equations as follows:

$$\frac{dV^*(T|t,j)}{dt}=\theta(t)V^*(T|t,j)-\theta(t)(S^*(T|t,j+1)+1)$$

Algorithm 3 Time-Dependent Threshold Policy

$t_K^* = T$

Calculate $\Theta(0,t) = \int_0^t \theta(s)ds$

for $j = K-1, K-2, \ldots, 0$ **do**

  for $i = 0, 1, \ldots, K-1$ **do**

    Calculate $p(i, \Theta(t_{j+1}^*, y))(t_{j+1}^* \le y \le t_{j+2}^*)$

  $r_0(t_{j+1}^*|t_{j+1}^*, j) = 1$

  for $k = 2, 3, \ldots, K-j$ **do**

    for $n = 0, 1, \ldots, k-1$ **do**

      $r_n(t_{j+k}^*|t_{j+1}^*, j) = \sum_{m=0}^{n \wedge (k-2)} r_m(t_{j+k-1}^*|t_{j+1}^*, j) p(n-m, \Theta(t_{j+k-1}^*, t_{j+k}^*))$

      $r_n(t_{j+k}^*|t, j) = p(0, \Theta(t, t_{j+1}^*)) r_n(t_{j+k}^*|t_{j+1}^*, j)$

  Calculate $q_1(y|t,j) = \theta(y)p(0, \Theta(t,y))$

  for $i = 2, 3, \ldots, K-j$ **do**

    $q_i(y|t,j) = \sum_{n=0}^{i-2} r_n(t_{j+i-1}^*|t,j)\theta(y)p(i-1-n, \Theta(t_{j+i-1}^*, y))$

  Calculate $V^*(T|t,j) = \int_t^{t_{j+1}^*}(1+\Gamma(T|y,j+1))q_1(y|t,j)dy + \sum_{i=2}^{K-j}\int_{t_{j+i-1}^*}^{t_{j+i}^*}(i+\Gamma(T|y,j+i))q_i(y|t,j)dy + \sum_{n=0}^{K-j-1} n \cdot r_n(T|t,j)$

  Find t_j^* as the largest t such that $\Gamma(T|t,j) \ge V^*(T|t,j)$. If no such t_j^* exists, define $t_j^* = t_{j-1}^* = \ldots = t^*_{)} = 0$ and jump out of the outer loop.

return $t_0^*, t_1^*, \ldots, t_K^*$

Here $S^*(T|t,j+1)$ denotes the optimal sales-to-go under the time-dependent threshold policy when $j+1$ units are sold. We need to find $V^*(T|t,j)$ from the above differential equation and find t_j^* which is the largest t such that $\Gamma(T|t,j) \ge V^*(T|t,j)$. Based on this, we can derive Algorithm 4.

We can analyze the complexity of Algorithm 4 as follows. First, to calculate $\Theta(0,t) = \int_0^t \theta(s)ds$, we need to simulate N_t time points, so the time complexity is $O(N_t)$. Next, we analyze what happens at every j. We calculate $V^*(T|t,j)$ for t from T to 0 until we find the largest t such that $\Gamma(T|t,j) \ge V^*(T|t,j)$ and define it as t_j^*. This is a cost of $O(N_t)$. Record values for $S^*(T|t,j)$ is another cost of $O(N_t)$. Thus, for every j, it takes $O(N_t)$ to perform all calculations. The complexity of the algorithm is $O(N_tK)$. Even if we add the time of calculating $\Gamma(T|t,j)$'s, which is a cost of another $O(N_tK)$, the complexity is still $O(N_tK)$. This shows that our time-dependent threshold policy algorithm is extremely efficient, reaching a complexity the same as the fixed threshold policy algorithm.

