

ΠΑΝΕΠΙΣΤΗΜΙΟ ΘΕΣΣΑΛΙΑΣ

ΠΟΛΥΤΕΧΝΙΚΗ ΣΧΟΛΗ

ΤΜΗΜΑ ΜΗΧΑΝΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ

Διδακτορική Διατριβή

**ΠΟΛΙΤΙΚΕΣ ΠΡΟΜΗΘΕΥΤΩΝ ΥΠΟ ΕΞΑΡΤΩΜΕΝΗ ΑΠΟ ΤΗΝ
ΕΞΥΠΗΡΕΤΗΣΗ ΖΗΤΗΣΗ ΑΓΟΡΑΣΤΩΝ: ΠΑΡΑΓΓΕΛΙΟΔΟΣΙΑ,
ΑΝΤΑΓΩΝΙΣΜΟΣ ΚΑΙ ΕΠΙΛΟΓΗ ΑΓΟΡΑΣΤΩΝ**

υπό

ΜΙΧΑΗΛ ΔΕΛΗΓΙΑΝΝΗ

Διπλωματούχου Μεταπτυχιακών Σπουδών στα Μαθηματικά στην Ειδικευση Στατιστική και
Μοντελοποίηση του Αριστοτελείου Πανεπιστημίου Θεσσαλονίκης, 2016

Πτυχιούχου Μαθηματικού
του Αριστοτελείου Πανεπιστημίου Θεσσαλονίκης, 2013

Υπεβλήθη για την εκπλήρωση μέρους των
Απαιτήσεων για την απόκτηση του
Διδακτορικού Διπλώματος

2022

© 2022 Μιχαήλ Δεληγιάννης

Η έγκριση της διδακτορικής διατριβής από το Τμήμα Μηχανολόγων Μηχανικών της Πολυτεχνικής Σχολής του Πανεπιστημίου Θεσσαλίας δεν υποδηλώνει αποδοχή των απόψεων του συγγραφέα (Ν. 5343/32 αρ. 202 παρ. 2).

Εγκρίθηκε από τα Μέλη της Επταμελούς Εξεταστικής Επιτροπής:

Πρώτος Εξεταστής
(Επιβλέπων)

Δρ. Γεώργιος Λυμπερόπουλος, Καθηγητής
Τμήμα Μηχανολόγων Μηχανικών
Πανεπιστήμιο Θεσσαλίας

Δεύτερος Εξεταστής
(Μέλος Συμβουλευτικής
Επιτροπής)

Δρ. Δημήτριος Παντελής, Καθηγητής
Τμήμα Μηχανολόγων Μηχανικών
Πανεπιστήμιο Θεσσαλίας

Τρίτος Εξεταστής
(Μέλος Συμβουλευτικής
Επιτροπής)

Δρ. Απόστολος Μπουρνέτας, Καθηγητής
Τμήμα Μαθηματικών
Εθνικό και Καποδιστριακό Πανεπιστήμιο Αθηνών

Τέταρτος Εξεταστής

Δρ. Γεώργιος Ταγαράς, Καθηγητής
Τμήμα Μηχανολόγων Μηχανικών
Αριστοτέλειο Πανεπιστήμιο Θεσσαλονίκης

Πέμπτος Εξεταστής

Δρ. Κωνσταντίνος Αμπουντώλας, Αναπληρωτής Καθηγητής
Τμήμα Μηχανολόγων Μηχανικών
Πανεπιστήμιο Θεσσαλίας

Έκτος Εξεταστής

Δρ. Γεώργιος Κοζανίδης, Αναπληρωτής Καθηγητής
Τμήμα Μηχανολόγων Μηχανικών
Πανεπιστήμιο Θεσσαλίας

Έβδομος Εξεταστής

Δρ. Γεώργιος Σαχαρίδης, Αναπληρωτής Καθηγητής,
Τμήμα Μηχανολόγων Μηχανικών,
Πανεπιστήμιο Θεσσαλίας

ΠΟΛΙΤΙΚΕΣ ΠΡΟΜΗΘΕΥΤΩΝ ΥΠΟ ΕΞΑΡΤΩΜΕΝΗ ΑΠΟ ΤΗΝ ΕΞΥΠΗΡΕΤΗΣΗ ΖΗΤΗΣΗ ΑΓΟΡΑΣΤΩΝ: ΠΑΡΑΓΓΕΛΙΟΔΟΣΙΑ, ΑΝΤΑΓΩΝΙΣΜΟΣ ΚΑΙ ΕΠΙΛΟΓΗ ΑΓΟΡΑΣΤΩΝ

ΜΙΧΑΗΛ ΔΕΛΗΓΙΑΝΝΗΣ

Τμήμα Μηχανολόγων Μηχανικών, Πανεπιστήμιο Θεσσαλίας, 2022

Επιβλέπων Καθηγητής: Δρ. Γεώργιος Λυμπερόπουλος, Καθηγητής

Περίληψη

Η παρούσα διατριβή πραγματεύεται δυναμικές πολιτικές προμηθευτών των οποίων η ζήτηση εξαρτάται από την ποιότητα εξυπηρέτησης που παρέχουν. Η διατριβή διαρθρώνεται σε τρεις ενότητες. Η Ενότητα 1 διερευνά βέλτιστες δυναμικές πολιτικές αποθέματος όταν η συχνότητα επισκέψεων των αγοραστών καθορίζεται από προηγούμενες εξυπηρετήσεις. Η Ενότητα 2 εξετάζει τον ανταγωνισμό και τη συνεργασία προμηθευτών για την αφοσίωση των αγοραστών βάσει της εξυπηρέτησης που τους παρέχουν. Η Ενότητα 3 εξετάζει τη δυναμική παραγγελιοδοσία και επιλογή αγοραστών όταν η εξυπηρέτηση επηρεάζει τη μελλοντική ζήτηση. Ακολουθούν οι περιλήψεις των τριών ενοτήτων.

Ενότητα 1. Ένας αγοραστής που εκτίθεται σε έλλειψη αποθέματος μπορεί να χάσει την εμπιστοσύνη του και να είναι λιγότερο διατεθειμένος να επιλέξει τον ίδιο προμηθευτή στην επόμενη προμήθειά του. Αντίστροφα, μια θετική αγοραστική εμπειρία με διαθέσιμο απόθεμα μπορεί να αποκαταστήσει την προοπτική του προμηθευτή να επιλεγεί στο μέλλον. Ποια πρέπει να είναι η πολιτική ελέγχου αποθεμάτων του προμηθευτή σε αυτήν την περίπτωση; Για να αντιμετωπίσουμε αυτό το ερώτημα, αναπτύσσουμε ένα μοντέλο πολλαπλών περιόδων ενός αγοραστή που επιλέγει έναν προμηθευτή με πιθανότητα που εξαρτάται από την αξιολόγηση του προμηθευτή. Η αξιολόγηση αυτή αντικατοπτρίζει την εμπιστοσύνη του αγοραστή προς τον προμηθευτή βάσει της προηγούμενης εξυπηρέτησης που μετριέται με όρους περιστατικών διαθεσιμότητας/έλλειψης αποθέματος και ενημερώνεται από τον αγοραστή μετά από κάθε εξυπηρέτηση. Η βέλτιστη πολιτική αποθέματος του προμηθευτή διαχωρίζει τον χώρο αποθέματος σε διαστήματα παραγγελίας και μη παραγγελίας για κάθε επίπεδο αξιολόγησης. Η

βέλτιστη απόφαση εξαρτάται από το εάν η παραγγελία μειώνει τον κίνδυνο υποβάθμισης του προμηθευτή αρκετά ώστε να αντισταθμίσει την αύξηση του κόστους παραγγελίας και αποθέματος. Βρίσκουμε και αξιολογούμε τα όρια της βέλτιστης πολιτικής και παρουσιάζουμε ορισμένες από τις ιδιότητές της. Δίνουμε προϋποθέσεις για τη βελτιστότητα πολιτικών αποθέματος βάσης (basestock) και δείχνουμε ότι τέτοιες πολιτικές είναι βέλτιστες εάν υπάρχουν μόνο δύο αξιολογήσεις (καλή και κακή) ή εάν η ζήτηση του αγοραστή είναι σταθερή. Χρησιμοποιώντας το μοντέλο μας, υπολογίζουμε το κόστος αποθεμάτων στο πλαίσιο ενός προτύπου εφημεριδοπώλη (newsvendor). Τα αριθμητικά πειράματα υποδεικνύουν ότι (i) ο προμηθευτής μπορεί να επωφεληθεί από τη διατήρηση υψηλότερων αποθεμάτων όταν έχει μέτριες αξιολογήσεις παρά ακραίες αξιολογήσεις και από τη συναλλαγή με έναν αγοραστή που ανταποκρίνεται λιγότερο ακανόνιστα στην εξυπηρέτηση, (ii) οι πολιτικές basestock είναι αποτελεσματικές και (iii) χρησιμοποιώντας ένα αυθαίρετο κόστος έλλειψης αποθέματος στο πλαίσιο του προτύπου του εφημεριδοπώλη μπορεί να βλάψει σημαντικά τα κέρδη.

Ενότητα 2. Η συμπεριφορά μιας εταιρείας που αλλάζει εύκολα προμηθευτές (always-a-share) με γνώμονα την εξυπηρέτηση μπορεί να έχει σημαντική επίδραση στις ανταγωνιστικές και συνεργατικές πολιτικές αποθεμάτων των προμηθευτών της. Για να διερευνήσουμε αυτή την επίδραση, εξετάζουμε ένα μοντέλο ενός επαναλαμβανόμενου αγοραστή που μοιράζει την προτίμησή του μεταξύ δύο ετερογενών προμηθευτών τύπου εφημεριδοπώλη σε έναν άπειρο ορίζοντα. Για να πετύχει το μέγιστο πλεονέκτημα εξυπηρέτησης, ο αγοραστής επιβραβεύει τη διαθεσιμότητα του προϊόντος με επαναγορά (εμπιστοσύνη) και τιμωρεί την έλλειψη αποθέματος με αλλαγή (δυσπιστία) την επόμενη περίοδο. Για την αντιμετώπιση αυτής της συμπεριφοράς, η βέλτιστη πολιτική παραγγελιών κάθε προμηθευτή είναι μια πολιτική basestock με μη αρνητικό «ενεργό» επίπεδο basestock όταν ο προμηθευτής έχει την εμπιστοσύνη του αγοραστή και μηδενικό όταν δεν την έχει. Κάτω από ανταγωνισμό, το βέλτιστο ενεργό επίπεδο basestock κάθε προμηθευτή είναι μεγαλύτερο από το μυωπικό του επίπεδο basestock και είναι αύξον ως προς το ενεργό επίπεδο basestock του άλλου προμηθευτή. Κάτω από μια ελάχιστα περιοριστική συνθήκη, τα ενεργά επίπεδα basestock και των δύο προμηθευτών έχουν τουλάχιστον μία λύση ισορροπίας Nash καθαρής στρατηγικής (pure strategy). Εάν οι προμηθευτές συνεργάζονται, το βέλτιστο ενεργό επίπεδο basestock του προμηθευτή με το υψηλότερο/χαμηλότερο μυωπικό κέρδος είναι μεγαλύτερο/μικρότερο από το αντίστοιχο επίπεδο του μυωπικού επιπέδου basestock. Για να κατανοήσουμε καλύτερα αυτά τα αποτελέσματα, τα εφαρμόζουμε στην περίπτωση που η ζήτηση του αγοραστή έχει εκθετική

κατανομή. Αυτό μας επιτρέπει να λάβουμε ακριβείς εκφράσεις για τα βέλτιστα ενεργά επίπεδα basestock και τις συναρτήσεις κερδοφορίας, τις οποίες στη συνέχεια χρησιμοποιούμε σε μια αριθμητική ανάλυση ευαισθησίας. Ολοκληρώνουμε την ενότητα αυτή με μια συζήτηση για την επέκταση των αποτελεσμάτων σε περισσότερους από δύο προμηθευτές.

Ενότητα 3. Οι προμηθευτές που κατασκευάζουν προϊόντα για αποθήκευση (make-to-stock) και έχουν τακτικούς αγοραστές πρέπει να εξισορροπήσουν το κόστος των υπεραποθεμάτων έναντι του κόστους που προκύπτει από τις αντιδράσεις των αγοραστών όταν τα προϊόντα δεν είναι διαθέσιμα. Σε περίπτωση ελλείψεων, η επιλογή των αγοραστών που θα εξυπηρετηθούν και συνεπώς θα ικανοποιηθούν πρέπει να αντισταθμίζει τα τρέχοντα έσοδα από τους ικανοποιημένους αγοραστές έναντι των απωλειών της μελλοντικής ζήτησης από τους δυσαρεστημένους αγοραστές. Για να παράσχουμε κατανόηση και υποστήριξη αποφάσεων σχετικά για το παραπάνω πρόβλημα, (i) αναπτύσσουμε ένα καινοτόμο μοντέλο τύπου εφημεριδοπώλη μιας εταιρείας με ετερογενείς αγοραστές με ζήτηση που εξαρτάται από την εξυπηρέτηση, (ii) παρέχουμε ιδιότητες των βέλτιστων αποφάσεων παραγγελίας και επιλογής αγοραστών και (iii) αναπτύσσουμε μια καινοτόμα πολιτική τύπου δείκτη (index policy) για την επιλογή αγοραστών που βασίζεται στη Λαγκρανζιανή Χαλάρωση (Lagrangian Relaxation (LR)) και τη συγκρίνουμε με δύο άλλες πολιτικές δεικτών LR. Το μοντέλο αφορά μια εταιρεία που παραγγέλλει προϊόντα για μια ομάδα επανερχόμενων αγοραστών που αφήνουν διαφορετικά έσοδα και επισκέπτονται την εταιρεία με διαφορετικούς μέσους ρυθμούς που εξαρτώνται από το αν είναι ικανοποιημένοι ή δυσαρεστημένοι με την τελευταία τους εξυπηρέτηση. Η επιχείρηση επιλέγει ποιους αγοραστές θα εξυπηρετήσει εάν η ζήτηση υπερβαίνει την ποσότητα παραγγελίας (τρέχουσα δυναμικότητα). Για δύο αγοραστές, δείχνουμε ότι μια πολιτική Σταθερής Ποσότητας Παραγγελίας (Fixed Order Quantity (FOQ)) και μια πολιτική επιλογής αγοραστών τύπου δείκτη είναι βέλτιστες. Για περισσότερους αγοραστές, η βέλτιστη πολιτική περιλαμβάνει υπερπαραγγελιοδοσία (υποπαραγγελιοδοσία) όταν η συνολική ικανοποίηση του αγοραστή είναι υψηλή (χαμηλή) και επιλογή αγοραστών που μεγιστοποιούν τα τρέχοντα έσοδα (μελλοντική ζήτηση) όταν η συνολική ικανοποίηση των αγοραστών μετά τη ζήτηση είναι υψηλή (χαμηλή). Για να αντιμετωπίσουμε το πρόβλημα, δημιουργούμε τρεις πολιτικές δεικτών LR: μια πολιτική δείκτη Lagrange που χρησιμοποιεί μια ομοιόμορφη τιμή δυναμικότητας, μια πολιτική δείκτη Whittle που χρησιμοποιεί μια διακριτική τιμή δυναμικότητας και είναι μυωπικά βέλτιστη και μια καινοτόμα πολιτική δείκτη «ενεργού περιορισμού» που χρησιμοποιεί διακριτική τιμή δυναμικότητας όταν ο περιορισμός

δυναμικότητας είναι ενεργός. Τα αριθμητικά αποτελέσματα δείχνουν ότι η τελευταία πολιτική είναι σχεδόν βέλτιστη και υπερσχύει των άλλων δύο και ότι ο συνδυασμός της με τη κατάλληλη πολιτική FOQ, μπορεί να είναι πολύ αποτελεσματικός.

UNIVERSITY OF THESSALY
SCHOOL OF ENGINEERING
DEPARTMENT OF MECHANICAL ENGINEERING

Dissertation

**SUPPLIER POLICIES UNDER
SERVICE-DEPENDENT BUYER DEMAND:
ORDERING, COMPETITION, AND BUYER SELECTION**

by

MICHALIS DELIGIANNIS

M.Sc.in Mathematics with Specialization in Statistics and Modeling,
Aristotle University of Thessaloniki, 2016

B.S. in Mathematics,
Aristotle University of Thessaloniki, 2013

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2022

ix

© Copyright by Michalis Deligiannis, 2022

All Rights Reserved

Approved by

First Reader
(Supervisor)

Dr. George Liberopoulos, Professor
Department of Mechanical Engineering
University of Thessaly

Second Reader
(Member of Supervisory
Committee)

Dr. Dimitrios G Pandelis, Professor
Department of Mechanical Engineering
University of Thessaly

Third Reader
(Member of Supervisory
Committee)

Dr. Apostolos Burnetas, Professor
Department of Mathematics
National and Kapodistrian University of Athens

Fourth Reader

Dr. George Tagaras, Professor
Department of Mechanical Engineering
Aristotle University of Thessaloniki

Fifth Reader

Dr. Kostantinos Ampountolas, Associate Professor
Department of Mechanical Engineering
University of Thessaly

Sixth Reader

Dr. George Kozanidis, Associate Professor
Department of Mechanical Engineering
University of Thessaly

Seventh Reader

Dr. Georgios KD Saharidis, Associate Professor
Department of Mechanical Engineering
University of Thessaly

Acknowledgements

In the Spring of 2018, personal reasons brought me back to Greece from TU Eindhoven in the Netherlands, where I had spent two semesters of Ph.D. studies in Stochastic Operations Research. At that time, I was thinking about starting a Ph.D. in Finance, as I had already worked in Stochastic Modeling in Finance for my M.Sc. thesis in Statistics and Modeling. Then, I met my supervisor, Professor George Liberopoulos, who spent an extensive amount of time initiating me to the exciting world of Operations Research and Operations Management, especially, Dynamic Programming and Inventory Control, emphasizing the importance of quality and depth in research. I am very grateful for his devotion all these years and the great effort he put into our work. I was always impressed by the amount of time that he spent on our research—sometimes I noticed him editing our papers in Overleaf until late at night—and I am indebted for his guidance.

Apart from my supervisor, I would like to thank the other two members of the Supervisory Committee of my thesis, Professors Dimitrios Pandelis and Apostolos Burnetas, who have provided valuable guidance, especially in some of the ongoing work that we have started together. I am also thankful to the other members of the Examination Committee of my thesis, Professors Kostas Ampountolas, George Kozanidis, George Saharidis, and George Tagaras for their willingness to read and evaluate my work.