Appendix E: Robustness Checks

We first consider the probability distribution of $\widetilde{q}_x$ and $\widetilde{q}_{TQ}$. In our model, we assume that $\widetilde{q}_x$ and $\widetilde{q}_{TQ}$ both follow normal distribution. We make this assumption as the conjugate prior of normal distribution is normal distribution, which makes the process of Bayesian learning tractable. Here, we change the distribution of $\widetilde{q}_x$ to a Gamma distribution $Gamma(\alpha, \beta)$ and the distribution of $\widetilde{q}_{TQ}$ to a Poisson distribution $Poisson(q_{TQ})$. This means that the mean of

Algorithm 4 Time-Dependent Threshold Policy (Efficient Version)

1: $t_K^* = T$

2: $S^*(T|t,K) = 0$

3: Calculate $\Theta(0,t) = \int_0^t \theta(s)ds$

4: **for** $j = K-1, K-2, \ldots, 0$ **do**

5: Calculate $V^*(T|t,j) = e^{-\Theta(t,T)} \int_t^T e^{\Theta(s,T}\theta(s)(S^*(T|s,j+1)+1)ds$

6: Find t_j^* as the largest t such that $\Gamma(T|t,j) \geq V^*(T|t,j)$. If no such t_j^* exists, define $t_j^* = t_{j-1}^* = \ldots = t_)^* = 0$ and jump out of the loop.

7: $S^*(T|t,j) = \begin{cases} \Gamma(T|t,j), & 0 \leq t \leq t_j^* \\ V^*(T|t,j), & t_j^* < t \leq T \end{cases}$

8: **return** $t_0^*, t_1^*, \ldots, t_K^*$

customers' prior belief of the product quality is $q_{x,0} = \frac{\alpha}{\beta}$, and the mean of signals is q_{TQ}. As Gamma distribution is the conjugate prior of Poisson distribution, we can derive tractable results for customers' posterior belief. To be specific, after observing i units of sales, customers' posterior belief $\widetilde{q}_x(i)$ is updated to a Gamma distribution with mean:

$$q_x(i) = \frac{\alpha + iq_{TQ}}{\beta + i} = \frac{\beta}{\beta + n} \cdot \frac{\alpha}{\beta} + \frac{i}{\beta + i} q_{TQ} = (1-\rho(s))q_{x,0} + \rho(s)q_{TQ}$$

Here, we define $\rho(s) = \frac{i}{\beta+i} = s/\frac{\beta}{K} + s$. We can then find out that this updating rule is almost the same with the current updating rule assuming $\widetilde{q}_x$ and $\widetilde{q}_{TQ}$ both follow normal distribution. As using different distribution assumptions of parameters leads to similar results, we can conclude that our updating rule in the main model can represent customers' belief update in the real world.

We then discuss the distribution of the mean of customers' initial belief $q_{x,0}$. In our model, we assume that it follows a uniform distribution. Here, we assume that it is a unimodal distribution. We choose an adaptation of Beta distribution $Beta(2,2)$ as the distribution of $q_{x,0}$. $Beta(2,2)$ is defined on a closed interval $[0,1]$, so we let $\frac{q_{x,0}-q_{min}}{q_{max}-q_min}$ follow $Beta(2,2)$ to make sure that $q_{x,0} \in [q_{min}, q_{max}]$. We explore the impact of various factors in this case, and we find that the curves do not change much, and our main results still hold. The unimodal distribution of $q_{x,0}$ makes the customers near the median of the distribution very important, thus reshaping the curves. Figures E.1-E.4 shows these curves.

Next, we modify the distribution of utility uncertainty ε_x. Similar to the case of $q_{x,0}$, we change its distribution from a uniform distribution to a unimodal distribution which can be adapted from $Beta(2,2)$. Specifically, we let $\frac{\varepsilon_x}{2\delta_\varepsilon} + \frac{1}{2}$ follow the $Beta(2,2)$ distribution, which guarantees that $\varepsilon_x \in [-\delta_\varepsilon, \delta_\varepsilon]$. The curves of improvements are shaped by the new distribution of ε_x, and most curves are similar to those in our main model. However, for the cases of b and δ_ε, the distribution of ε_x affects them much and the curves are showing some different results, such as that now R_2 and R_4 do not peak at the point where two simple policies lead to the same expected sales. The reason is that when b is small or δ_ε is large, the uncertainty term ε_x in customers' utility becomes very important. Switching from a simple policy to the fixed threshold policy can lower the level of utility required for a customer to buy the product. As ε_x now follows a unimodal distribution, this makes the lowest level of mean utility required for customers to buy decrease much more than the case where ε_x is uniformly distributed, and as a result R_2 becomes higher than the case in our main model. The curve of R_4 is also shaped by similar reasons. Though some curves are significantly affected by the distribution

Figure E.1 **Measures of improvements as b varies when $q_{x,0}$ follows an adaptation of $Beta(2,2)$.**

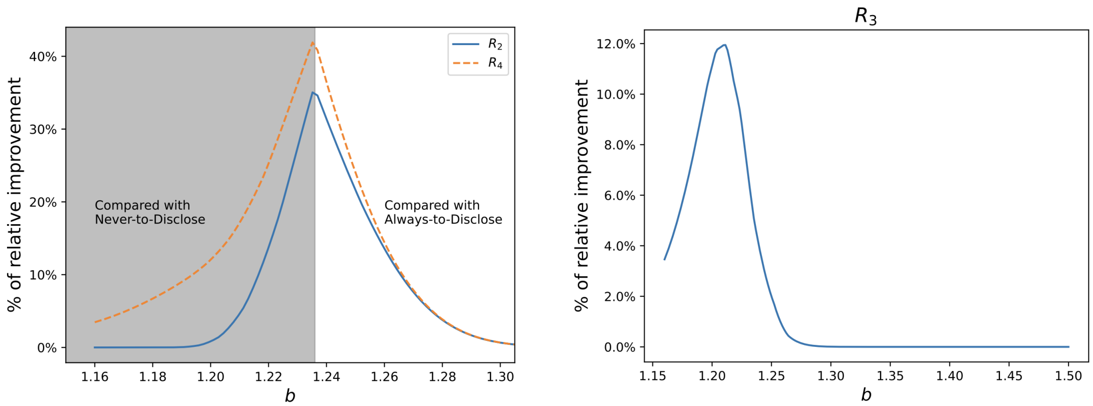

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.2 **Measures of improvements as ρ_{max} varies when $q_{x,0}$ follows an adaptation of $Beta(2,2)$.**

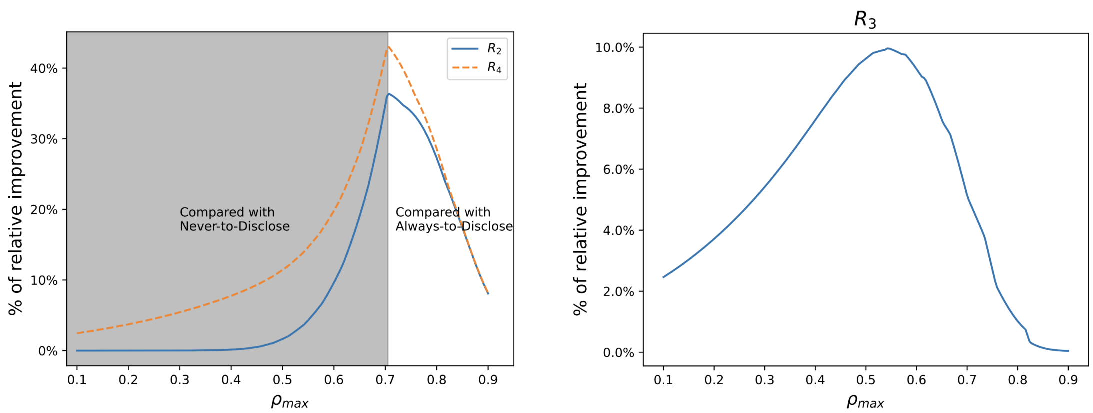

Note. $b = 1.235$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.3 **Measures of improvements as $E\eta$ varies when $q_{x,0}$ follows an adaptation of $Beta(2,2)$.**