I want to thank my close friends whom I met during my studies in Thessaloniki: George Argyris (my best man as well), Aris Lazaridis, Antonis Frontistis, Dimitris Gkogkos, Kostas Kariotis, Kimon Gaitanis, Panagiotis Logothetis, Nikos Karabatza-kis, Polyzois Bountzis, Sotiris Skaperas, and Vasilis Apidopoulos. I am also fortunate

to have made new friends during my Ph.D. studies with whom we have shared similar experiences. I thank Apostolis Gkountas, Giannis Papachristos, Ilias Bellas, George Aretis, Kostis Papadimitropoulos, and Myron Benioudakis for the fruitful discussions we had and for their advice and friendship.

I am also grateful to my parents Giorgos and Panagiota, my sister Katerina, and my brothers Kostas and Vasilis for their love and support.

Words cannot express my gratitude to my wife Alexandra and my daughter Panagiota for the happiness and joy they bring me. I dedicate this thesis to them.

Michalis Deligiannis
Volos, December 2022

Abstract

This thesis deals with dynamic supplier policies under service-dependent demand. The thesis is structured in three parts. Part 1 explores optimal dynamic inventory policies when buyer purchase incidence is driven by past service, Part 2 looks at supplier competition and cooperation for buyer loyalty on service, and Part 3 tackles dynamic ordering and buyer selection when service affects future demand. A summary of each part follows.

Part 1. A buyer exposed to a stockout may lose goodwill and be less inclined to select the same supplier in his next procurement. Reversely, an in-stock experience may restore the supplier's prospect of being selected in the future. What should the supplier's inventory control policy be in this situation? To address this question, we develop a multiperiod model of a buyer who selects a supplier with a probability that depends on the supplier's rating. This rating reflects the buyer's goodwill towards the supplier based on past service, measured in terms of in-stock/out-of-stock incidents, and is updated by the buyer after each service. The supplier's optimal inventory policy partitions the inventory space in order-up-to and do-not-order intervals for each rating. The optimal decision depends on whether ordering reduces the supplier's risk of being downgraded enough to offset the increase in her ordering and inventory costs. We derive and evaluate bounds on the optimal policy and expose some of its properties. We obtain conditions for the optimality of basestock policies and show that such policies are optimal if there are only two ratings or if the buyer's demand is constant. Using our model, we impute the stockout cost in a newsvendor setting. Numerical experiments suggest that (i) the supplier may benefit from holding more inventory in intermediate than in extreme ratings, and from dealing with a buyer who

responds less erratically to service, (ii) basestock policies are efficient, and (iii) using an arbitrary stockout cost in the newsvendor setting can significantly hurt profits.

Part 2. A firm's service-driven always-a-share behavior may have a significant effect on the competitive and cooperative inventory policies of its suppliers. To explore this effect, we consider a model of a repeat buyer (she) sharing her patronage among two heterogeneous newsvendor-type suppliers over an infinite horizon. To enjoy the best service advantage, the buyer plays one supplier (him) against the other by rewarding product availability with repurchase (loyalty) and punishing stockouts with switching (disloyalty) in the next period. Faced with this behavior, the optimal ordering policy of each supplier is a basestock policy with a non-negative "active" basestock level when the supplier has the buyer's loyalty and a zero basestock level when he does not. Under competition, the optimal active basestock level of each supplier is greater than his myopic basestock level and increases in the other supplier's active basestock level. Under a mild condition, the active basestock levels of both suppliers have at least one pure-strategy Nash equilibrium solution. If the suppliers cooperate, the optimal active basestock level of the supplier with the highest/lowest myopic profit is greater/smaller than his myopic basestock level. To better comprehend these results, we apply them to the case where the buyer's demand is exponentially distributed. This allows us to obtain exact expressions for the optimal active basestock levels and payoff functions, which we then use in a numerical sensitivity analysis. We conclude with a discussion of the extension of the results to more than two suppliers.

Part 3. Make-to-stock suppliers with regular buyers must balance the cost of overstocking against the cost arising from the buyers' reactions when items are unavailable. In selecting which buyers to satisfy when shortages occur, they must weigh the current revenue from the satisfied buyers against the loss in future demand from the dissatisfied buyers. To provide insight and decision support on these trade-offs, (i) we develop a novel newsvendor model of a firm with heterogeneous buyers with service-dependent demand, (ii) we provide properties of the optimal ordering and buyer selection decisions of the firm, and (iii) we derive a novel Lagrangian Relaxation (LR)-based index policy for selecting buyers and compare it with two other LR-based index policies. The model concerns a firm that orders items for a group

of repeat buyers who generate different revenues and visit the firm with different average rates that depend on whether they are satisfied or dissatisfied with their last visit. The firm selects which buyers to serve if the demand exceeds the order quantity (current capacity). For two buyers, we show that a fixed order quantity (FOQ) policy and an index buyer selection policy are optimal. For more buyers, the optimal policy involves overstocking (understocking) when the overall buyer satisfaction is high (low) and selecting buyers that maximize the current revenue (future demand) when the overall buyer satisfaction after the demand is high (low). To tackle the problem, we derive three LR-based index policies: a Lagrangian index policy that uses a uniform capacity price, a Whittle index policy that uses a discriminatory capacity price and is myopically optimal, and a novel active-constraint index policy that uses a discriminatory capacity price when the capacity constraint is active. Numerical results indicate that the latter policy is near-optimal and outperforms the other two and that combining it with the right FOQ policy can be very efficient.

Contents

Acknowledgements	xiii
Abstract	xv
1 Introduction	1
1.1 Background and motivation	1
1.2 Thesis questions	6
1.3 Literature review	8
1.4 Thesis organization and contributions	12
2 Inventory policies with Service-Driven Demand	20
2.1 Introduction	20
2.2 Model description	21
2.3 Properties and structure of the optimal policy	26
2.4 When is a basestock policy optimal?	31
2.5 The case of constant buyer demand	34
2.6 Imputing the fixed stockout cost newsvendor model	38
2.7 Numerical results	42
2.7.1 Evaluation of bounds on optimal policy	42
2.7.2 Verification of optimal policy structure and effect of problem parameters on optimal policy	44
2.7.3 Performance evaluation of FS-equivalent SD model	49
2.8 Discussion and future research	52

3	Competition/Cooperation for Buyer demand	53
3.1	Introduction	53
3.2	Model formulation	54
3.3	The suppliers' optimal policy and payoff	56
3.4	Supplier competition	58
3.4.1	Best response function	59
3.4.2	Nash equilibrium	62
3.5	Supplier cooperation	64
3.5.1	The suppliers' gain	64
3.5.2	The buyer's perspective	69
3.6	Exponentially distributed demand	70
3.6.1	Analytical results	70
3.6.2	Numerical example	75
3.7	Extension to multiple sourcing	80
3.8	Discussion and future research	87
4	Ordering and Buyer Selection Policies	89
4.1	Introduction	89
4.2	Model formulation	90
4.3	On the optimal policy	96
4.3.1	Myopic policy	96
4.3.2	Properties of the optimal policy	98
4.4	The two-buyer problem	102
4.4.1	Optimal policy	102
4.4.2	Effect of the index on the optimal FOQ and the average expected profit	104
4.5	Current revenue vs. loss in future demand	107
4.6	Lagrangian relaxation	109
4.7	Index policies	111
4.7.1	Whittle index policy	112
4.7.2	Lagrangian index policy	112

4.7.3	Active-constraint index policy	115
4.8	Numerical results	117
4.8.1	Evaluation of index policies	117
4.8.2	Evaluation of FOQ policy	121
4.8.3	Effect of number of buyers and revenue rates	123
4.9	Discussion and future research	125
5	Thesis Summary	127
A	Chapter 2 Supplemental Material	130
B	Chapter 3 Supplemental Material	146
C	Chapter 4 Supplemental Material	156

List of Tables

4.1	Optimal FOQ and resulting average expected profit for $z_i > z_j$, $i \neq j$, and $n = 2$	102
4.2	Optimal FOQ and resulting average expected profit under $\mathbf{u}^x$; difference in average expected profit under $\mathbf{u}^z$ and $\mathbf{u}^x$ when $y^z = 1$, for $x_i > x_j$, $z_i < z_j$, $i \neq j$, and $n = 2$	104
4.3	Average performance (sample mean and standard deviation) of all policies, for 250 instances of a five-buyer problem.	120
4.4	Percent difference (sample mean and standard deviation) of the average performance of the optimal FOQ policy (under all selection policies) from the optimal policy, for 250 instances of a five-buyer problem. . .	122
4.5	Average performance (sample mean and standard deviation) of all policies, for 150 instances of a ten-buyer problem.	124

List of Figures

2.1	Optimal inventory control policy.	29
2.2	Optimal inventory control policy for $p = 4.8$, $h = 3$, $b = 1$ and different values of q_{M-1}	36
2.3	LBR, LBM, LBH, UB, and UBR, for 200 instances.	43
2.4	Four different profiles of q_α vs. α for $M = 5$	44
2.5	$V_\alpha(x)$ vs. $x, x \geq 0$, and optimal order-up-to and do-not-order regions for $h = 0.5$, $b = 0.4$, $r = 3.6$, $\beta = 0.9$, and $\theta = 2.2$, for the four q_α profiles in Figure 2.4.	45
2.6	S_α^0 vs. α for $b = 0.8$, $r = 3.4$, red $M = 5$, four values of h , two demand distributions, and the four q_α profiles in Figure 2.4.	47
2.7	S_α^0 vs. q_α for $b = 0.8$, $r = 3.4$, the linear q_α profile, two values of h , four values of M , and two demand distributions.	48
2.8	Six different profiles of q_α vs. α for $M = 5$	49
2.9	$\tilde{\Pi}(S^*)$, S^* , $\tilde{q}(S^*)$, and $\hat{b}^*$, for the six q_α profiles in Figure 2.6.	50
2.10	Left: $\Delta\tilde{\Pi}(\hat{b})$, for $\hat{b} = m\hat{b}^*$, $m = 0, 0.5, 1, 2, 4$, for q_α profile 3 in Figure 2.10. Right: $\Delta\tilde{\Pi}(\hat{b}^*)$, for q_α profiles 1–5 in Figure 2.10.	51
3.1	Sample trajectory of the suppliers' inventory levels.	58
3.2	Best response functions of the two suppliers and Nash equilibrium (s_i^e, s_j^e)	63
3.3	Position of the optimal active basestock level pair when the two suppliers cooperate.	67
3.4	Optimal basestock levels and performance measures vs. c_j	76
3.5	Optimal basestock levels and performance measures vs. I_j	77
3.6	Optimal basestock levels and performance measures vs. r_j	77

3.7	Optimal basestock levels and performance measures vs. b_j	78
3.8	Close-up of optimal active basestock levels and performance measures vs. b_j	78
3.9	Optimal active basestock levels and performance measures vs. r	79
3.10	Optimal active basestock levels and performance measures vs. b	79
3.11	Adjustment backorder penalty cost b^c vs. c , I , and r , for $b = 0$	80
3.12	Markov chain transition diagram of the ranking vector for $n = 4$	82
4.1	Optimal FOQ policy under the optimal index policy $\mathbf{u}^z$ for $n = 2$ ($c = 1$).	103
4.2	Optimal FOQ policy under index policy $\mathbf{u}^r$ for $n = 2$ ($c = 1$).	105
4.3	Optimal FOQ policy under index policy $\mathbf{u}^s$ for $n = 2$ ($c = 1$).	105
4.4	State transition diagram (omitting self-transitions) corresponding to the optimal ordering policy $y^r(\alpha)$ under the revenue-greedy policy $\mathbf{u}^r$, for a two-buyer example with $r_1 > r_2$, $\gamma_1 > \gamma_2$, and $\bar{\gamma}_1 < \bar{\gamma}_2$	106
4.5	State transition diagram (omitting self-transitions) corresponding to the optimal ordering and buyer selection policy, y^* and $\mathbf{u}^*$, for a three- buyer example with $r_1 > r_2 > r_3$, $\gamma_1 = \gamma_2 = \gamma_3$ and $\bar{\gamma}_1 = \bar{\gamma}_2 < \bar{\gamma}_3$	108
4.6	Optimal Lagrange multiplier λ^* vs. $y(\alpha)$	111
4.7	Optimal Lagrangian index $l_{(i)}^*$ vs. buyer index (i) using the optimal Lagrange multiplier λ^*	114
4.8	Average expected profit and order quantity under different policies for 100 instances of a problem with five buyers.	119
4.9	Average expected fill rate for buyers 1, 3, and 5 under different policies for 100 instances of a problem with five buyers.	119
4.10	Percent profit loss of the FOQ policy under the optimal selection, for different FOQ values, for seven representative instances of a problem with five buyers.	123
4.11	Average expected profit and order quantity under different policies for 100 instances of a problem with ten buyers.	124

A.1	Reduction in the probability that the supplier stocks out, if she orders a small quantity ϵ , for different demand density functions and initial inventory levels.	131
B.1	Constant demand-share contour lines $\pi_i(s_i, s_j)$ for different values of $\pi_i \in (0, 1)$	150

Chapter 1

Introduction

In this chapter, we introduce the topic of this thesis. In Section 1.1, we present the background and motivation behind this work, in Section 1.2, we pose the main questions that we address, in Section 1.3, we review the related literature, and in Section 1.4, we summarize the main contributions of the thesis.

1.1 Background and motivation

Buyer reactions to service. Manufacturers and suppliers of industrial market goods must balance the cost of overstocking against the cost arising from the buyers' reactions when not enough products are available on demand. These reactions can vary significantly depending on the extent of the inconvenience that stockouts cause, which ranges from increased administrative costs to production disruptions to lost sales damages. In the worst-case scenario, stockouts can cause an extraordinary upheaval in entire industrial sectors, even the global economy, with factories around the world limiting or even halting their operations, despite powerful demand for their wares. This has been the case with the recent shortages of goods, such as computer chips, construction materials, and many others, reflecting the disruption of the COVID-19 pandemic and other catastrophic events, combined with decades of companies limiting their inventories in the pursuit of Just-in-Time (Goodman and Chokshi, 2021).

Although the adverse impact of retailer stockouts on consumer behavior has been researched extensively for over fifty years, the literature on buyer reactions to stockouts in B2B markets is scarce. According to an earlier survey, most buyers return to their original suppliers after experiencing a stockout, but firms need to assess the losses of those buyers who remain with an alternate source of supply that was pursued because of a stockout (Dion and Banting, 1995). Another study reports empirical evidence based on corporate financial data that firms are more inclined to avoid stockouts with more inventory when profit margins are higher and that the likelihood of losing the demand of disenfranchised buyers depends on the alternative sources of supply that are available (Blazencko and Vandezande, 2003).

Multiple sourcing. Many firms use multiple sourcing to hedge against operational and disruption risks and stimulate competition among their suppliers (Minner, 2003; Tang, 2006; Svoboda, Minner and Yao, 2021). The trend toward multiple or dual sourcing has been intensified in the aftermath of the COVID-19 pandemic and other disasters (Alicke, Barriball, Foster, Mauhourat and Trautwein, 2022). Having established multiple supply channels, they can easily switch their patronage from one supplier to another, especially for products for which supplier switching costs are low. Though having multiple suppliers for the same products can add complexity and cost to a buyer, many of the risks of multiple sourcing can be mitigated by making sure that the buyer is collaborating with high-quality suppliers.

To use multiple sourcing, buyers develop procurement strategies based on which they select their supply partners. Using scorecard-based assessment methods, which are readily available in most ERP systems (e.g., SAP (2022b); Oracle (2021)), they evaluate potential suppliers against strategic criteria related to cost, quality, service, social responsibility, risk, agility, etc., to create a shortlist of certified suppliers that best meet the criteria. While cost optimization is still a critically important criterion for supplier selection, in many cases, it is no longer at the top of the list (Tang, 2006). Companies are looking for partnerships that meet high standards in several areas, with product quality and on-time delivery being among the leading criteria (Bosch, 2020; Intel, 2020; Samsung, 2020).

Product availability in particular has emerged as a decisive selection factor in the

wake of major catastrophic events that have exposed the vulnerability of companies that rely heavily on one or a limited number of trading partners. Indicatively, during the COVID-19 pandemic, nearly one-fifth of Wendy's restaurants in the US run out of beef due to severe meat shortages caused by COVID-19 outbreaks in meat processing plants (Valinsky, 2020).

Often, the shortlist of suppliers is limited to two suppliers, because the additional effort required to track the performance and manage the relationship with more suppliers can become counterproductive. Also, as the number of suppliers increases, the relationship developed with each supplier weakens, and the suppliers tend to pay more attention to their other business partners, including the buyer's competitors. The predominance of dual sourcing over multiple sourcing is reflected in the inventory control literature, where 70% of the publications consider dual sourcing and only 30% look at multiple sourcing (Svoboda et al., 2021).

Supplier selection. Once a supplier base has been created at the strategic level, the buyer must decide how to divide the demand among the suppliers at the operational level. Depending on the market context, the buyer may allocate the demand to all the suppliers, select one supplier to fill it, or even announce the demand to all the suppliers and award it to the first supplier who fills it (Armony and Plambeck, 2005). If the demand allocated to a supplier is not fully met, the missing items may be substituted with compatible items from the same supplier, backordered, procured via transshipment from another supplier, or canceled.

In some situations, the supplier uses flexible long-term supply contracts that impose some restriction on the buyer, usually in the form of a commitment to purchase certain minimum quantities (Tang and Tomlin, 2008). The most flexible commitment is to specify a total minimum quantity where the buyer is allowed to place any order in any period, as long as the cumulative orders across all periods exceed this quantity (Bassok and Anupindi, 1997). This type of flexibility allows the buyer to occasionally call on all suppliers and reassess their service. At the same time, it guarantees a minimum level of trade for the suppliers, helping them to stay in business—therefore, also keeping the buyer's supply channels open—and giving them the opportunity

to compete for a larger share of the demand by improving their service level. Examples of companies that preferentially allocate all their business to suppliers that performed well in the past but hesitate to withdraw business from suppliers who performed poorly are presented in Andrews and Barron (2016) in the context of a dynamic *avored supplier allocation* rule.