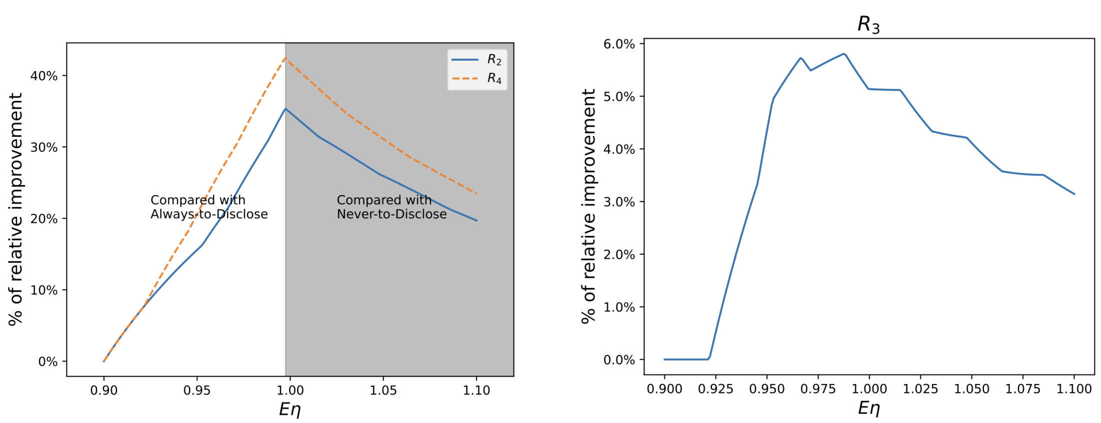

Note. $b = 1.235$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\Delta\eta = 0.2$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.4 **Measures of improvements as δ_ε varies when $q_{x,0}$ follows an adaptation of $Beta(2,2)$.**

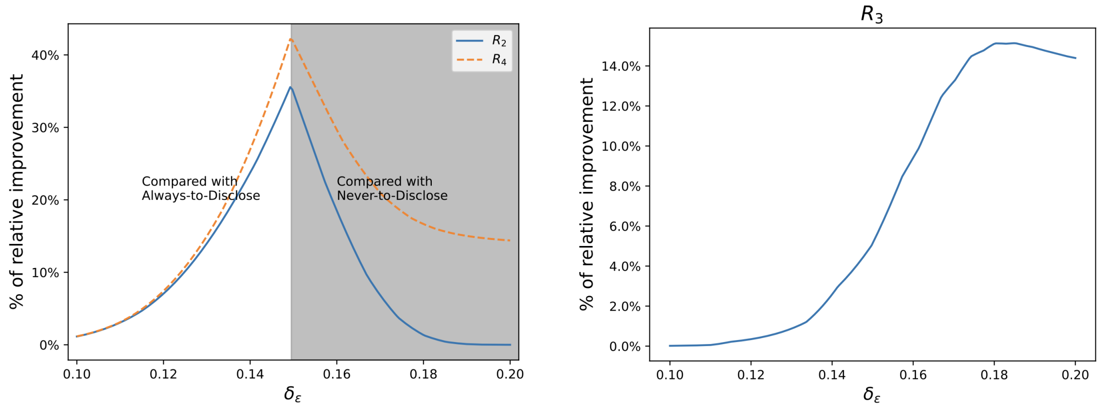

Note. $b = 1.235$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

of ε_x, our analysis in the main model still provides the basic shape of these curves. Figures E.5-E.8 show the curves which are shaped by the distribution of ε_x.

Figure E.5 **Measures of improvements as b varies when ε_x follows an adaptation of $Beta(2,2)$.**

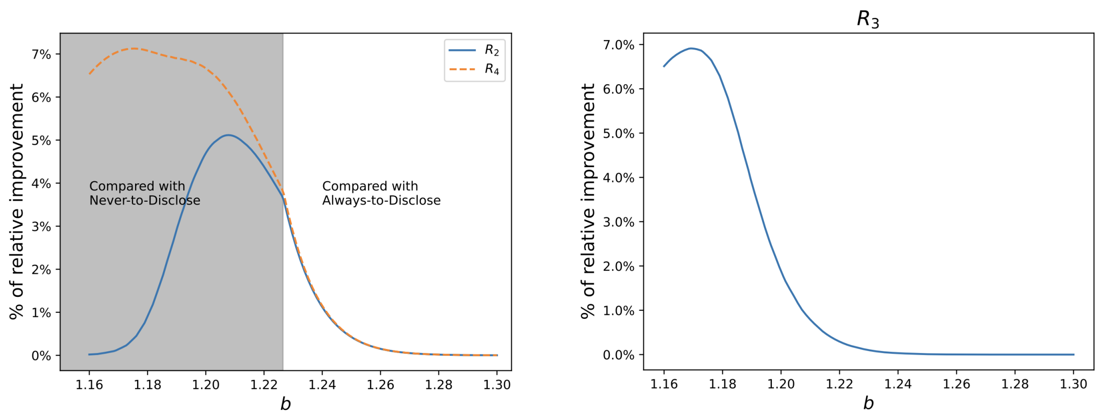

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Finally, we explore how the distribution of η affect our results. As η is uniformly distributed in a closed interval in our main model, we also change its distribution to an adaptation of a Beta distribution. To be specific, we let $\frac{\eta-\underline{\eta}}{\overline{\eta}-\underline{\eta}}$ follow the $Beta(2,2)$ distribution, which guarantees that $\eta \in [\underline{\eta}, \overline{\eta}]$. The curves of improvements are almost the same as in our main model, which indicates that the distribution of η has little impact on our main results. Figures E.9-E.12 show the curves which are shaped by the distribution of η.

Figure E.6 **Measures of improvements as ρ_{max} varies when ε_x follows an adaptation of $Beta(2,2)$.**

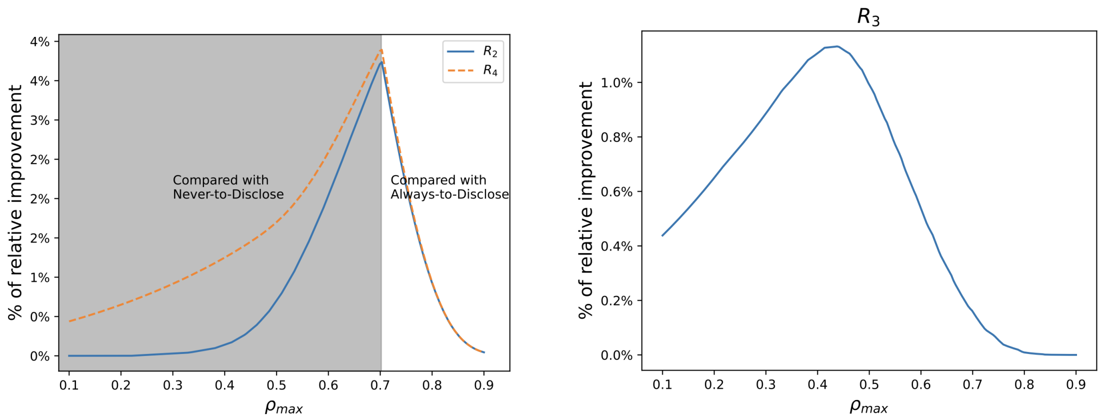

Note. $b = 1.226$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.7 **Measures of improvements as $E\eta$ varies when ε_x follows an adaptation of $Beta(2,2)$.**

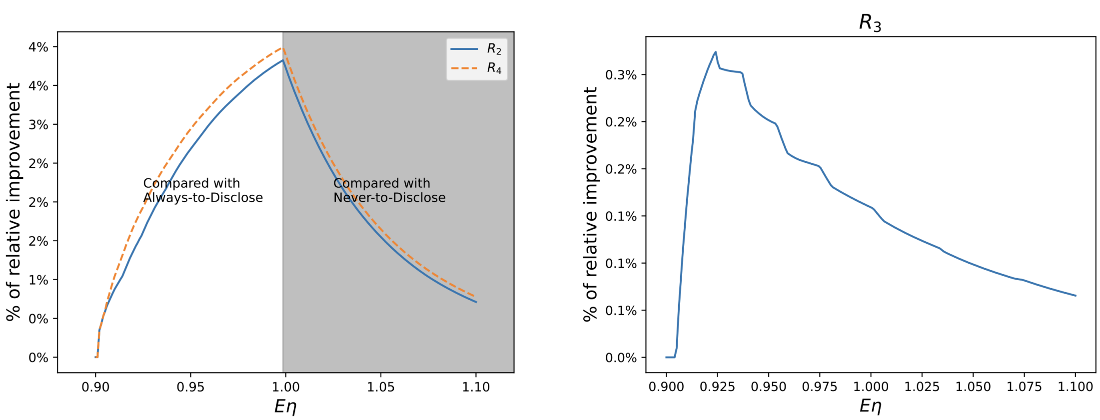

Note. $b = 1.226$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\Delta\eta = 0.2$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.8 **Measures of improvements as δ_ε varies when ε_x follows an adaptation of $Beta(2,2)$.**