In other situations, the buyer can leverage the demand allocation decision to foster competition between the suppliers. This can be done by allocating the demand to the suppliers based on either promised or past performance. The first approach has been studied extensively in the following context. The buyer sets up a rule for allocating the demand, and the suppliers compete for a share or all of the demand by offering their price, capacity, service level, lead time, or some other performance measure, depending on the rule. The allocation of demand based on past performance has been less researched. Most of the literature in this area concerns firms that compete for market share on service in a B2C environment. The market share of each firm evolves smoothly over time as customers flow in and out of its customer pool depending on the service they receive from the firm and its competitors.

To monitor suppliers and stimulate competition between them, many companies use formal quantitative rating schemes (Li, Zhang and Fine, 2013). Such schemes are commonplace in many ERP systems. For example, the Supplier Rating System of Oracle's PeopleSoft enables companies to group and weigh the KPIs of suppliers into categories that are further weighted and grouped into an overall composite supplier score. This score is then compared to a rating scale and assigned a rating, much like a report card in school. For instance, a score 0–100 can be transformed to a rating A–F using the following rating rule: $A \geq 90$, $B \geq 80$, $C \geq 70$, $D \geq 60$, $F \leq 59$. The supplier scorecard is constantly updated and accessible across the company and to the supplier (Oracle, 2020).

While a supplier with a high rating clearly has an advantage over a competitor with a lower rating, it is not obvious that the former supplier will always be selected over the latter supplier, because the selection process can be highly uncertain due to the buyers' perceived difficulty in predicting supplier performance (Riedl, Kaufmann, Zimmermann and Perols, 2013). Tang and Tomlin (2008) attribute this difficulty

to the low visibility and control level of supply chain partners which reduces the confidence of each partner in the information provided by the other partners. Despite the seamless access to information enabled by technological advances, suppliers may disguise the on-hand inventory or lead time that they quote to their buyers, and buyers may inflate the demand forecasts that they provide to their suppliers (Christopher and Lee, 2004). The uncertainty in supplier selection may be further amplified due to the multiple decision makers it involves, e.g., it may be the outcome of an opinion poll by the buyer's purchasing managers (Benjaafar, Elahi and Donohue, 2007). As a result of this uncertainty, the buyer's request incidents may appear as random to the supplier.

The issue of supplier selection has been extensively studied in the literature, mostly from the buyer's perspective, in the context of strategic supply chain contracting and coordination. An important aspect that has been overlooked from the supplier's viewpoint is the buyers' dynamic behavior in response to supplier service and in particular stockout incidents. Designing inventory control policies that account for the adverse effect of stockouts on buyer (or customer) goodwill and future demand has long been a challenging issue for OR/OM researchers and practitioners. Traditionally, a penalty cost or a service level constraint has been used to address this issue, but this approach ignores the dependence of demand on stockouts. In the last two decades, a stream of research has emerged, whose origins can be traced to the 1970s, that endogenizes customer reaction to stockouts into the demand dynamics, predominantly in B2C environments and at an aggregate-demand level.

Buyer selection. In addition to diverse reactions to stockouts, buyers also have different margins due to customized pricing arising from differences in their market power, agreement with the supplier, sales volume, location, etc. A recent study on wholesale price discrimination reports empirical evidence from a market where some buyers pay up to 70% more than others for the same good on the same day (Marshall, 2020).

In the face of the buyers' heterogeneous demand dynamics and margins, firms must dynamically decide how many items to order in advance of demand, given that buyers may be at different satisfaction (goodwill) levels from previous encounters,

and which buyers to select to satisfy if the order falls short of demand. To address the adverse effect of stockouts on buyer goodwill and future demand, firms typically use a penalty cost or a service level constraint. This cost is supposed to reflect the impact on future demand due to the loss of buyer goodwill following a stockout, yet the demand is almost always considered to be independent of past service.

Moreover, in practice, firms often prioritize buyers based on their past sales (Cachon and Lariviere, 1999). This approach is also reflected in many ERP systems, such as SAP (SAP, 2022a). Prioritization based on past sales has been found to positively affect high-priority buyers (Homburg, Droll and Totzek, 2008) although it may also potentially undermine profitability by inducing important buyers to feel more entitled than grateful (Wetzel, Hammerschmidt and Zablah, 2014). When the buyers' demand is sensitive to past service, prioritizing buyers based on past sales has the risk of becoming a self-fulfilling prophecy. That is, if up to a certain period, buyer i happens to have higher past sales than buyer j and both buyers compete for the same product, i will be selected, her total sales will further increase, and she will be satisfied, positively impacting her expected future sales. On the other hand, the total sales of buyer j will remain unchanged, and she will be dissatisfied, adversely affecting her expected future sales. If buyer j happens to have higher sales than i up to the same period, the tables will be turned and j instead of i will be selected.

Sheffi (2020) discusses several product-allocation schemes used at times of scarcity, on the occasion of the global shortage of semiconductor chips triggered by the COVID-19 pandemic. Among them are the fair treatment of all buyers and the prioritization of powerful buyers such as Apple and Samsung, high-margin buyers, or vulnerable buyers, especially when the product is essential to the buyer's (or the buyer's customers') survival. As is pointed out, the downside of these approaches is that they ignore the long-term importance of a buyer to the firm.

1.2 Thesis questions

In view of the issues discussed in the previous section, and motivated by recent technological advances that enable the collection and analysis of big data on individual

customers, allowing for identifying their unique repurchase behavior at different levels of rated satisfaction (Mittal and Kamakura, 2001) or following stockouts (Fitzsimons, 2000), different types of decision-making problems and questions arise for suppliers facing service-dependent demand. In the main part of this thesis (Chapters 2–4), we consider three such problems, each focusing on a different setting and set of questions.

In **Chapter 2**, we focus on the asymmetric responses to good and bad service of a buyer with memory of past service and its implications on the inventory policy of the supplier. More specifically, we consider the problem of a buyer who uses a rating scheme with a finite number of ratings, e.g., A–F, as mentioned earlier, and visits a supplier with a rating-dependent probability, uprating/derating the supplier after an in-stock/out-of-stock incidence. The questions that we ask are:

- What is the structure of the supplier’s optimal inventory policy as a function of her inventory and rating?
- Can a basestock policy be optimal, and if so, under what conditions?
- Should the supplier stock more when her rating is low or high?
- What is the imputed cost of a stockout under the buyer’s rating-dependent visit behavior?

In **Chapter 3**, we focus on the switching behavior of a buyer from one supplier to another following poor service and its implication on the suppliers’ competitive inventory policy, in a B2B setting. The setting that we consider fits the description of the *always-a-share* model introduced in Jackson (1985), which assumes that a firm making purchases of some product category repeatedly can easily switch its patronage from one supplier to another, therefore sharing its patronage among multiple suppliers. Jackson notes that in some situations suggesting always-a-share behavior, a firm may make a series of purchases each from a single supplier but share its patronage among suppliers over time. As examples of always-a-share firms, she lists buyers of simple machine tools, commodity chemicals, carbon steel, and apartment building owners who purchase major appliances, among others. Jackson (1985) contrasts the always-a-share model with the *lost-for-good* model, where a firm faces high costs of

switching suppliers and therefore is either totally committed to one supplier or totally lost and committed to some other supplier. That setting is outside the scope of this paper. Given the buyer's always-a-share behavior, we pose the following questions:

- What is the optimal inventory policy of each supplier in response to the other supplier's decision?
- Do the suppliers' inventory policies reach equilibrium and if so, is it unique, and how is it related to their myopic inventory policy?
- What is the optimal joint inventory policy and gain for the suppliers if they team up?
- What are the implications for the buyer if the suppliers cooperate instead of competing?

In **Chapter 4**, we focus on the dynamic ordering and buyer selection decisions of a supplier with many buyers with service-dependent demand. These decisions require the careful balancing of the ordering cost, the current revenue from the satisfied buyers, and the loss in future demand from the dissatisfied buyers, raising several important questions for the supplier:

- What is the interaction between ordering and buyer selection decisions?
- How sensitive is performance to each decision?
- When does future demand matter more than the current revenue in buyer selection?
- How efficient is it to order a fixed quantity and how efficient is it to select buyers based on a fixed prioritization?

1.3 Literature review

In this section, we review the literature that is related to our work. For ease of presentation, we organize it into three parts.

Service-driven demand. The first model in which demand changes dynamically as a function of the service level can be traced to Fergani (1976). In the last two decades, there has been a renewed interest in similar dynamic demand models, matching the increasing evidence from data-driven marketing studies that stockouts have an adverse effect on long-term demand Campo, Gijbrenchts and Nisol (2003); Anderson, Fitzimons and Simester (2006); Jing and Lewis (2011). Dynamic models can be divided into two categories. The first considers demand at an aggregate level and the second focuses on the individual customer level.

Notable papers in the first category are Fergani (1976), Hall and Porteus (2000), Liu, Shang and Wu (2007), and Olsen and Parker (2008). These studies consider single-supplier models or duopolies, where the demand of each supplier in each period is a linear function of the market size. This assumption is key for ensuring the optimality of basestock policies with basestock levels that are proportional to the market size. They also assume that individual customers behave homogeneously toward each supplier and have no memory of past service. Robinson (2016) considers a more general demand model where in each period the mean demand changes linearly in the number of satisfied and unsatisfied customers. The optimal policy for this model is not in general stationary and will vary with the mean demand, which may increase or decrease unboundedly; therefore, finding it is computationally intractable.

Notable papers in the second category are Gans (2002), Gaur and Park (2007), Liberopoulos and Tsikis (2007), and Deng, Shen and Shanthikumar (2014). The first two papers consider multiperiod models with multiple customers and suppliers, where each supplier maintains a constant service level, and the customers learn about this level from experience. Liberopoulos and Tsikis (2007) introduce a duopoly model of two suppliers competing for one customer. Each supplier can be in any of several “credibility levels” that affect the probability of being chosen by the customer. The evolution of these levels depends on the service experiences of the customer. Based on a restricted numerical study, they find that for geometrically distributed demand, the optimal stationary policy of the two suppliers at equilibrium is a basestock policy. Deng et al. (2014) consider a similar model to that in Liberopoulos and Tsikis (2007) involving a single supplier with several customers. In each period, each customer

demands exactly one unit with a probability that depends on his contentment level, which can be in any of two states: satisfied or unsatisfied. They show that the optimal inventory level always increases in the number of satisfied customers.

Supplier competition. Numerous papers explore how a single buyer can stimulate competition among multiple suppliers by allocating her demand to the suppliers based on their price, service quality, or other competitive dimensions. In Kalai, Kamien and Rubinovitch (1992), Gilbert and Weng (1998), Cachon and Zhang (2007), Benjaafar et al. (2007), and Elahi (2013), the suppliers are modeled as Make-to-Order (MTO) or Make-to-Stock (MTS) service systems, and the buyer allocates her demand to the suppliers based on their service quality, measured in terms of service rate, lead time, or service level, or assigns all the demand to a randomly selected supplier, where the probability of selecting a supplier is based on the supplier's service quality. In Ha, Li and Ng (2003) and Jin and Ryan (2012), the suppliers are modeled as EOQ firms and MTS queues, respectively, and demand is allocated based on price and delivery frequency or price and service level, respectively. In many cases, it turns out that allocating the demand to one supplier, i.e., selecting a supplier, is optimal for the buyer. A review of some of these models can be found in Wang, Eallace, Shen and Choi (2015).

Some papers investigate the sourcing strategy of a buyer and the pricing strategies of unreliable suppliers under an environment of supply disruption (Babich, Burnetas and Ritchken, 2007; Li, Wang and Cheng, 2010). There is also a sizable literature on newsvendor competition as is manifested by the review articles that have appeared in the last two decades (Cachon and Zhang, 2006; Nagarajan and Sošić, 2008; Chinchuluun, Karakitsiou and Mavrommati, 2008; Silbermayr, 2020). Much of this work involves lateral transshipments, consolidation of inventories at a central location, and product substitution or complementarity in case of a stockout. In all of these works, the buyer has no memory of service, so the models are essentially single-period. Dynamic volume allocation in an infinite-horizon setting is considered in Li et al. (2013) in a problem in which the buyer induces the desired supplier behavior through business share allocation based on supplier performance.

Buyer selection. The problem of heterogeneous customer selection has been

extensively studied under the assumption that service does not affect demand. Two paradigms of selection problems are the *inverse newsvendor* Carr and Lovejoy (2000); Choi and Ketzenberg (2018); Bavafa, Leys, Örmeci and Savin (2019) and the *selective newsvendor* Taaffe, Romeijn and Tirumalasetty (2008); Taaffe et al. (2008); Chahar and Taaffe (2009); Abdel-Aal and Selim (2019). In the first problem, a firm with a given service level and several customer classes, each with a predefined priority and a random demand, must choose the fraction of each class to serve, i.e., it must choose the demand distribution. In the second problem, a newsvendor serving multiple buyers must decide the order quantity and select which buyers to serve. The demand of each buyer is influenced by the marketing or pricing effort, and buyer selection takes place before the demand realization. In a related paper, Durango-Cohen and Li (2017) consider a supplier who must decide her order quantity and allocate it to several heterogeneous customers with contracts to demand within a specified range and the right to receive a penalty for any unmet demand within that range.

Another stream of research focuses on the strategic competition of customers under a given capacity allocation policy of the supplier. An example is the *turn-and-earn* policy where the supplier allocates capacity to customers based on past sales, motivating customers to influence their future allocations by increasing their sales Cachon and Lariviere (1999); Lu and Lariviere (2012).

Adelman and Mersereau (2013) consider the problem of a supplier who must dynamically allocate capacity among a finite number of heterogeneous customers with different margins and different demands that depend on past fill rates. They investigate when and how goodwill matters and they demonstrate that an approximate dynamic programming policy that rationalizes the fill rates that the firm provides to each customer can achieve higher rewards than margin-greedy and Lagrangian-derived policies. They interpret this policy as prioritizing each customer using an “adjusted margin” that augments the customer’s margin by an amount that values the goodwill impact of meeting current demand.

Moreover, Klein and Kolb (2015) considers a provider with several customers, each belonging to one of a finite number of segments defined by a combination of customer properties, recency, and purchase intention. The provider must decide which

customers with positive purchase intention to accept and which to deny. The problem is formulated as an MDP where the provider's action determines the transition probability of each customer from one segment to another, and the overall state is the number of customers in each segment. A myopic policy that maximizes the current revenue is compared with the optimal MDP solution in a numerical study of a problem with up to two customer properties and two recency states. In both papers, the firm's capacity is fixed. None of the two papers provides analytical results on the optimal policy, although Adelman and Mersereau (2013) shows that the margin-greedy policy is asymptotically optimal when the number of customers tends to infinity, and optimal when the demand is deterministic.

1.4 Thesis organization and contributions

The main contribution of this thesis is the development and analysis of three novel stochastic models, presented in Chapters 2–4, that provide insight and decision support for firms (suppliers) facing service-dependent demand. The main questions that each model addresses were presented in Section 1.2. In this section, we briefly describe each model and the main conclusions that we draw from its analysis.

In **Chapter 2**, we develop a multiperiod model of a supplier (she) selling items to a buyer (he) who rates the supplier based on the history of her service, measured in terms of in-stock/out-of-stock incidents. At the beginning of each period, the supplier orders a quantity, ahead of the demand, based on her rating and inventory surplus/backlog, and receives it before the end of the period. At the end of the period, the buyer generates a random demand and selects the supplier to fill this demand with a probability that depends on her rating. If the supplier fails to meet all the demand at once, the buyer backorders the unmet demand with her but downgrades her. In addition, the supplier incurs a backorder penalty cost that is proportional to the shortage. This cost is a direct measurable cost of the shortage, e.g., a price discount per item short or an overtime cost for the procurement and handling of the backlogged items. If the supplier meets all the demand at once, she is upgraded by the buyer and carries over any leftover inventory to the next period. Using dynamic programming

principles and Markov chain analysis, we draw the following conclusions regarding the supplier's inventory control policy.

Myopic policy. The supplier's myopic (single-period) inventory control policy is a basestock policy with non-negative, rating-dependent basestock levels that are non-decreasing in her rating.

Structure of optimal policy. For the infinite-horizon discounted expected profit problem, the optimal policy partitions the inventory space in multiple order-up-to and do-not-order intervals, defined by successive order-up-to and reorder points, for each rating. The optimal decision—order up to the next point or do not order—depends on whether ordering reduces the supplier's risk of being downgraded enough to offset the increase in her ordering and inventory holding costs. This tradeoff depends both on the supplier's inventory level and the probability density function of the buyer's demand. We show that the smallest order-up-to-point for each rating is greater than or equal to the basestock level of the myopic policy for that rating. This implies that it is optimal to satisfy all backorders and that using the myopic policy will lead to profit losses. Unlike the basestock levels of the myopic policy, the smallest order-up-to points are not necessarily non-decreasing in the rating, even though the discounted expected profit is. Numerical results show that it can be optimal for the supplier to hold more inventory in intermediate ratings than in extreme ratings and that the more erratic the buyer's response to service, the higher the inventory level.

Bounds on optimal policy. We derive upper and lower bounds on the optimal inventory control policy and numerically evaluate them and compare them against existing bounds. Our results show that in many instances, a heuristic basestock policy with rating-dependent basestock levels that are higher than the respective myopic levels and are non-decreasing in the rating, is near-optimal.

Optimality of basestock policies. We show that under a certain condition on the buyer demand distribution, the optimal policy reduces to a basestock policy with non-negative, rating-dependent basestock levels. This condition is always satisfied if the demand density function is non-increasing. We also present a condition for the optimality of a policy that is effectively basestock. Two special cases of this condition arise when either the smallest order-up-to points or the smallest reorder points are

increasing in the rating. We show that the smallest order-up-to points for the first two ratings are always increasing in the ratings, implying that in the case of two ratings, a basestock policy is effectively optimal. For this case, we also derive analytical expressions for computing these points and evaluate them for two distributions of buyer demand (exponential and uniform).

Constant buyer demand. For the case of constant buyer demand, we show that for the average expected profit problem, the optimal policy is a basestock policy in which the supplier: (i) always operates in a make-to-stock mode with no backlogs and her rating is absorbed in the largest level, (ii) always operates in a make-to-order mode only with backlogs, and her rating is absorbed in the smallest level, or (iii) operates in a make-to-stock mode in all ratings, except for the largest rating where she operates in a make-to-order mode, and her rating is absorbed in the largest two values, alternating between them. We derive conditions for determining which of the three above cases is optimal. These conditions depend on the supplier selection probabilities in the largest, second largest, and smallest rating only, as well as on the revenue and cost parameters. We show that in the case of two ratings, the policy of alternating between the two ratings is never optimal.

Fixed stockout cost. Finally, we consider a variant of the newsvendor model studied in Çetinkaya and Parlar (1998), with a fixed stockout cost, representing the buyer's loss of goodwill due to a stockout, in addition to the variable backorder cost. To estimate the fixed cost, we relate this model to the service-driven demand model developed in this chapter, operated under a basestock policy with a common basestock level for all ratings. Numerical results for different functional forms of the supplier selection probability w.r.t. to the rating show that if the imputed fixed cost is used in the newsvendor model to compute the optimal basestock level, the drop in the average expected profit with respect to the maximum profit is limited on average. However, choosing an arbitrary value for the fixed stockout cost can lead to significant profit losses, especially if this value is smaller than the imputed value.

In **Chapter 3**, we develop a stylized model of two newsvendor-type suppliers with inventory carryover and backordering (we also discuss the extension to multiple

suppliers) who provide the same product to an always-a-share buyer (she) and compete for the buyer's business over an infinite horizon. There is little room for price differentiation because the product is standard and the suppliers have anyway been shortlisted among a larger group of candidates based on their more or less equally competitive prices. Therefore, the suppliers compete on the service they provide.

Among the various determinants of service quality, we restrict our attention to *product availability* which has emerged as a decisive factor in the wake of severe global shortages that have exposed the vulnerability of supply chains to the disruption of major catastrophic events.

To enjoy the best availability advantage, the buyer in our model plays one supplier (him) against the other by rewarding availability with repurchase (loyalty) and punishing stockouts with switching (disloyalty) in the next period. Faced with this “carrot-and-stick” behavior, each supplier must decide his ordering policy to maximize his long-run expected average profit by balancing his current inventory and backorder cost against his future profit loss resulting from ceding the buyer's loyalty to his competitor.

Using stochastic optimization and game-theoretic analysis, we characterize the optimal inventory policy of the suppliers under competition and relate it to their myopic policy. To measure the service level gain of the buyer and the respective profit loss of the suppliers brought about by competition, we also characterize the optimal joint policy of the suppliers if they decide to cooperate. To better comprehend the results and their implications, we apply them and evaluate them numerically in the case where the buyer's demand is exponentially distributed. Finally, we extend the results to more than two suppliers under a round-robin supplier selection policy. Based on our results, we draw the following conclusions.

Optimal inventory policy. The myopic policy of each supplier is identical to the basestock policy of a multi-period newsvendor who seeks to minimize his expected inventory and backorder cost. The optimal inventory policy of each supplier under competition and cooperation is also a basestock policy with a non-negative “active” basestock level when the supplier has the buyer's loyalty and a zero basestock level when he does not.

Competition. Under competition, each supplier raises his active basestock level above his myopic level, sacrificing his myopic profit to extend his stay at the top of the buyer's list. The optimal active basestock level of each supplier is increasing in the other supplier's active level, sparking inventory competition between the suppliers to the buyer's advantage. Under a mild condition on the buyer's demand distribution, implying that the demand density function does not increase sharply above each supplier's myopic basestock level, the best response of each supplier has a unique global maximizer above his myopic basestock level that guarantees the existence of a pure-strategy Nash equilibrium which is symmetric for symmetric suppliers. The equilibrium is unique if the suppliers' best response functions are contraction mappings or if the suppliers are symmetric (under a stricter condition on the buyer's demand distribution).

Cooperation. Under cooperation, each supplier sets his active basestock level at his myopic level, if the myopic profits of both suppliers are the same. Otherwise, the supplier who has the smallest myopic profit, sets his active basestock level below his myopic level, ceding a part of his long-term demand share to the more profitable supplier, who sets his active basestock level above his myopic level but below his active basestock level at equilibrium under competition. Under a condition that again involves the buyer's demand density function, the active basestock level of the less profitable supplier drops to zero, meaning that this supplier cedes all his demand share to the more profitable supplier, except for the occasional times when the buyer returns to him for one period following a stockout by the more profitable supplier.

The buyer's perspective. Cooperation benefits the suppliers as it results in reduced inventories for them. This cancels out the high-fill rate advantage that the buyer enjoys thanks to her carrot-and-stick behavior when the suppliers compete. To counter this setback, the buyer can charge the cooperating suppliers an extra backorder penalty cost rate every time she faces a stockout. For symmetric suppliers, the penalty rate that makes the buyer fully recover her fill rate under competition is increasing in the symmetric active basestock level of the suppliers at equilibrium, which hinges on the buyer's demand distribution and the suppliers' margin-to-interest-rate ratio.

Exponentially distributed demand. When the buyer's demand is exponentially distributed, the conditions guaranteeing the uniqueness of the maximizer of the best response function of each supplier and the Nash equilibrium are satisfied. In this case, we obtain exact expressions for the active basestock levels at equilibrium and under cooperation. The former expressions depend mainly on the tradeoff between the supplier's inventory cost rate and profit margin, while the latter expressions depend on the tradeoff between the suppliers' inventory and backorder cost rates.

Multiple sourcing. Most of the general results for two suppliers extend to multiple suppliers if the buyer uses a *round-robin* policy where she switches suppliers on a circular basis after each stockout.

In **Chapter 4**, we study a newsvendor model of a firm that orders items for a group of repeat buyers. The buyers generate different revenues and have different average visit rates that depend on whether they are satisfied or dissatisfied with their last visit. If the demand exceeds the order quantity (current capacity), the firm must select which buyers to serve without violating capacity. We formulate the firm's problem as an average-profit Markov decision process (MDP) whose state is the vector of buyer satisfaction states and where the decisions are made in two stages: Before the demand is realized (ex-ante), the firm must decide its order quantity, and after the demand takes place (ex-post), it must select which buyers to serve.

Using stochastic analysis, we characterize the myopic policy and the optimal policy for two buyers, and we provide some properties and conjectures on the optimal policy for multiple buyers. We also numerically compare three Lagrangian relaxation-based index policies for selecting buyers, where an index policy is defined as an ex-ante prioritization of buyers based on the value of some function (index). The three policies are the Lagrangian index, the Whittle index, and the active-constraint index policy. The index in each policy is derived by relaxing the capacity constraint of the original problem and solving a separate problem for each buyer using a penalty price that internalizes the relaxed constraint.

In the Lagrangian index, the price is uniform (common for all buyers) and arbitrary. The price that we use in our numerical experiments is derived in closed form as the solution of the Lagrangian dual. This price depends on capacity and yields the

tightest bound of the original problem. In the Whittle index, the price is discriminatory (buyer-specific) and independent of capacity. In the active-constraint index, the price is also discriminatory but is applied only when demand exceeds capacity, so it depends on both capacity and the demand characteristics of all buyers. Based on our results, we draw the following conclusions.

Optimal buyer selection. When the order quantity suffices to cover the demand of all but one buyer, the optimal selection policy is an index policy where the index of each buyer (she) is increasing in three terms: her revenue rate, the loss in her future demand (average visit rate) if she is not served, and the type-I service level of all other buyers if she is served. This result enables the full characterization of the optimal selection policy for two buyers. In general, however, the optimal selection is not index-based but depends on the realization of demand. Our analysis shows that it tends to maximize the current revenue if the buyers' ex-post satisfaction level is high and maximize future demand if it is low.

Index-based selection. The three Lagrangian relaxation-based index policies that we compare have varying degrees of efficiency depending on how well they internalize the relaxed constraint into the index. The Whittle index is simply the revenue rate, so it does not internalize capacity. Prioritizing buyers based on their revenue rates, while myopically optimal, can be arbitrarily bad in the long term because it ignores future demand. The Lagrangian index outperforms the Whittle index because it accounts for the loss in future demand and leads to the required usage of capacity on average through the uniform price that it uses.

The active-constraint index depends on the same three terms as the optimal index for the above-mentioned case where the order quantity is enough to cover the demand of all but one buyer and is optimal in that case. The third term in particular accounts for the effect of the selection of one buyer on the stockout probability of the other buyers, based on their satisfaction states. It makes the firm dynamically readjust its goal between maximizing the current revenue and future demand based on the ex-ante satisfaction state vector, echoing the observed optimal policy, and leading to more well-balanced satisfaction states and service levels among the buyers. Our numerical results show that the active-constraint index policy is near-optimal.

Optimal and fixed order quantities. Under optimal buyer selection, the optimal order quantity is non-decreasing in the satisfaction state, matching supply to demand. If buyers are selected red suboptimally, the firm may benefit from ordering fewer items in higher satisfaction states if this allows it to reach more profitable states which cannot be approached with the suboptimal selection policy.

For two buyers, the optimal order quantity is fixed for all satisfaction states, under the optimal buyer selection. If the firm selects buyers inefficiently, the fixed order quantity may increase or decrease as the firm tries to make up for the loss of efficiency by overstocking or understocking, respectively. This means that it may prefer to satisfy both buyers all the time or not satisfy any buyer at all rather than prioritize the wrong buyer. For more buyers, using a fixed order quantity can be quite efficient if this quantity is optimally chosen, but can lead to severe losses if it is not.

Finally, **Chapter 5** provides a summary of our main findings.

Supplemental material for Chapters 2–4, including proofs, can be found in Appendices A–C, respectively.

Chapter 2

Inventory policies when buyer demand is driven by past service

2.1 Introduction

In this chapter, we develop a multiperiod model of a supplier (she) selling items to a buyer (he) who rates the supplier based on the history of her service, measured in terms of in-stock/out-of-stock incidents. In Section 2.2, we develop the service-driven demand model. In Section 2.3, we determine the myopic policy, derive bounds on the optimal policy for the infinite-horizon problem, and present the structure and properties of the optimal policy. In Section 2.4, we explore the optimality of basestock policies, and in Section 2.5, we characterize the optimal policy for the constant-demand case. In Section 2.6, we impute the fixed stockout cost in the newsvendor model from the service-driven demand model developed in Section 2.2. In Section 2.7, we present numerical results, and in Section 2.8, we summarize the results and propose directions for future work. Supplemental material for this chapter, including proofs, can be found in Appendix A.

2.2 Model description

A profit-maximizing supplier sells items to a buyer. The buyer rates the supplier based on the history of her service, defined in terms of in-stock/out-of-stock incidents. In each period t , the supplier orders a non-negative quantity based on her current rating α_t and inventory level x_t ; α_t belongs to a finite set of discrete values $A = \{1, \dots, M\}$, e.g., as in the A–F score system mentioned earlier, and x_t can be positive or negative, indicating surplus or backlog, respectively. Due to her lead time, the supplier places her order at the beginning of the period, ahead of the buyer’s demand, in a make-to-stock mode Benjaafar et al. (2007). This type of ordering process is common in practice including the computer and apparel industries Tang and Tomlin (2008), where manufacturers “preposition” (produce or purchase prior to demand and hold inventory) buyer-specific, semi-finished components with long lead times that are incorporated into end-products with much shorter lead times, e.g., ICs for specific types of printers, or greige fabric (a fabric that has been woven or knitted but not yet dyed or printed) for specific types of sports garments. The order arrives before the end of the period, raising the supplier’s inventory level to $y_t \geq x_t$.

At the end of the period, the buyer demands a quantity w_t and selects the supplier, with probability q_{α_t} , or an outside source, with probability $\bar{q}_{\alpha_t} = 1 - q_{\alpha_t}$, to fill this demand; q_{α_t} is referred to as the rating-dependent selection probability of the supplier. The demands $\{w_t, t = 0, 1, \dots\}$ are based on the buyer’s needs and are independent of the supplier’s past service. We assume that they i.i.d. continuous random variables with p.d.f., c.d.f., and mean, $f(\cdot)$, $F(\cdot)$, and θ , respectively. Based on these assumptions, the demand seen by the supplier in period t , $d_t(\alpha_t)$, is given by

$$d_t(\alpha_t) = \begin{cases} w_t, & \text{w.p. } q_{\alpha_t}, \\ 0, & \text{w.p. } \bar{q}_{\alpha_t}. \end{cases} \quad (2.1)$$

For notational simplicity, henceforth we will drop the dependence of $d_t(\alpha_t)$ on α_t . If the supplier is not selected by the buyer, her rating remains unchanged. If she is selected, she fills all the demand or the part of it that she can cover from inventory. If she fails to meet all the demand at once, the buyer backorders the unmet demand

with her to ensure the uniformity and traceability of his order and receives the missing items in the next period. In this case, the supplier incurs a backorder penalty cost that is proportional to the shortage. At the same time, the buyer downgrades the supplier by one rating point (unless her rating is already at its lowest value) due to the overall disruption that the missing items cause him.

The notion that a stockout incident has a fixed adverse effect on the supplier's standing, irrespectively of the shortage quantity or time, has been addressed in the literature by assuming a "lumpsum," "red-tape," or "negative image" fixed cost per stockout occasion Çetinkaya and Parlar (1998) or a "type-1" service constraint imposing a minimum probability that demand will be immediately served from inventory. An example of this effect is when a production line is stopped whether 1 unit or 100 units are short Nahmias and Olsen (2015). Another example is when a supplier faces buyer loss if the demand from a buyer cannot be met a certain number of times, as in the case of a pharmaceutical distributor selling medicines to pharmacies Saracoglu, Topaloglu and Keskindurk (2014). In our model, we incorporate this effect into the demand dynamics through the rating process.

If the supplier meets all the demand at once, she is upgraded by one rating point (unless her rating is already at its highest value, M), and carries over any leftover inventory to the next period. By holding inventory, she expects to fully meet the buyer's demand in the next period and improve or maintain her rating. Alternatively, she may choose not to hold inventory, in which case she will operate in a make-to-order mode, compromising her rating. Reserving inventory for a specific buyer is not that unusual, especially if this buyer has agreed to purchase a minimum order quantity per period on average or has a predominant position among the supplier's partners. Competition that breeds demand for customized make-to-stock is also reported in cases where a buyer places duplicate orders to several suppliers and buys from the supplier who fills the order first, canceling all other orders Li (1992). This practice is common for microchip suppliers in the semiconductor industry where yields and processing times are unpredictable, lead times are long, and products are highly customized. A similar practice for network product suppliers is reported in Armony and Plambeck (2005).

Based on the above assumptions, the supplier's inventory state is updated as follows,

$$x_{t+1} = y_t - d_t = \begin{cases} y_t - w_t, & \text{w.p. } q_{\alpha_t}, \\ y_t, & \text{w.p. } \bar{q}_{\alpha_t}, \end{cases} \quad (2.2)$$

and the supplier's rating state is updated as follows,

$$\alpha_{t+1} = \begin{cases} \alpha_t + \delta_{\alpha_t}^+ - \delta_{\alpha_t}^-, & \text{w.p. } q_{\alpha_t}, \\ \alpha_t, & \text{w.p. } \bar{q}_{\alpha_t}, \end{cases} \quad (2.3)$$

where

$$\delta_{\alpha_t}^+ = 1_{\{w_t \leq y_t, \alpha_t < M\}} \text{ and } \delta_{\alpha_t}^- = 1_{\{w_t > y_t, \alpha_t > 1\}}, \quad (2.4)$$

and $1_{\{\cdot\}}$ is the indicator function.

As is natural to assume, the probability that the buyer selects the supplier is non-decreasing in the supplier's rating. We also assume that the buyer may select the supplier even if her rating is at the lowest level. This would be the case, e.g., if the two parties had agreed on a minimum average order quantity per period, $q_1\theta$. In mathematical terms,

$$q_{\alpha+1} \geq q_{\alpha}, \alpha \in \{1, \dots, M-1\} \text{ and } q_1 > 0. \quad (2.5)$$

While the selection probability is non-decreasing in the supplier's rating, its exact functional form is not restricted and depends on the buyer's response to service. Such a response may in general be asymmetric, as has been documented in the behavioral economics literature in a B2C context Kahneman and Tversky (1979). Indicatively, in a recent study of supermarket consumers, Koos and Shaikh (2019) reports an asymmetric S-shape relationship between customer dissatisfaction due to stockouts and customer response. A different large-scale study of automotive customers finds that the functional form relating rated satisfaction to repurchase behavior exhibits increasing returns Mittal and Kamakura (2001). On the other hand, under loss-aversion Tversky and Kahneman (1991), the functional form should exhibit decreasing

returns. These behaviors are different and depend on the buyer's characteristics and the market context. The proposed model is flexible and can accommodate different behavioral patterns whether the setting is B2C or B2B.

In each period, the supplier incurs a cost c per item ordered and receives a revenue r per item sold. The quantity sold is $\min(y_t, d_t)$. She also incurs an inventory holding cost h per item in inventory and backorder cost b per item short, at the end of the period. We assume discounting with rate $\beta < 1$. To ensure that the supplier can be profitable even with backorders, we also assume that $\beta p > b$, where p is the per unit profit defined as $p = r - c$.