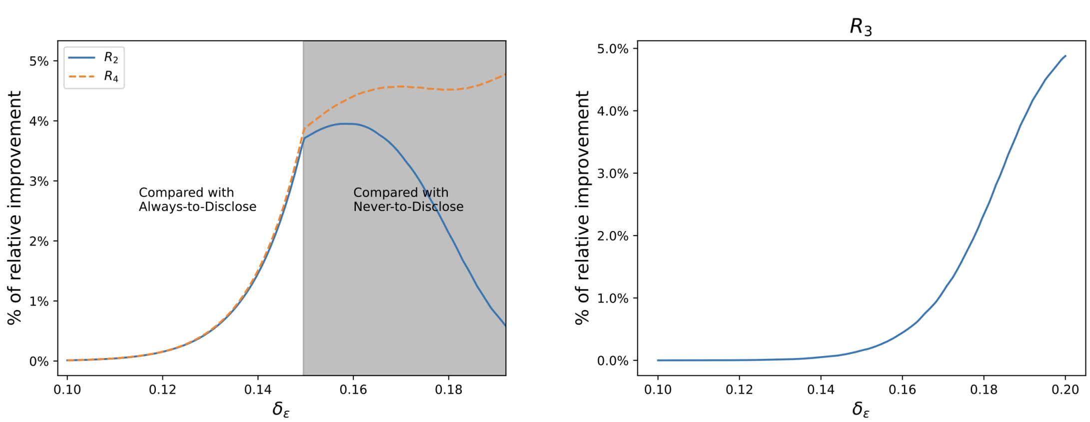

Note. $b = 1.226$, $\rho_{max} = 0.7$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.9 **Measures of improvements as b varies when η follows an adaptation of $Beta(2,2)$.**

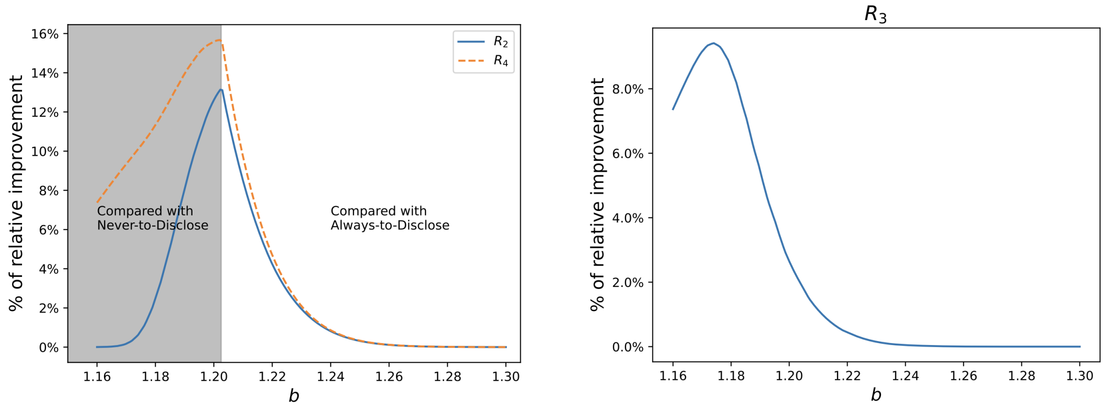

Note. $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.10 **Measures of improvements as ρ_{max} varies when η follows an adaptation of $Beta(2,2)$.**