The profit of the supplier in period t is $r[(x_t)^- + \min(y_t, d_t)] - c(y_t - x_t) - h(y_t - d_t)^+ - b(d_t - y_t)^+$, where $(x)^+ = \max(x, 0)$ and $(x)^- = (-x)^+$, $x \in \mathbb{R}$. After replacing $\min(y_t, d_t)$ and $(y_t - d_t)^+$ with $d_t - (d_t - y_t)^+$ and $y_t - d_t + (d_t - y_t)^+$, respectively, and rearranging terms, the profit can be written as $r(x_t)^- + cx_t + (r + h)d_t - (c + h)y_t - (r + b + h)(d_t - y_t)^+$. Rolling back the terms $r(x_t)^-$ and cx_t into period $t - 1$ using (2.2) and discounting them at rate β , the profit can be redefined as $\beta r(x_{t+1})^- + \beta cx_{t+1} + (r + h)d_t - (c + h)y_t - (r + b + h)(d_t - y_t)^+$ (see Olsen and Parker (2008) and Robinson (2016) for similar treatments). For $t = 0$, $r(x_0)^-$ and cx_0 are not rolled back but are added to the total profit as an extra term $r(x_0)^- + cx_0$, which can be rewritten as $c(x_0)^+ + p(x_0)^-$, after replacing x_0 with $(x_0)^+ - (x_0)^-$. Finally, after replacing x_{t+1} and $(x_{t+1})^-$ with $y_t - d_t$ and $(d_t - y_t)^+$, respectively, and collecting terms, the redefined profit in period t reduces to:

$$K_3 d_t - K_1 y_t - K_2 (d_t - y_t)^+,$$

where K_1 , K_2 , and K_3 are positive constants given by

$$K_1 = (1 - \beta)c + h, \tag{2.6}$$

$$K_2 = (1 - \beta)r + b + h = K_1 + (1 - \beta)p + b, \tag{2.7}$$

$$K_3 = (r - \beta c + h) = K_1 + p. \tag{2.8}$$

Essentially, the redefined profit in period t refers to the profit in the interval starting

after the arrival of the supplier's order in period t and ending before the arrival of her order in period $t + 1$; so, it is expressed in terms of y_t instead of x_t . From the above definitions and the assumption $\beta p > b$, K_1, K_2 , and K_3 are ordered as follows:

$$0 < K_1 < K_2 < K_3. \quad (2.9)$$

The problem of the supplier is to select order-up-to levels $y_t \geq x_t, t = 0, 1, \dots$, to maximize her discounted expected profit over an infinite horizon, $\Pi_{\alpha_0}(x_0)$, defined as

$$\Pi_{\alpha_0}(x_0) = c(x_0)^+ + p(x_0)^- + V_{\alpha_0}(x_0), \quad (2.10)$$

where $V_{\alpha_0}(x_0)$ is a value function given by

$$V_{\alpha_0}(x_0) = \max_{y_t \geq x_t} E_{d_t} \left\{ \sum_{t=0}^{\infty} \beta^t [K_3 d_t - K_1 y_t - K_2 (d_t - y_t)^+] \right\}.$$

The term in the square brackets is the redefined profit in period t . Its expected value, denoted by $\Lambda_{\alpha_t}(y_t)$, is given by

$$\Lambda_{\alpha_t}(y_t) = K_3 q_{\alpha_t} \theta - L_{\alpha_t}(y_t), \quad (2.11)$$

where $L_{\alpha_t}(y_t)$ is the expected cost of the supplier in period t and is given by

$$L_{\alpha_t}(y_t) = K_1 y_t + K_2 [q_{\alpha_t} B(y_t) + \bar{q}_{\alpha_t} (y_t)^-], \quad (2.12)$$

with $B(y)$ denoting the expected backlog, defined as $B(y) = E[(w - y)^+]$. With these definitions, $V_{\alpha_0}(x_0)$ can be rewritten as

$$V_{\alpha_0}(x_0) = \max_{y_t \geq x_t} \left\{ \sum_{t=0}^{\infty} \beta^t \Lambda_{\alpha_t}(y_t) \right\}.$$

The value function $V_{\alpha_t}(x_t)$, in any period t , satisfies the following dynamic programming (Bellman) equation:

$$V_{\alpha_t}(x_t) = \max_{y_t \geq x_t} \left\{ \Lambda_{\alpha_t}(y_t) + \beta E_{dt} [V_{\alpha_{t+1}}(x_{t+1})] \right\}, x_t \in \mathbb{R}, \alpha_t \in A. \quad (2.13)$$

Given the boundedness of $\Lambda_{\alpha_t}(y_t)$ and $V_{\alpha_t}(x_t)$, equation (2.13) has a unique solution. From (2.11) and (2.2)-(2.4), the Bellman equation can be written as follows, after dropping the time index:

$$V_{\alpha}(x) = K_3 q_{\alpha} \theta + \max_{y \geq x} H_{\alpha}(y), x \in \mathbb{R}, \alpha \in A, \quad (2.14)$$

where

$$\begin{aligned} H_{\alpha}(y) = & -L_{\alpha}(y) + \beta \left\{ q_{\alpha} \left[\int_0^y V_{\alpha+\delta_{\alpha}^+}(y-w) dF(w) \right. \right. \\ & \left. \left. + \int_y^{\infty} V_{\alpha-\delta_{\alpha}^-}(y-w) dF(w) \right] + \bar{q}_{\alpha} V_{\alpha}(y) \right\}. \end{aligned} \quad (2.15)$$

2.3 Properties and structure of the optimal policy

Before setting out to characterize the optimal stationary inventory control policy for the infinite-horizon problem, $y_{\alpha}^*(x)$, we determine the optimal policy for the single-period problem—henceforth referred to as the *myopic* policy—denoted by $y_{\alpha}^{my}(x)$.

Proposition 2.1. *The myopic policy is a basestock policy given by*

$$y_{\alpha}^{my}(x) = \max(x, S_{\alpha}^{my}), \quad (2.16)$$

where the rating-dependent basestock levels $S_{\alpha}^{my}, \alpha \in A$, are given by

$$S_{\alpha}^{my} = F^{-1} \left(\left[1 - \frac{K_1}{q_{\alpha} K_2} \right]^+ \right). \quad (2.17)$$

Note that K_1/K_2 in (2.17) is the well-known critical ratio in the newsvendor model with surplus cost K_1 and backlog cost $K_2 - K_1$. This ratio is independent of α . What makes S_{α}^{my} dependent on α is q_{α} in (2.17). From (2.5) and (2.17), S_{α}^{my} is non-decreasing in α . The following proposition provides properties and bounds on

$\Pi_{\alpha_0}(x_0)$ and $V_{\alpha_0}(x_0)$ for the infinite-horizon problem, where for notational simplicity we have dropped the time index.

Proposition 2.2. *For $\alpha \in A$ and $x \in \mathbb{R}$, $\Pi_\alpha(x)$ and $V_\alpha(x)$ satisfy*

$$\lim_{x \rightarrow \infty} \Pi_\alpha(x) = \lim_{x \rightarrow \infty} V_\alpha(x) = -\infty, \quad (2.18)$$

$$V_\alpha(0) > 0, \quad (2.19)$$

$$V_{\alpha'}(x) \geq V_\alpha(x), \alpha' > \alpha, \quad (2.20)$$

$$\Pi_{\alpha'}(x) \geq \Pi_\alpha(x), \alpha' > \alpha, \quad (2.21)$$

$$\frac{-h(x - S_\alpha)^+}{1 - \beta} + V_\alpha^L(S_\alpha) \leq V_\alpha(x) \leq V_\alpha^U(S_\alpha^{my}), \quad (2.22)$$

$$c(x)^+ + p(x)^- - \frac{h(x - S_\alpha)^+}{1 - \beta} + V_\alpha^L(S_\alpha) \leq \Pi_\alpha(x) \leq c(x)^+ + p(x)^- + V_\alpha^U(S_\alpha^{my}), \quad (2.23)$$

where S_α are arbitrarily chosen non-negative basestock levels that are non-decreasing in α , and $V_\alpha^L(S_\alpha)$ and $V_\alpha^U(S_\alpha^{my})$ are constants that are also non-decreasing in α and are given by expressions (A.3) and (A.6) in Appendix A.

Inequalities (2.20)-(2.21) state that the higher the initial rating, the higher the expected discounted profit over an infinite horizon, for the same initial inventory; however, this does not mean that the optimal inventory level is increasing in the rating. The upper bounds in (2.22)-(2.23) are constructed by considering the myopic policy under an ideal scenario in which α_t increases by $\delta_{\alpha_t}^+$ whenever the buyer selects the supplier, irrespectively of whether the demand is met or not, and remains unchanged, otherwise. In the proof, we note that the upper bound developed for a similar problem in Robinson (2016) is obtained by further allowing the supplier to order after observing the demand. The lower bound is constructed by considering an order-up-to policy with rating-dependent order-up-to points S_α that are non-decreasing in α . Intuitively, S_α values satisfying $S_\alpha \geq S_\alpha^{my}$ are likely to produce tighter lower bounds, because the order-up-to point under the optimal policy is greater than that under the myopic policy, as we show later in Proposition 2.3. We also note that the lower bound developed

in Robinson (2016) is $V_\alpha^L(0)$, where $0 \leq V_\alpha^L(0) \leq V_\alpha^L(S_\alpha^{my})$, implying (2.19). Finally, observe that in the newsvendor model, where $q_\alpha = q, \alpha \in A$, the myopic policy is optimal and independent of α , and (2.22) becomes $V^L(S^{my}) = V(x) = V^U(S^{my})$, $x \leq S^{my}$, after dropping index α . A quantity that plays an important role in our analysis is the difference $\beta[V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)]$ that represents the supplier's discounted future profit loss following a stockout when her rating is α . Although evaluating this difference is in general computationally intractable, it can be bounded using expressions (2.20) and (2.22), as follows:

$$0 \leq V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0) \leq \Delta_\alpha, \quad (2.24)$$

where

$$\Delta_\alpha = V_{\alpha+\delta_\alpha^+}^U(S_{\alpha+\delta_\alpha^+}^{my}) - V_{\alpha-\delta_\alpha^-}^L(S_{\alpha-\delta_\alpha^-}). \quad (2.25)$$

From our discussion following Proposition 2.2, in the newsvendor model, $\Delta_\alpha = 0$. From the Bellman equation, we can derive the following properties regarding the structure of the optimal policy, $y_\alpha^*(x)$, and $V_\alpha(x)$.

Lemma 2.1. *The optimal inventory control policy is to satisfy all backorders, i.e.,*

$$y_\alpha^*(x) \geq (x)^+. \quad (2.26)$$

Lemma 2.1 implies that $y_\alpha^*(x) \geq 0$. For this reason, henceforth, we will restrict our attention to the case where $y \geq 0$. In this case, $L_\alpha(y)$ in (2.12) becomes

$$L_\alpha(y) = K_1 y + K_2 q_\alpha B(y). \quad (2.27)$$

From the above analysis, $H_\alpha(y)$ is continuous and tends to $-\infty$ as $y \rightarrow \infty$. Moreover, it is bounded from above and its global maximum is non-negative. In general, it may have several local maxima, even though the per period profit is concave. This is a major deviation—and a source of substantial difficulty in the analysis—from the newsvendor model where the equivalent function preserves the concavity of the per-period profit. As a result, under the optimal stationary inventory control policy, for

each rating α , the inventory space is partitioned into regions that are separated by multiple threshold points, $S_\alpha^0, s_\alpha^1, S_\alpha^1, \dots, s_\alpha^i, S_\alpha^i, \dots, s_\alpha^n, S_\alpha^n$, with the following properties:

- (i) $0 \leq S_\alpha^0 < s_\alpha^1 < S_\alpha^1 < \dots < s_\alpha^i < S_\alpha^i < \dots < s_\alpha^n < S_\alpha^n$,
- (ii) $H'_\alpha(S_\alpha^i) = 0$ and $H''_\alpha(S_\alpha^i) < 0, i = 1, 2, \dots, n$; if $S_\alpha^0 = 0$, then $H'_\alpha(S_\alpha^0) \leq 0$,
- (iii) $H_\alpha(S_\alpha^{i-1}) > H_\alpha(S_\alpha^i)$ and $H_\alpha(s_\alpha^i) = H_\alpha(S_\alpha^i), i = 1, 2, \dots, n$.

The optimal policy is to order up to S_α^i in region $R_\alpha^i = [s_\alpha^i, S_\alpha^i]$ and not order in region $\bar{R}_\alpha^i = [S_\alpha^i, s_\alpha^{i+1}], i = 0, \dots, n$, where by convention, $s_\alpha^0 = -\infty$ and $s_\alpha^{n+1} = \infty$. Note that at $s_\alpha^i, i = 1, \dots, n$, it is optimal both to order up to S_α^i and not order. Figure 2.1 illustrates the optimal policy for $n = 2$.




Figure 2.1: Optimal inventory control policy.

In mathematical terms, the optimal policy is given by the following expression:

$$y_\alpha^*(x) = \sum_{i=0}^n S_\alpha^i 1_{\{x \in R_\alpha^i\}} + x 1_{\{x \in \bar{R}_\alpha^i\}}. \quad (2.28)$$

The intuition behind the optimal policy is discussed in Section A in Appendix A using two examples of demand density functions $f(w)$, which suggest that n depends on the shape of $f(w)$ and is bounded by the number of its local maxima, which typically is only one or two at most.

Lemma 2.2. *The derivative of the value function is non-positive, i.e.,*

$$V'_\alpha(x) \leq 0, x \in \mathbb{R}. \quad (2.29)$$

Proof. In view of (2.28), $V_\alpha(x)$ in (2.14) and its first derivative can be written as follows:

$$V_\alpha(x) = K_3 q_\alpha \theta + \sum_{i=0}^n H_\alpha(S_\alpha^i) 1_{\{x \in R_\alpha^i\}} + H_\alpha(x) 1_{\{x \in \bar{R}_\alpha^i\}}. \quad (2.30)$$

$$V'_\alpha(x) = H'_\alpha(x) 1_{\{x \in \cup_{i=0}^n \bar{R}_\alpha^i\}} \leq 0. \quad (2.31)$$

Clearly, $V'_\alpha(x) = 0$, if $x \in \cup_{i=0}^n R_\alpha^i$, as is shown in Figure 2.1. $\square$

Expression (2.30) states that $V_\alpha(x)$ is constant and equal to $K_3 q_\alpha \theta + H_\alpha(S_\alpha^i)$ in region R_α^i . In region $\bar{R}_\alpha^i$, $V_\alpha(x)$ drops parallelly to $H_\alpha(y)$, $x = y$, as is shown in Figure 2.1. For $x < 0$, (2.30) and (2.31) imply that $V_\alpha(x) = K_3 q_\alpha \theta + H_\alpha(S_\alpha^0) = V_\alpha(0)$ and $V'_\alpha(x) = 0$, respectively. Consequently, the second integral in (2.15) becomes $V_{\alpha-\delta_\alpha^-}(0) \bar{F}(y)$, and $H_\alpha(y)$ can be simplified. The simplified form of $H_\alpha(y)$ and its first two derivatives are

$$H_\alpha(y) = -L_\alpha(y) + \beta \left\{ q_\alpha \left[\int_0^y V_{\alpha+\delta_\alpha^+}(y-w) dF(w) + V_{\alpha-\delta_\alpha^-}(0) \bar{F}(w) \right] + \bar{q}_\alpha V_\alpha(y) \right\}, \quad (2.32)$$

$$H'_\alpha(y) = -L'_\alpha(y) + \beta \left\{ q_\alpha [f(y)(V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)) + \int_0^y V'_{\alpha+\delta_\alpha^+}(y-w) dF(w)] + \bar{q}_\alpha V'_\alpha(y) \right\}, \quad (2.33)$$

$$H''_\alpha(y) = -L''_\alpha(y) + \beta \left\{ q_\alpha \left[f'(y) (V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)) + f(y) V'_{\alpha+\delta_\alpha^+}(0^+) + \int_0^y V''_{\alpha+\delta_\alpha^+}(y-w) dF(w) \right] + \bar{q}_\alpha V''_\alpha(y) \right\}, \quad (2.34)$$

The following proposition provides lower and upper bounds on S_α^0 and S_α^n , respectively.

Proposition 2.3. $S_\alpha^{my} \leq S_\alpha^0 \leq S_\alpha^n \leq \bar{S}_\alpha^n, \alpha \in A$, where S_α^{my} is given by (2.17) and $\bar{S}_\alpha^n$ is given by

$$\bar{S}_\alpha^n = \arg \min_{y \geq (x)^+} \left\{ K_2 \bar{F}(y) + \beta \Delta_\alpha f(y) \leq \frac{K_1}{q_\alpha} \right\}, \alpha \in A, \quad (2.35)$$

where Δ_α is given by (2.25).

Proposition 2.3 indicates that the order quantity under the optimal policy is greater than or equal to that under the myopic policy; hence, using the myopic policy will lead to profit losses. Proposition 2.3 also provides an upper bound on the maximum order-up-to level which is useful for designing storage capacity.

2.4 When is a basestock policy optimal?

As we saw in Section 2.3, this is not in general the case in our service-driven demand model. However, if S_α^0 is the unique maximizer of $H_\alpha(y)$ each α , the optimal policy is a basestock policy with rating-dependent basestock levels S_α^0 , defined as follows:

$$y_\alpha^*(x) = \max(x, S_\alpha^0). \quad (2.36)$$

The following proposition provides a sufficient condition for the concavity of $H_\alpha(y)$ which ensures the existence of a unique maximum and the optimality of a basestock policy.

Proposition 2.4. *If $f(y)$ satisfies*

$$\frac{f'(y)}{f(y)} \leq \frac{K_2}{\beta \Delta_\alpha}, y \geq 0, \forall \alpha \in A, \quad (2.37)$$

where Δ_α is given by (2.25), then $H_\alpha(y)$ is concave for all α , and therefore the optimal stationary inventory control policy is a basestock policy with rating-dependent basestock levels given by (2.36).

Obviously, if f is non-increasing (e.g., exponential, uniform, Weibull and Gamma with shape parameter ≤ 1 , etc.), then (2.37) immediately holds for every $y \geq 0$. Even

if $f(y)$ is increasing for some y , however, condition (2.37) may still hold if its r.h.s. is large enough. Intuitively, a large value of $K_2/\beta\Delta_\alpha$ means that the myopic backorder cost is more important than future profit losses following a stockout, suggesting that the optimal policy should be similar in structure to the myopic policy, which is a basestock policy. Note that in the newsvendor model, $\Delta_\alpha = 0$; therefore, (2.37) always holds, verifying that the optimal inventory control policy is a basestock policy, as was also mentioned earlier. If $f(y)$ is log-concave (e.g., normal, logistic, Weibull, Beta, Gamma with shape parameter ≥ 1 , etc.), then $f'(y)/f(y)$ decreases in y Bagnoli and Bergstrom (2005). In this case, (2.37) holds for every $y > 0$, as long as it holds for $y = 0$.

A sufficient condition for the optimality of a basestock policy that does not require the concavity of $H_\alpha(y)$ is $H'_\alpha(y) \leq 0, y \geq S_\alpha^0$. Evaluating this condition, however, is practically impossible, because no analytical expression for S_α^0 exists. An exception is the case $\bar{S}_\alpha^n = 0$, where all order-up-to points including S_α^0 collapse onto zero, making a basestock policy with zero basestock level (equivalent to a make-to-order policy) optimal for rating α , and all ratings above α transient.

Finally, Proposition 2.5 provides a condition under which the optimal policy effectively is a basestock policy with rating-dependent basestock levels S_α^0 .

Proposition 2.5. *The optimal stationary inventory control policy effectively is a basestock policy with rating-dependent basestock levels S_α^0 , if*

$$x_0 \leq \min_{\alpha' \geq \alpha_0} (s_{\alpha'}^1) \quad \text{and} \quad S_\alpha^0 \leq \min_{\alpha' \geq \alpha} (s_{\alpha'}^1), \forall \alpha \in A, \quad (2.38)$$

where x_0 and α_0 are the initial inventory level and rating, respectively.

The proof is straightforward and is based on showing that (2.38) ensures that $x_t \leq s_{\alpha_t}^1, t \geq 0$, implying that x_t always belongs in regions $R_{\alpha_t}^0$ and $\bar{R}_{\alpha_t}^0$, where $y_{\alpha_t}^* = \max(x_t, S_{\alpha_t}^0)$ from (2.28). Two special cases where (2.38) holds are when the global maximizers S_α^0 (respectively, the reorder points s_α^1) are non-decreasing in α . These cases are given by Corollary 2.1.

Corollary 2.1. *The optimal stationary inventory control policy effectively is a base-stock policy with rating-dependent basestock levels S_α^0 , if either*

$$x_0 \leq S_{\alpha_0}^0 \text{ and } S_\alpha^0 \leq S_{\alpha+\delta_\alpha^+}^0, \forall \alpha \in A, \text{ or} \quad (2.39)$$

$$x_0 \leq s_{\alpha_0}^1 \text{ and } s_\alpha^1 \leq s_{\alpha+\delta_\alpha^+}^1, \forall \alpha \in A, \quad (2.40)$$

where x_0 and α_0 are the initial inventory level and rating, respectively.

Using the Bellman equation and the first-order conditions, we can derive conditions under which S_α^0 is non-decreasing in α for all α , but these conditions are hard to verify as they involve the simultaneous solution of many non-linear equations and are too complicated to provide any useful insights. In fact, in some cases, it is easier to show that S_α^0 is decreasing in α for some α . For example, it can be argued that for $M > 2$, if $q_M = q_{M-1} > q_{M-2}$, then $S_{M-1}^0 > S_M^0$. Intuitively, this happens because in both ratings $M-1$ and M , the demand distribution seen by the supplier is the same, making the myopic profits equal. Rating $M-1$, however, is “riskier” than M , because it borders a lower rating, $M-2$. To hedge against this risk, the supplier needs to hold more inventory; hence, $S_{M-1}^0 > S_M^0$. An exception is the case of the first two ratings, where it can be shown that $S_2^0 \geq S_1^0$ always. For $M=2$, this further implies that (2.39) holds, and therefore, the optimal policy effectively is a basestock policy. This is stated in Proposition 2.6, where we also provide expressions that lead to the computation of S_1^0 and S_2^0 .

Proposition 2.6. S_1^0 and S_2^0 satisfy

(i) $S_2^0 \geq S_1^0$ for $M \geq 2$.

(ii) If $M=2$, then, for any initial rating $\alpha = 1, 2$ and inventory level $x \leq S_\alpha^0$, the optimal stationary inventory control policy effectively is a basestock policy with rating-dependent basestock levels S_1^0 and S_2^0 satisfying the first-order conditions,

$$\frac{K_1 - K_2 q_\alpha \bar{F}(S_\alpha^0)}{\beta q_\alpha f(S_\alpha^0)} = \frac{K_3(q_2 - q_1)\theta - K_1(S_2^0 - S_1^0) - K_2[q_2 B(S_2^0) - q_1 B(S_1^0)]}{1 - \beta + \beta[q_2 \bar{F}(S_2^0) + q_1 \bar{F}(S_1^0)]} \quad (2.41)$$

for $\alpha = 1, 2$.

In the proof, we also provide conditions for the two special cases where S_1^0 or both S_1^0 and S_2^0 are zero and do not satisfy first-order conditions. Note that the r.h.s. of (2.41) is independent of α . Hence, the l.h.s. is the same for both $\alpha = 1, 2$, i.e.,

$$[K_1 - K_2 q_1 \bar{F}(S_1^0)] q_2 f(S_2^0) = [K_1 - K_2 q_2 \bar{F}(S_2^0)] q_1 f(S_1^0). \quad (2.42)$$

From (2.41) and (2.42), we can derive closed-form expressions for S_1^0 and S_2^0 for different demand distributions. For example, if w_t is exponentially distributed with mean θ , equation (2.42) yields $S_2^0 - S_1^0 = \theta \ln(q_2/q_1)$, implying that $S_2^0 - S_1^0$ depends only on the average demand and the relative selection probabilities. Substituting S_1^0 from this expression into (2.41) and solving (2.42) for $\alpha = 1$, yields $S_1^0 = \theta \ln \{ \max [q_1(K_2 + [K_3(q_2 - q_1) - K_1 \ln(q_2/q_1)]) / (1 - \beta \bar{q}_1) K_1, 1] \}$. Similarly, if w_t is uniformly distributed in $[0, 2\theta]$, equation (2.42) yields $S_2^0 - S_1^0 = 2\theta(K_1/K_2)(q_2 - q_1)/q_1 q_2$, implying that $S_2^0 - S_1^0$ depends on all problem parameters. More specifically, from (2.6) and (2.7), it is decreasing in r and b and increasing in c and h . Substituting S_2^0 from the above expression into (2.41), and solving (2.41) for $\alpha = 1$, yields a complicated expression for S_1^0 , which we omit for space considerations.

2.5 The case of constant buyer demand

In this section, we consider the case where the buyer demand is a constant $\theta > 0$. In this case, the demand seen by the supplier is still random (Bernoulli) and given by (2.1), where $w_t = \theta, \forall t$. For notational and computational simplicity, we consider the average instead of the discounted profit criterion. Specifically, the problem of the supplier is to select order-up-to levels $y_t \geq x_t, \forall t$, to maximize her average expected profit over an infinite horizon, denoted by $\tilde{\Pi}$, defined as

$$\tilde{\Pi} = \lim_{T \rightarrow \infty} \frac{1}{T} \max_{y_t \geq x_t} \left\{ \sum_{t=0}^T \Lambda_{\alpha_t}(y_t) \right\}, \quad (2.43)$$

where $\Lambda_{\alpha_t}(y_t)$ is the expected profit in period t and is given from (2.11)-(2.12), after substituting K_1, K_2 , and K_3 from (2.6)-(2.8) for $\beta = 1$, as follows:

$$\Lambda_{\alpha_t}(y_t) = (p + h)q_{\alpha_t}\theta - hy_t - (h + b) [q_{\alpha_t}B(y_t) + \bar{q}_{\alpha_t}(y_t)^-]. \quad (2.44)$$

For this problem, we have the following results.

Lemma 2.3. *If the buyer demand is a constant θ , then the optimal stationary inventory control policy, $y_{\alpha}^*(x)$ is a basestock policy with rating-dependent basestock levels S_{α}^0 satisfying*

- (i) $S_{\alpha}^0 \leq \theta, \alpha \in A$.
- (ii) If $S_{\alpha}^0 < \theta$ for some α , then all ratings $\alpha', \alpha' > \alpha$, are not accessible from α .
- (iii) If $S_{\alpha}^0 = \theta$ for some α , then all ratings $\alpha', \alpha' < \alpha$, are not accessible from α .

The proof is straightforward and is omitted. Lemma 2.3 implies that the optimal policy must be searched among the following candidate basestock policies which differ in the values of S_{α}^0 : (i) Policy $P_1 : 0 \leq S_{\alpha}^0 < \theta, \alpha \in A$. Under P_1 , the supplier's rating will eventually be absorbed in the lowest value 1, because she will never immediately satisfy the demand. (ii) Policy $P_M : S_{\alpha}^0 = \theta, \alpha \in A$. Under P_M , the supplier's rating will eventually be absorbed in the highest value M , because she will always immediately satisfy the demand. (iii) Policy $P_{\alpha-1, \alpha}, \alpha \in \{2, \dots, M\} : S_{\alpha'}^0 = \theta, \alpha' \leq \alpha - 1$ and $0 \leq S_{\alpha'}^0 < \theta, \alpha' \geq \alpha$. Under $P_{\alpha-1, \alpha}$, the supplier's rating will eventually be absorbed in the set $\{\alpha - 1, \alpha\}$, because she will always immediately satisfy the demand, when her rating is at or below $\alpha - 1$, and never immediately satisfy the demand when her rating is at or above α . Theorem 2.1 provides the conditions under which each of the candidate policies is optimal and the resulting maximum average expected profit.

Theorem 2.1. *If the buyer demand is a constant θ , the optimal basestock levels S_{α}^0 and the resulting maximum average expected profit $\tilde{\Pi}$ are given by the following table,*

$$\text{where } Q_1 = \frac{h}{b+h}, Q_3 = \frac{(p-b)q_{M-1} + \sqrt{[(p-b)q_{M-1}]^2 + 4(p+h)hq_{M-1}}}{2(p+h)},$$

$$Q_2 = \frac{h}{p+h} + \frac{p-b}{p+h}q_1, Q_4 = \frac{(p-b)q_1q_{M-1}}{(p-b)(q_{M-1} - q_1) + (p+h)q_{M-1} - h}.$$

Optimal policy	S_α^0	$\tilde{\Pi}$	Condition
P_M	θ	$[(p+h)q_M - h]\theta$	$q_M > \max[Q_2, \min(Q_1, Q_3)]$
P_1	0	$(p-b)q_1\theta$	$q_M < \min(Q_2, Q_4)$
$P_{M-1,M}$	$\theta 1_{\{\alpha \neq M\}}$	$\frac{[(2p-b+h)q_{M-1} - h]q_M\theta}{q_{M-1} + q_M}$	$Q_4 < q_M < \min(Q_1, Q_3)$

Theorem 2.1 states that the only policy that can be optimal, besides P_1 and P_M , is $P_{M-1,M}$. The conditions under which each policy is optimal have the form of inequalities involving the selection probabilities of the extreme ratings, q_1, q_{M-1} , and q_M , and are independent of $q_2, \dots, q_{M-2}$. These conditions partition the q_1 - q_{M-1} - q_M space into regions where only one policy is optimal. On the boundaries separating two regions, the policies that are optimal on either side of the boundary yield the same expected profit, hence they are both optimal.

Figure 2.2 shows three representative graphs in the q_1 - q_M space, displaying the regions where the three policies are optimal, for $p = 4.8, h = 3, b = 1$, and three different values of q_{M-1} , respectively. The regions where P_M , $P_{M-1,M}$, and P_1 are optimal are filled in dark gray, light gray, and white, respectively. Note that not all parts of these regions are feasible. The feasible parts are in the top left quadrant of each graph where $q_1 \leq q_{M-1} \leq q_M$. The other three quadrants are non-feasible and are shaded with stripes.




Figure 2.2: Optimal inventory control policy for $p = 4.8, h = 3, b = 1$ and different values of q_{M-1} .

From Theorem 2.1, it follows that if $q_M > h/(b+h)$, policy P_M outperforms the

other two policies, irrespectively of the supplier's profit margin p and the buyer's reaction to stockouts expressed by the other selection probabilities (see also the proof). This condition holds if h/b is relatively small and/or q_M is relatively large. In both cases, the expected inventory holding cost is limited. Graph (a) in Figure 2.2 represents a case where $q_{M-1} > h/(b+h)$ and therefore $q_M > h/(b+h)$. In this case, P_M outperforms the other two policies for all feasible values of q_1 and q_M ; therefore, region P_M entirely covers the feasible quadrant.

Graph (b) shows a case where $h/(2p+h-b) < q_{M-1} < h/(b+h)$. In this case, the feasible quadrant contains all three regions. If q_M is smaller than $h/(b+h)$ but relatively larger than q_{M-1} , the supplier would incur a significant profit loss if she allowed her rating to drop below M ; therefore, the overall optimal policy is P_M , irrespectively of q_1 . If q_M and q_{M-1} are close to each other but significantly larger than q_1 , $P_{M-1,M}$ is overall optimal, because it keeps a better balance between inventory and backlog costs than P_M does, without sacrificing revenues too much, since q_{M-1} is close to q_M . Finally, if q_M , q_{M-1} , and q_1 are close to each other and are not too high, P_1 is overall optimal, because it eliminates inventory holding costs, which can be quite high, given that the selection probabilities are not too high.

Graph (c) shows a case where $q_{M-1} < h/(2p+h-b)$. In this case, the feasible quadrant is covered by regions P_M and P_1 only. $P_{M-1,M}$ is never optimal, because q_{M-1} is too small and close to q_1 and is overtaken by P_1 when q_M is also small, as was explained earlier. If q_M is large, on the other hand, P_M is overall optimal.

Finally, it is straightforward to show that in the case of two ratings ($M = 2$), Theorem 2.1 reduces to Corollary 2.2.

Corollary 2.2. *If $M = 2$ and the buyer demand is a constant θ , the optimal base-stock levels S_α^0 and the resulting maximum average expected profit $\tilde{\Pi}$ are given by the following table:*

Optimal policy	S_α^0	$\tilde{\Pi}$	Condition
P_2	θ	$[(p+h)q_2 - h]\theta$	$q_2 > Q_2$
P_1	0	$(p-b)q_1\theta$	$q_2 < Q_2$

Corollary 2.2 states that when $M = 2$, the option of alternating between ratings 1 and 2, i.e., policy $P_{1,2}$, is never optimal. Indeed, under such a policy, the supplier would order up to θ when $\alpha = 1$ and order up to zero when $\alpha = 2$, which would imply that $S_2^0 < S_1^0$. However, as we have shown in Proposition 2.6 for the general case where the buyer demand is stochastic when $M = 2$, $S_2^0 \geq S_1^0$.

2.6 Imputing the fixed stockout cost newsvendor model

As was mentioned earlier, Çetinkaya and Parlar (1998) studied an extension of the newsvendor model with a fixed stockout cost in addition to the proportional backorder cost. One of the interpretations of the fixed cost is that it is a penalty for the buyer's loss of goodwill and hence future demand, due to a stockout. How to estimate this cost, however, remains questionable. To address it, we consider a similar model, which we refer to as FS, and impute the fixed stockout cost in this model by relating it to the service-driven demand (SD) model developed in Section 2.2. In the FS model, we assume that the demand seen by the supplier in each period t is given by:

$$d_t = \begin{cases} w_t, & \text{w.p. } q, \\ 0, & \text{w.p. } \bar{q}. \end{cases} \quad (2.45)$$

where w_t is the buyer demand and q is the probability that the buyer selects the supplier; w_t has the same distribution as in the SD model, but unlike in that model, q is constant and independent of past service. The per period profit is identical to that in the SD model, with the addition of a fixed penalty cost per stockout incident $\hat{b}$, and is given by $r[(x_t)^- + \min(y_t, d_t)] - c(y_t - x_t) - h(y_t - d_t)^+ - b(d_t - y_t)^+ - \hat{b}1_{\{d_t > (y_t)^+\}}$. Çetinkaya and Parlar (1998) consider the same expression for the per period profit without the term $r(x_t)^-$. Moreover, the fixed stockout cost term in their model is $\hat{b}1_{\{d_t > y_t\}}$ instead of $\hat{b}1_{\{d_t > (y_t)^+\}}$, because they assume that the buyer always selects the supplier, hence, $P(d_t = 0) = 0$. Similarly to the SD model, the redefined profit in the

FS model is $K_3d_t - K_1y_t - K_2(d_t - y_t)^+ - \hat{b}1_{\{d_t > (y_t)^+\}}$. Its expected value, $\Lambda(y_t)$, is:

$$\Lambda(y_t) = K_3q\theta - L(y_t), \quad (2.46)$$

where $L(y_t)$ denotes the expected cost of the supplier in period t , and is given by:

$$L(y_t) = K_1y_t + K_2 [qB(y_t) + \bar{q}(y_t)^-] + \hat{b}q\bar{F}(y_t). \quad (2.47)$$

Expression (2.47) is the same as (2.12) with the addition of the last term and without the dependence on the rating. Proposition 2.7 gives the optimal single-period (myopic) policy for the FS model.

Proposition 2.7. *If $f(y)$ satisfies:*

$$\frac{f'(y)}{f(y)} \leq \frac{K_2}{\hat{b}}, \quad (2.48)$$

then $\Lambda(y), y \geq 0$, is concave, and therefore the myopic inventory control policy in the FS model is a basestock policy with basestock level:

$$S^{my} = \arg \min_{y \geq 0} \left\{ K_1 - K_2q\bar{F}(y) - \hat{b}qf(y) \geq 0 \right\}. \quad (2.49)$$

Expressions (2.48) and (2.49) are similar to expressions (20) and (21) in Çetinkaya and Parlar (1998), except that there, $q = 1$ and K_2 equals $r + b + h$ instead of $(1 - \beta)r + b + h$, because of the omission of the term $r(x_t)^-$ in the per period profit, as was mentioned earlier. Condition (2.48) is also similar to condition (2.37) in Proposition 2.4, except that the latter contains $\beta\Delta_\alpha$ in place of $\hat{b}$. This implies that $\hat{b}$ can be interpreted as the supplier's maximum discounted future profit loss following a stockout. Assuming (2.48) holds, it follows from (2.49) that if $\hat{b}f(0) > K_1/q - K_2$, then S^{my} is the unique positive solution of $K_1 - K_2q\bar{F}(y) - \hat{b}qf(y) = 0$; otherwise, $S^{my} = 0$. Note that if $f(y)$ is non-increasing, (2.48) immediately holds, and S^{my} is the unique solution of (2.49). For example, if w_t is exponentially distributed with mean θ , (2.49) yields $S^{my} = \theta \ln\{\max[q(K_2 + \hat{b}/\theta)/K_1, 1]\}$. Similarly, if w_t is uniformly distributed in $[0, 2\theta]$, (2.49) yields $S^{my} = [\hat{b} + (K_2 - K_1/q)2\theta]^+/K_2$.

Note that under condition (2.48), the myopic policy in Proposition 2.7 is optimal also for the infinite-horizon problem. For the discounted expected profit criterion, this has been shown in Çetinkaya and Parlar (1998). For the average expected profit criterion, it can be shown, e.g., by using the vanishing discount method Beyer, Chang, Sethi and Taksar (2010). In the latter case, the optimal basestock level is given by (2.49) for $\beta = 1$. For computational and notational simplicity, here, we consider the infinite-horizon expected average profit criterion. In the FS model, $\hat{b}$ is supposed to reflect the cost from the loss in future demand due to the loss of goodwill following a stockout; yet, the demand is assumed to be stationary and independent of past service. Moreover, there are no guidelines on how to select $\hat{b}$. Choosing $\hat{b}$ in an arbitrary way may lead to potentially significant profit losses. To address this issue, we propose to estimate $\hat{b}$ by linking the FS model to the SD model which explicitly incorporates the buyer's response to service into the demand dynamics. Because in the FS model the supplier uses a basestock policy with basestock level S^{my} given by (2.49) (assuming (2.48) holds), we presume that the linked SD model is also operated under a basestock policy with a common basestock level S for all ratings. We refer to this model as the FS-equivalent SD model. To estimate $\hat{b}$ in the FS model, we compute the optimal basestock level S^* and the corresponding average selection probability $\tilde{q}(S^*)$ in the FS-equivalent SD model. Then, we set $S^{my} = S^*$ and $q = \tilde{q}(S^*)$ in the FS model and solve (2.49) for $\hat{b}$.