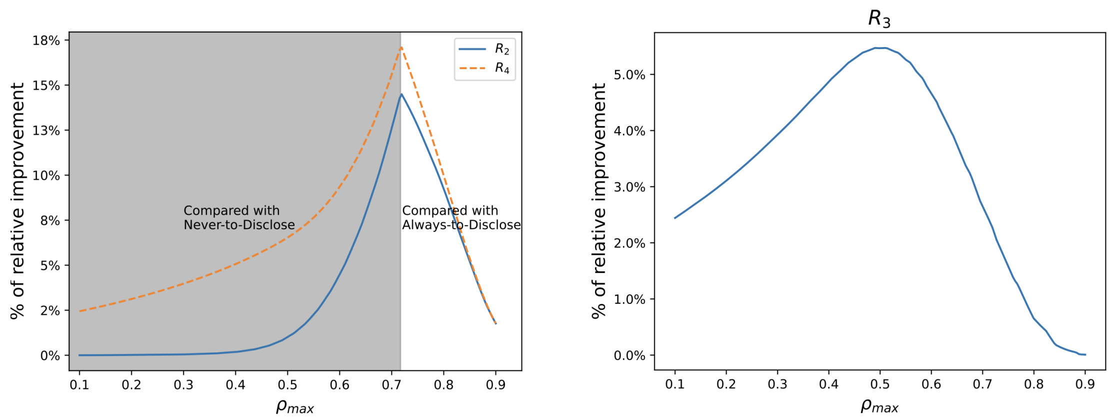

Note. $b = 1.2$, $\delta_\varepsilon = 0.15$, $\underline{\eta} = 0.9$, $\overline{\eta} = 1.1$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.11 **Measures of improvements as $E\eta$ varies when η follows an adaptation of $Beta(2,2)$.**

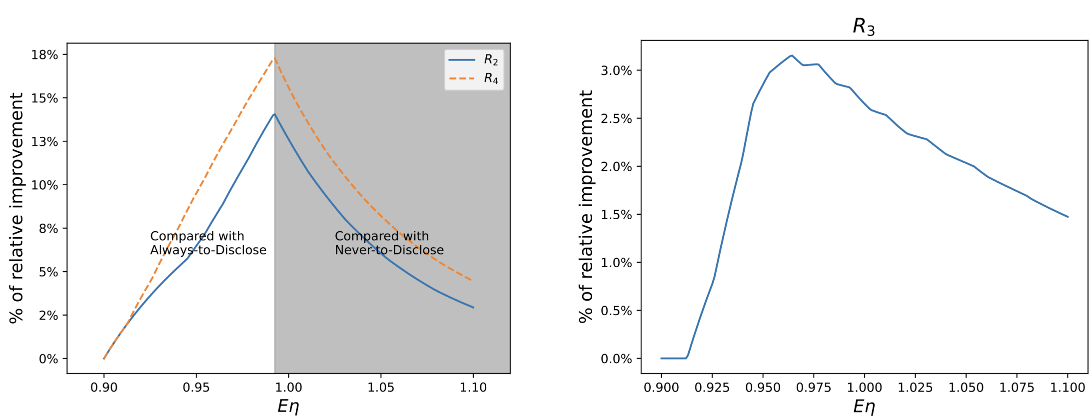

Note. $b = 1.2$, $\rho_{max} = 0.7$, $\delta_\varepsilon = 0.15$, $\Delta\eta = 0.2$, $a_0 = 0$, $\Lambda_0 = 200$, $K = 20$, and $T = 1$.

Figure E.12 **Measures of improvements as δ_ε varies when η follows an adaptation of $Beta(2,2)$.**

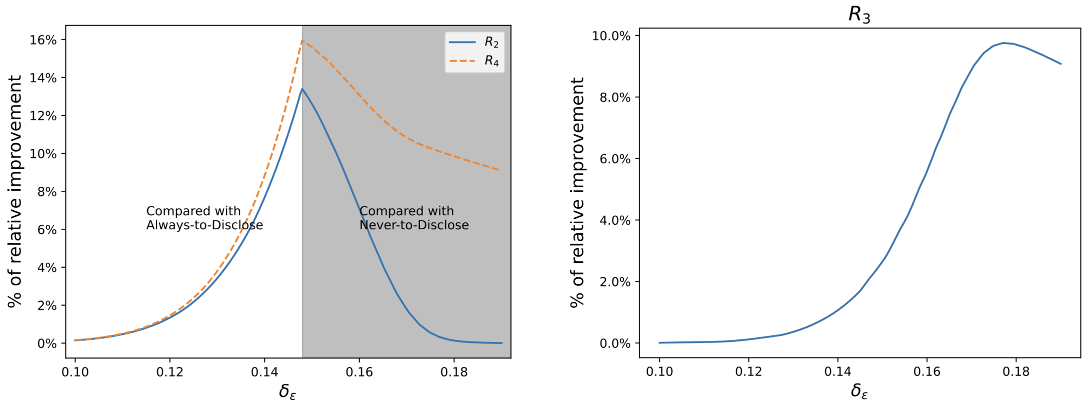

Note. $b=1.2$, $\rho_{max}=0.7$, $\underline{\eta}=0.9$, $\overline{\eta}=1.1$, $a_0=0$, $\Lambda_0=200$, $K=20$, and $T=1$.