Proposition 2.8. *The optimal basestock level S^* in the FS-equivalent SD model and the imputed fixed stockout cost $\hat{b}^*$ in the FS model are given by*

$$S^* = \arg \max_{S \geq 0} \left\{ \tilde{\Pi}(S) \right\}, \quad (2.50)$$

$$\hat{b}^* = \begin{cases} \frac{1}{f(S^*)} \left(\frac{K_1}{\tilde{q}(S^*)} - K_2 \bar{F}(S^*) \right), & S^* > 0, \\ \hat{b} \in \left[0, \frac{1}{f(0)} \left(\frac{K_1}{q_1} - K_2 \right) \right], & S^* = 0, \end{cases} \quad (2.51)$$

where

$$\tilde{\Pi}(S) = \tilde{q}(S) [K_3\theta - K_2B(S)] - K_1S, \quad (2.52)$$

$$\tilde{q}(S) = \frac{1 - \Phi(S)^M}{1 - \Phi(S)} \left[\sum_{\alpha \in A} \frac{\Phi(S)^{\alpha-1}}{q_\alpha} \right]^{-1}, \quad (2.53)$$

with $\Phi(S) = F(S)/\bar{F}(S)$ and $K_1 = h, K_2 = h + b, K_3 = h + p$, since $\beta = 1$.

Proposition 2.8 gives expressions for S^* and the imputed fixed stockout cost $\hat{b}^*$, along with expressions for the average expected profit $\tilde{\Pi}(S)$ and average selection probability $\tilde{q}(S)$ in the FS-equivalent SD model. As was mentioned above, expression (2.51) is derived by solving (2.49) for $\hat{b}$, using $S^{my} = S^*$ and $q = \tilde{q}(S^*)$. If we reverse the problem and use $\hat{b} = \hat{b}^*$ from (2.51) and $q = \tilde{q}(S^*)$ from (2.53) to compute S^{my} from (2.49), the solution may not be unique. Certainly, one solution is S^* , but there may be other solutions too. A sufficient condition for S^* to be the unique solution is that $\hat{b}^*$ satisfies (2.48).

In the method described above, the supplier in the FS-equivalent SD model is restricted to operate under a basestock policy with a common basestock level for all ratings, to match the operation of the newsvendor in the FS model. This policy, besides being useful for estimating $\hat{b}^*$, is of interest in itself, because of its simplicity and ease of implementation. An obvious question is, how well does it perform compared to the optimal policy, which can be found by solving for the maximum average expected profit over an infinite horizon $\tilde{\Pi}$ given by (2.43), where $\Lambda(y_t)$ is given by (2.11) for $\beta = 1$. If its optimality gap is small, this would make it attractive for practical purposes. Moreover, it would justify the adoption of the simpler FS model, provided that the supplier uses the imputed fixed cost $\hat{b}^*$. In this case, a natural follow-up question is, how sensitive is the average expected profit to errors in $\hat{b}^*$. We address this questions in Section 2.7.3

2.7 Numerical results

In this section, we present numerical results on the bounds of the optimal policy developed in Section 2.3, on the verification of the optimal policy structure and the effect of problem parameters on the optimal policy, and on the performance of the FS-equivalent SD model developed in Section 2.6.

2.7.1 Evaluation of bounds on optimal policy

To numerically evaluate the bounds that we derived in Proposition 2.2 and compare them against the bounds in Robinson (2016), we tested 400 instances of a problem with $M = 5$. The parameters for each instance were generated randomly within the following ranges: $c \in [1, 5]$, $h, b \in [0, 1]$, $\beta \in [0.85, 0.95]$, $p = b/\beta + \Delta p$, $\Delta p \in [1, 5]$; r was computed as $p + c$. The selection probabilities $q_\alpha, \alpha \in A$, were generated as the order statistics of M random variates uniformly distributed in the interval $[0.15, 0.95]$. The buyer demand distribution was a mixture of two normal distributions with means 2.5 and 5.0, variances equal to the means, and weights $\xi_{2.5}$ and $\xi_{5.0} = 1 - \xi_{2.5}$, respectively, where $\xi_{2.5}$ was randomly generated in the interval $[0.77, 0.87]$, i.e., $f(w)$ was bimodal with mean $\theta = \xi_{2.5}2.5 + \xi_{5.0}5.0$.

For each instance, we computed $V_\alpha(0)$ by numerically solving the dynamic programming equation (2.14) using state-space discretization and value iteration. To implement our numerical scheme, we approximated the two normal distributions with two Poisson distributions with means 25 and 50, respectively, we discretized the inventory space using step size 1, truncated it in the interval $[-90, 80]$, and scaled it by a factor of 0.1. We also computed the upper and lower bounds in Robinson (2016), as well as the bounds in (2.22), for $x = 0$. For the lower bounds, $V_\alpha^L(S_\alpha)$, we used $S_\alpha = 0$ (Robinson's lower bound) and $S_\alpha = S_\alpha^{my}$. Our numerical experiments showed that when the myopic policy was used, the drop in profit was significant, ranging on average from 27.66% to 40.76%, depending on the initial rating. As was shown in Proposition 2.3, $S_\alpha^0 \geq S_\alpha^{my}$, suggesting that order-up-to points that are larger than S_α^{my} are likely to produce tighter bounds. With this in mind, we also tested $S_\alpha > S_\alpha^{my}$, for several order-up-to points S_α , such that $S_{\alpha'} \geq S_\alpha, \alpha' > \alpha$. After some experimentation, we

observed that S_α values that are computed by the following heuristic formula yield reasonably tight bounds:

$$S_\alpha^{heur} = S_M^{my} \left(1 + \frac{q_{\alpha+\delta_\alpha^+} - q_1}{q_{\alpha+\delta_\alpha^+}} \right) = S_M^{my} \left(2 - \frac{q_1}{q_{\alpha+\delta_\alpha^+}} \right). \quad (2.54)$$

According to (2.54), S_α^{heur} equals S_M^{my} plus a term which is proportional to the percent increase of the selection probability after a good service w.r.t. to the selection probability in the lowest rating.

To compare the five bounds of $V_\alpha(0)$ discussed above, we computed the percent difference $100 \times (X - V_\alpha(0))/X$ for each bound X . For notational simplicity, we denote these differences by LBR for $X = V_\alpha^L(0)$ (lower bound in Robinson (2016)), LBM for $X = V_\alpha^L(S_\alpha^{my})$, LBH for $X = V_\alpha^L(S_\alpha^{heur})$, UB for $X = V_\alpha^U(S_\alpha^{my})$, and UBR for $X =$ upper bound in Robinson (2016). Figure 2.3 shows plots of LBR, LBM, LBH, UB, and UBR, for $\alpha = 2$ and 4, for the first 200 instances. In each plot, the instances are sorted in ascending order of the LBR values, for ease of exposition.




Figure 2.3: LBR, LBM, LBH, UB, and UBR, for 200 instances.

The plots demonstrate the superiority—in terms of tightness—of UB over UBR and of LBM over LBR and show that LBH is a much tighter lower bound than LBM and LBR in almost all instances. Only in a few instances where the optimal policy is to order up to zero for some α , LBH is lower than LBM and even LBR, because LBH is constructed by considering a policy where the supplier stocks a multiple of S_M^{my} , which is positive when $S_M^{my} > 0$.

We also computed the mean values of LBR, LBM, LBH, UB, and UBR over all 400 instances. These values show that on average, LBH, UB, and UBR drop significantly with α , because they are constructed by considering policies that drive α_t towards higher values. Indicatively, the mean LBH value drops from -4.28 for $\alpha = 5$ to -9.66 , for $\alpha = 1$. LBR and LBM, on the other hand, are U-shaped in α . The mean LBM value is 3.91% higher than the mean LBR value, for $\alpha = 1$, but this difference rises to 22.38% for $\alpha = 5$.

2.7.2 Verification of optimal policy structure and effect of problem parameters on optimal policy

To verify the structure of the optimal policy, we numerically solved the dynamic programming equation (2.14) using state-space discretization and value iteration, for several problem instances. In each instance, we varied h, b, r , and the selection probabilities q_α , for which we considered the four different profiles shown in Figure 2.4, representing different buyer responses to stockouts.




Figure 2.4: Four different profiles of q_α vs. α for $M = 5$.

In all instances, $M = 5, \beta = 0.9, c = 1$, and the buyer demand distribution is a mixture of two normal distributions with means 1.0 and 5.0, variances equal to the means, and weights $\xi_{1.0} = 0.7$ and $\xi_{5.0} = 0.3$, respectively, i.e., $f(w)$ is bimodal with mean $\theta = (0.7)(1.0) + (0.3)(5.0) = 2.2$. To implement our numerical scheme, we approximated the two normal distributions with two Poisson distributions with

means 10 and 50, respectively, we discretized the inventory space using step size 1, truncated it in the interval $[-90, 80]$, and scaled it by a factor of 0.1.

Figure 2.5 shows graphs of $V_\alpha(x)$ vs. $x, x \geq 0$ for all α , for a representative instance where $h = 0.5, b = 0.4$, and $r = 3.6$. These graphs verify that the value function for each rating has the shape shown in Figure 2.1, leading to the partitioning of the inventory space in order-up-to and do-not-order regions, drawn with black and gray color, respectively. For the intensely convex q_α profile, rating 4 has two order-up-to regions, and all other ratings have one such region; for rating 1, this region is $(-\infty, 0]$, so $S_1^0 = 0$. For the convex, linear, and concave q_α profiles, ratings 1 and 5 have one order-up-to region, and ratings 2, 3, and 4 have two such regions. This verifies our intuition following equation (2.28) that the number of order-up-to regions is bounded by the number of local maxima of $f(w)$, which, for the instances tested, is 2. The




Figure 2.5: $V_\alpha(x)$ vs. $x, x \geq 0$, and optimal order-up-to and do-not-order regions for $h = 0.5, b = 0.4, r = 3.6, \beta = 0.9$, and $\theta = 2.2$, for the four q_α profiles in Figure 2.4.

graphs also show that although $V_\alpha(x)$ varies significantly with α , it is quite flat in the region between the first and last order-up-to points. This suggests that basestock policies with rating-dependent basestock levels may perform quite well. In fact, the policy used to construct $V_\alpha^L(S_\alpha)$ with $S_\alpha = S_\alpha^{heur}$, where S_α^{heur} is given by (2.54), gives

reasonably tight bounds, as was discussed in Section 2.7.1. Also, our numerical results in Section 2.7.3 indicate that using a basestock policy with a common basestock level for all ratings can be quite efficient.

To explore the effect of problem parameters on the optimal policy, we numerically solved the dynamic programming equation (2.14) using state-space discretization and value iteration, for a large number of problem instances, with $\beta = 0.9$ and $c = 1$, where we varied h , b , r , M , q_α , and the buyer demand distribution. We run two sets of experiments. In the first set, we considered the values $h \in \{0.2, 0.3, 0.4, 0.5\}$, $b \in \{0.7, 0.8, 0.9\}$, $r \in \{3.3, 3.4, 3.5\}$, $M = 5$, and the four profiles of q_α in Figure 2.4. In the second set, we considered the same values for h , b , and r as in the first set, and in addition the values $M \in \{2, 3, 5, 7\}$, for a linear q_α profile ranging between $q_1 = 0.2$ and $q_M = 0.9$. In both sets of experiments, we considered exponential and normal buyer demand distributions with mean $\theta = 5.0$ and, in the case of the normal distribution, variance equal to the mean. Note that the coefficient of variation of the exponential distribution is 1, whereas that of the normal distribution is $1/\sqrt{\theta} = 1/\sqrt{5} = 0.447$. To implement our numerical scheme, we approximated the exponential and normal distributions with geometric and Poisson distributions, respectively, with mean 50, we discretized the inventory space using step size 1, truncated it in the interval $[-550, 200]$ for the geometric case, and $[-90, 80]$ for the Poisson case, and scaled it by a factor of 0.1. In all instances, the optimal policy is basestock with rating-dependent basestock levels S_α^0 . Figure 2.6 shows indicative plots of S_α^0 vs. α for $b = 0.8$, $r = 3.4$, $M = 5$, and all the h values, q_α profiles, and demand distributions tested, for the first set of experiments. The plots show that for the intensely convex q_α profile, S_α^0 is increasing and concave in α ; the concavity breaks only in two instances where $S_1^0 = 0$, for the exponential distribution. For the other three q_α profiles, S_α^0 has a skewed inverted U shape as a function of α . As we move from the intensely convex to the concave q_α profile, S_α^0 increases for lower values of α and decreases for higher values, and the skewness of S_α^0 shifts from higher to lower values of α , reflecting the corresponding shift in the elasticity of q_α . Note however that in all instances $S_1^0 \leq S_2^0$, verifying Proposition 2.6 (i). These results suggest that the supplier tends to maintain higher inventory in intermediate ratings, where




Figure 2.6: S_α^0 vs. α for $b = 0.8$, $r = 3.4$, red $M = 5$, four values of h , two demand distributions, and the four q_α profiles in Figure 2.4.

she has to gain if she meets the demand and lose if she does not, than in the lowest and highest ratings, where she has nothing to lose and nothing to gain, respectively. Even under the convex q_α profile, where q_α is increasingly more elastic in α , S_α^0 has an inverted U shape instead of being increasing in α . If the buyer's intention is to create a supplier rating system to improve service, then the probability with which he selects the supplier must be sharply increasing in α , as is the case with the intensely convex q_α profile.

The plots also show that the optimal basestock level for the lowest rating is significantly lower for the exponential distribution than it is for the normal distribution and that this ordering is sharply reversed for larger ratings. This is due to the fact that the two distributions differ both in shape and variability. In the lowest rating, under the exponential distribution, the supplier has a good chance of meeting the demand and increasing her rating even if her basestock level is lower than the mean demand; under the normal distribution, her basestock level must be closer to the

mean demand to match this chance. On the other hand, for higher ratings, her risk of stocking out and being downgraded is higher under the exponential distribution, than it is under the normal distribution because the former distribution has higher variability. This results in significantly higher optimal basestock levels for the exponential distribution. In fact, the span of S_α^0 values for the exponential distribution is an order of magnitude larger than that for the normal distribution. These results suggest that the shape and, most importantly, the variability of the buyer's demand dramatically amplifies the effect of the problem parameters on the optimal basestock levels of the supplier.

As expected, in all instances, S_α^0 is decreasing in h . Similar results were observed when we independently varied b and r , except that S_α^0 is increasing in both b and r , although its sensitivity to b is quite low.

Figure 2.7 shows indicative plots of S_α^0 vs. q_α for $b = 0.8$, $r = 3.4$, a linear q_α profile, two h values, and all the M values and demand distributions tested, for the second set of experiments.




Figure 2.7: S_α^0 vs. q_α for $b = 0.8$, $r = 3.4$, the linear q_α profile, two values of h , four values of M , and two demand distributions.

The plots show that S_α^0 has the same skewed inverted U shape that we saw in the first set of experiments, for all values of M , except $M = 2$, where $S_2^0 > S_1^0$. They also reconfirm that the variability of the buyer's demand dramatically amplifies the effect of the problem parameters on the optimal basestock levels of the supplier. More importantly, the plots reveal that the higher the value of M , the lower the supplier's

S_α^0 profile. This is because when M is large, the buyer reacts less erratically to bad and good service, i.e., with smaller swings in q_α , allowing the supplier to reduce her basestock levels. Therefore, from the buyer's perspective, a more erratic response induces better service.

2.7.3 Performance evaluation of FS-equivalent SD model

To numerically assess the performance of the FS-equivalent SD model and the sensitivity of the average expected profit to errors in $\hat{b}^*$, we tested 60 instances for a problem with $M = 5$. In each instance, $\beta = 1$ and the revenue and cost parameters were generated randomly as in the numerical study in Section 2.7.1. In all instances, the buyer demand distribution was exponential with mean $\theta = 3$, hence condition (2.48) immediately holds. All instances were repeated for six different profiles of q_α vs. α , shown in Figure 2.8, representing different buyer responses to stockouts, raising the total number of instances tested to 360 ($= 60 \times 6$). Note that profile 0 corresponds to the newsvendor model, where $q_\alpha = q, \alpha \in A$, and hence the demand is independent of the rating. For each instance, we computed S^* , $\tilde{\Pi}(S^*)$, and $\tilde{q}(S^*)$ from (2.50), (2.52), and (2.53), respectively, in the FS-equivalent SD model, and the imputed fixed cost $\hat{b}^*$ in the FS model from (2.51).




Figure 2.8: Six different profiles of q_α vs. α for $M = 5$.

Figure 2.9 shows plots of S^* , $\tilde{q}(S^*)$, $\tilde{\Pi}(S^*)$, and $\hat{b}^*$, for the 60 instances and six buyer response profiles tested. The instances in each plot are sorted in ascending order of the plotted values corresponding to the linear profile (profile 3), for ease of exposition. Not surprisingly, the rating-dependent buyer response profiles that yield

the highest $\tilde{\Pi}(S^*)$ values are those with the highest q_α values in high ratings (profiles 5, 4, 3, 2, and 1, in decreasing order). These same profiles result in the highest $\tilde{q}(S^*)$ and smallest S^* values. As a result, they yield the smallest imputed fixed cost $\hat{b}^*$. For all these profiles, $\tilde{q}(S^*)$ seems to converge to a value in the interval $[0.81, 0.86]$. In almost all instances, profile 0 yields the lowest $\tilde{\Pi}(S^*)$ value, because its constant selection probability is relatively low in high ratings. The few instances where profile 0 results in a higher $\tilde{\Pi}(S^*)$ value than other profiles do, are characterized by high h/p values. When h/p is high, it is optimal not to hold inventory, driving the supplier's rating downwards. In low ratings, profile 0 has the advantage of a higher selection probability compared to other profiles. Note that for profile 0, $S^* = 0$ in more than one-third of the instances. In these instances, from (2.50), $\hat{b}^* \in [0, (K_1/q_1 - K_2)/f(0)]$. In Figure 2.9, we plotted the upper bound of this interval.




Figure 2.9: $\tilde{\Pi}(S^*)$, S^* , $\tilde{q}(S^*)$, and $\hat{b}^*$, for the six q_α profiles in Figure 2.6.

To assess the performance of the FS-equivalent SD policy, we compared $\tilde{\Pi}(S^*)$ against the maximum expected average profit under the optimal policy in the SD model, $\tilde{\Pi}$, which was computed by numerically solving the corresponding dynamic programming equation. To examine the sensitivity of the average expected profit to

errors in $\hat{b}^*$, we computed the optimal basestock level in the FS model for values of $\hat{b} \neq \hat{b}^*$ from (2.49), denoted by $S^{my}(\hat{b})$, after substituting $q = \tilde{q}(S^{my}(\hat{b}))$ from (2.53). Then, we substituted $S^{my}(\hat{b})$ in (2.52) to compute the resulting average expected profit in the FS-equivalent SD model, $\tilde{\Pi}(S^{my}(\hat{b}))$. The resulting percent loss in the average expected profit is denoted by $\Delta\tilde{\Pi}(\hat{b})$, i.e., $\Delta\tilde{\Pi}(\hat{b}) = 100 \times [\tilde{\Pi} - \tilde{\Pi}(S^{my}(\hat{b}))]/\tilde{\Pi}$.

Figure 2.10 (left) shows $\Delta\tilde{\Pi}(\hat{b})$, for $\hat{b} = m\hat{b}^*$, for different multiplication factors m between 0 and 4, for profile 3, for the 60 instances. Figure 2.10 (right) shows $\Delta\tilde{\Pi}(\hat{b}^*)$ for all profiles. The instances in both plots are sorted in ascending order of $\Delta\tilde{\Pi}(\hat{b}^*)$ for the linear profile (profile 3), for ease of exposition.




Figure 2.10: Left: $\Delta\tilde{\Pi}(\hat{b})$, for $\hat{b} = m\hat{b}^*$, $m = 0, 0.5, 1, 2, 4$, for q_α profile 3 in Figure 2.10. Right: $\Delta\tilde{\Pi}(\hat{b}^*)$, for q_α profiles 1–5 in Figure 2.10.

From Figure 2.10 (left), we see that the smallest loss (less than 2% on average) is obtained when the supplier uses the imputed fixed cost, i.e., $\hat{b} = \hat{b}^*$. If $\hat{b} = 2\hat{b}^*$, the percent loss is on average approximately 7%, indicating that the average expected profit is not sensitive to overestimations of the fixed cost. If $\hat{b} = 0.5\hat{b}^*$, however, the percent loss rises to approximately 34% on average, because in one-third of the instances, the supplier ends up operating under a make-to-order policy, when such a policy may be far from optimal. Similar results hold for the other buyer response profiles. Figure 2.10 (right) indicates that the profiles that yielded the highest $\tilde{\Pi}(S^*)$ values (see Figure 2.9) more or less have the highest $\Delta\tilde{\Pi}(\hat{b}^*)$ values.

2.8 Discussion and future research

In this chapter, we developed a multiperiod model of a supplier selling items to a buyer who rates the supplier based on the history of her service, measured in terms of in-stock/out-of-stock incidents. There are several possible directions for future work. Our results were derived for continuous buyer demand (except for the case of deterministic demand), but most of them can potentially be extended to discrete demand.

We assume that M and $q_\alpha, \alpha \in A$, are known. In a real application, these parameters must be estimated from longitudinal data. If the buyer's rating of the supplier is real and observable by the supplier, then M is known, and q_α can be estimated as the number of periods that the buyer selects the supplier when her rating is α over the number of periods that the supplier's rating is α . If the rating is an imaginary construct for capturing the buyer's goodwill, then the only information that the supplier observes in each period is whether the buyer selects her or not and if he does, whether she meets the demand or not. In this case, the problem of choosing the appropriate M is a model order determination problem, for which there exist established statistical (e.g., likelihood, Bayesian), information-theoretic (e.g., AIC and BIC), and machine learning (e.g., cross-validation) solution methods Singer, Helic, Taraghi and Strohmaier (2014).

We assume a single buyer with multiple satisfaction levels reflected by the ratings. A more general model can include multiple non-homogeneous buyers with different demand distributions and selection probabilities, expressing diversity in buyer needs and responses to service. In such a setting, the supplier must decide not only how much to order but how to ration inventory in case of excess demand. Judging from the work of Adelman and Mersereau (2013), who addressed the rationing but not the ordering issue in a similar setting where goodwill is modeled as an exponential smoothing of utilities derived from past fill rates, this is a very challenging problem.

Chapter 3

Dynamic Supplier Competition and Cooperation for Buyer Loyalty on Service

3.1 Introduction

In this chapter, we focus on the switching behavior of a buyer from one supplier to another following poor service and its implication on the suppliers' competitive inventory policy, in a B2B setting. In Section 3.2, we formulate the model of the buyer and the two suppliers. In Section 3.3, we discuss the myopic policy of the suppliers and characterize their long-run optimal policy. In Section 3.4, we derive properties of their best response functions under competition and discuss their equilibria. In Section 3.5, we derive properties of the optimal joint inventory policy of the suppliers when they cooperate, and we estimate the backorder penalty rate that the buyer must charge the suppliers to recover the fill rate that she enjoys under competition. In Section 3.6, we apply the results to the case where the buyer demand is exponentially distributed, and we illustrate the results with a numerical example to investigate the effect of the suppliers' parameters on the optimal outcome. Finally, in Section 3.7, we discuss the extension of our model to multiple suppliers. We summarize our findings in Section 3.8 where we propose directions for future work. Supplemental material

for this chapter, including proofs, can be found in Appendix B.

3.2 Model formulation

Two suppliers sell the same product to a buyer in consecutive periods. The buyer arranges the suppliers in a rank order list based on the last service she received, which is either satisfactory or unsatisfactory, depending on whether her demand was fully met or not. Throughout this paper, we reserve index i to denote one supplier and j to denote the other, i.e., $j \neq i$; therefore, either $(i, j) = (1, 2)$ or $(i, j) = (2, 1)$.

At the beginning of period t , each supplier i orders a non-negative quantity ahead of demand, based on his inventory level, $x_{i,t} \in \mathbb{R}$, and his placement or ranking in the buyer's list, $\alpha_{i,t} \in \{1, 2\}$, where 1 indicates the top of the list or high ranking and 2 indicates the bottom or low ranking. The order arrives before the end of the period, raising the supplier's inventory level to $y_{i,t} \geq x_{i,t}$.

At the end of the period, the buyer selects the high-ranking supplier and demands from him a random quantity w_t . If the supplier meets all the demand at once, he is kept at the top of the list and carries any leftover inventory to the next period. If he fails to meet all the demand at once, the buyer backorders the unmet demand with him and moves him to the bottom of the list, thereby bringing his competitor to the top. In other words, the buyer rewards her suppliers with loyalty if they serve her well but punishes them by switching at the first service failure. We refer to the streak of periods during which the buyer selects the same supplier before she switches to the other supplier as a *supply run*.

The demands $\{w_t, t = 0, 1, \dots\}$ are based on the buyer's needs and are independent of the suppliers' past service. We assume that they are i.i.d. continuous random variables with p.d.f., c.d.f., and mean, $f(\cdot)$, $F(\cdot)$, and θ , respectively. Based on the above assumptions, the demand seen by supplier i in period t is $w_t 1_{\{\alpha_{i,t}=1\}}$, where $1_{\{\cdot\}}$ is the indicator function.

The idea that a stockout incident has a fixed adverse impact on the standing of the supplier who runs out of stock, irrespective of the shortage quantity or time, has been addressed in the literature for the most part by considering a fixed cost per stockout

occasion or a minimum type-I service level constraint. One of the interpretations of the fixed cost is that it is a penalty for the buyer's loss of goodwill and hence future demand, due to the stockout Çetinkaya and Parlar (1998). Yet, almost always, it is assumed that the demand is unaffected by the stockout. In our model, the effect of a stockout on the supplier's demand is the direct consequence of the buyer's carrot-and-stick selection policy which signals the buyer's discontent about the stockout and stimulates competition between the suppliers. Moreover, we assume that the suppliers know their ranking before they order so that they know what to expect; hence, we are in a full information setting.

Based on the above assumptions, supplier i 's inventory level and ranking are updated as follows:

$$x_{i,t+1} = y_{i,t} - w_t 1_{\{\alpha_{i,t}=1\}}, \quad (3.1)$$

$$\alpha_{i,t+1} = \alpha_{i,t} + 1_{\{\alpha_{i,t}=1, y_{i,t} < w_t\}} - 1_{\{\alpha_{i,t}=2, y_{j,t} < w_t\}}. \quad (3.2)$$

In each period, supplier i incurs an acquisition cost c_i per item ordered and receives a revenue (price) r_i per item sold. The quantity sold is $\min(y_{i,t}, w_t 1_{\{\alpha_{i,t}=1\}})$. We also assume that he incurs an inventory cost of h_i per item in inventory and a backorder cost of b_i per item short at the end of the period. Typically, b_i is a transaction or some other friction cost for managing the backorder. To ensure that the supplier can be profitable even with backorders, we also assume that $p_i > b_i$, where p_i is the per-unit profit margin defined as $p_i = r_i - c_i$.

The profit of supplier i in period t is $r_i[(x_{i,t})^- + \min(y_{i,t}, w_t 1_{\{\alpha_{i,t}=1\}})] - c_i(y_{i,t} - x_{i,t}) - h_i(y_{i,t} - w_t 1_{\{\alpha_{i,t}=1\}})^+ - b_i(w_t 1_{\{\alpha_{i,t}=1\}} - y_{i,t})^+$, where we use the notation: $(x)^+ = \max(x, 0)$ and $(x)^- = (-x)^+$, $x \in \mathbb{R}$. After rolling the $x_{i,t}$ terms backwards for one period, similarly to Liberopoulos and Deligiannis (2022), the profit in period t can be recast as the following function of $y_{i,t}$:

$$g_i(\alpha_{i,t}, y_{i,t}, w_t) = p_i w_t 1_{\{\alpha_{i,t}=1\}} - h_i(y_{i,t} - w_t 1_{\{\alpha_{i,t}=1\}})^+ - b_i(w_t 1_{\{\alpha_{i,t}=1\}} - y_{i,t})^+. \quad (3.3)$$

Given the suppliers' decisions $y_{i,t}, y_{j,t}, t = 0, 1, \dots$, the expected average profit of

supplier i is

$$\Pi_i(y_{i,0}, y_{j,0}, y_{i,1}, y_{j,1}, \dots) = \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\sum_{t=0}^{T-1} g_i(\alpha_{i,t}, y_{i,t}, w_t) \right],$$

where the dependence of Π_i on $y_{j,t}$ stems from the dependence of $\alpha_{i,t+1}$ on $y_{j,t}$ from (3.2).

3.3 The suppliers' optimal policy and payoff

Before analyzing the model of the two suppliers, it is worth noting that in the absence of supplier j , supplier i will behave as a single multi-period newsvendor with backorders, whose optimal ordering policy is a basestock policy with basestock level s_i , and whose period profit is $g_i(1, s_i, w)$. His expected average profit, as a function of s_i , denoted by $G_i(s_i) = E[g_i(1, s_i, w)]$, and its first two derivatives are given by

$$G_i(s_i) = p_i\theta - h_i E[(s_i - w)^+] - b_i E[(w - s_i)^+], \quad (3.4)$$

$$G'_i(s_i) = -h_i F(s_i) + b_i \bar{F}(s_i), \quad (3.5)$$

$$G''_i(s_i) = -(h_i + b_i)f(s_i). \quad (3.6)$$

From (3.6), $G_i(s_i)$ is concave, so the optimal basestock level of the newsvendor, denoted by s_i^m , is the solution of the first-order condition, $G'_i(s_i) = 0$, given by the well-known critical fractile formula,

$$s_i^m = F^{-1} \left(\frac{b_i}{h_i + b_i} \right). \quad (3.7)$$

We refer to s_i^m as the *myopic* basestock level of supplier i because it maximizes the single-period expected average profit $G_i(s_i)$. From (3.4) and (3.5), $G_i(0) = (p_i - b_i)\theta > 0$, $G'_i(0) = b_i > 0$, and $\lim_{s_i \rightarrow \infty} G_i(s_i) = -\infty$, implying that $G_i(s_i)$ crosses zero and

becomes negative at a finite point, denoted by s_i^M , satisfying

$$G_i(s_i^M) = 0, \quad (3.8)$$

such that $G(s_i) > 0, s_i < s_i^M$, and $G(s_i) < 0, s_i > s_i^M$. This means that for basestock levels larger than s_i^M , the newsvendor incurs losses.

Going back to the dual-sourcing model, the structure of the optimal policy for each supplier and the resulting expected average profit over an infinite horizon is given by the following proposition.

Theorem 3.1. *The optimal ordering policy of supplier i is a ranking-dependent base-stock policy, denoted by $y_i^*(\alpha_i)$, given by*

$$y_i^*(2) = 0 \text{ and } y_i^*(1) = s_i \geq 0. \quad (3.9)$$

Under this policy, the expected average profit (payoff) of supplier i , as a function of s_i and s_j , denoted by $\Pi_i(s_i, s_j)$, is

$$\Pi_i(s_i, s_j) = \pi_i(s_i, s_j)G_i(s_i), \quad (3.10)$$

where

$$\pi_i(s_i, s_j) = \frac{\bar{F}(s_j)}{\bar{F}(s_j) + \bar{F}(s_i)}, \quad (3.11)$$

and $G_i(s_i)$ is given by (3.4).

We refer to s_i and s_j as the *active* basestock levels of supplier i and j , respectively, because the suppliers use them when they are “active”, i.e., when they are at the top of the buyer’s list, enjoying her loyalty. Figure 3.1 shows a sample trajectory of the suppliers’ inventory levels under the optimal ordering policy. The buyer switches suppliers after every supply run. If we join together the segments of supplier i ’s inventory trajectory when $\alpha_i = 1$, i.e., during his supply runs where he is active, ignoring the segments when $\alpha_i = 2$, the resulting trajectory coincides with that of a multi-period newsvendor with backorders who in every period orders up to s_i . His expected average profit in this case is $G_i(s_i)$ given by (3.4). In the remaining segments




Figure 3.1: Sample trajectory of the suppliers' inventory levels.

of the trajectory when $\alpha_i = 2$, i.e., during supplier j 's runs, supplier i is inactive and has zero inventory and zero profits. Therefore, his overall payoff is $G_i(s_i)$ weighted by the fraction of time that $\alpha_i = 1$, denoted by $\pi_i(s_i, s_j)$ and given by (3.11). This fraction represents the expected average demand share of supplier i .

The expected average fill rate seen by the buyer is denoted by $q(s_i, s_j)$ and is given by $q(s_i, s_j) = \pi_i(s_i, s_j)F_i(s_i) + \pi_j(s_i, s_j)F_j(s_j)$. From (3.11), this can be written as

$$q(s_i, s_j) = 1 - \frac{2\bar{F}(s_j)\bar{F}(s_i)}{\bar{F}(s_j) + \bar{F}(s_i)}. \quad (3.12)$$

3.4 Supplier competition

If the suppliers compete for the buyer's patronage, the problem of each supplier i is to choose an active basestock level s_i that maximizes $\Pi_i(s_i, s_j)$ defined in (3.10).

From (3.10) and (3.11), the first partial derivative of $\Pi_i(s_i, s_j)$ with respect to s_i is

$$\frac{\partial \Pi_i(s_i, s_j)}{\partial s_i} = \frac{\bar{F}(s_j)}{(\bar{F}(s_j) + \bar{F}(s_i))^2} \phi_i(s_i, s_j), \quad (3.13)$$

where $\phi_i(s_i, s_j)$ and its first partial derivatives are given by

$$\phi_i(s_i, s_j) = (\bar{F}(s_j) + \bar{F}(s_i)) G'_i(s_i) + f(s_i) G_i(s_i), \quad (3.14)$$

$$\frac{\partial \phi_i(s_i, s_j)}{\partial s_i} = (\bar{F}(s_j) + \bar{F}(s_i)) G''_i(s_i) + f'(s_i) G_i(s_i), \quad (3.15)$$

$$\frac{\partial \phi_i(s_i, s_j)}{\partial s_j} = -f(s_j) G'_i(s_i). \quad (3.16)$$

3.4.1 Best response function

From (3.10), the optimal value of s_i that maximizes the payoff of supplier i , $\Pi_i(s_i, s_j)$, depends on s_j . Let $s_i^*(s_j)$ denote the optimal active basestock level of supplier i given s_j , henceforth referred to as the best response (function) of supplier i . The following proposition provides upper and lower bounds on $s_i^*(s_j)$.

Proposition 3.1. *The best response $s_i^*(s_j)$ is bounded as follows:*

$$0 < s_i^m < s_i^*(s_j) < s_i^M, \quad s_j \in [0, \infty), \quad (3.17)$$

where s_i^m and s_i^M satisfy (3.7) and (3.8).

Proposition 3.1 states that the best response of supplier i is higher than his myopic basestock level s_i^m . By setting s_i above s_i^m , the supplier compromises part of his expected myopic profit $G(s_i)$ to extend his stay at the top of the buyer's list, thus increasing his long-term average demand share $\pi_i(s_i, s_j)$ and the resulting profits. This means that using basestock level s_i^m , although myopically optimal, will lead to payoff losses in the long run.

Given that $\partial \Pi_i(s_i, s_j)/\partial s_i > 0$, for $0 \leq s_i \leq s_i^m$, and $\partial \Pi_i(s_i, s_j)/\partial s_i < 0$, for $s_i \geq s_i^M$, the first-order condition $\partial \Pi_i(s_i, s_j)/\partial s_i = 0$ has at least one solution in (s_i^m, s_i^M) .

The solution that corresponds to the global maximizer of $\Pi_i(s_i, s_j)$ is the best response $s_i^*(s_j)$. From (3.4), (3.5), and (3.13), note that all the solutions including $s_i^*(s_j)$ depend on $F(\cdot)$, h_i , and b_i , as does the myopic basestock level s_i^m . In addition, they depend on p_i and s_j . The bounds on $s_i^*(s_j)$ given by (3.17), however, are independent of s_j . The following proposition provides a condition under which the best response $s_i^*(s_j)$ is unique.

Theorem 3.2. *If the following condition holds:*

$$\frac{\partial \phi_i(s_i, s_j)}{\partial s_i} < 0, \quad s_i \in (s_i^m, s_i^M), \quad s_j \in [0, \infty), \quad (3.18)$$

the best response $s_i^*(s_j)$ is

- (i) *a global maximizer of the payoff $\Pi_i(s_i, s_j)$ uniquely satisfying $\partial \Pi_i(s_i, s_j)/\partial s_i = 0$, which reduces to*

$$\phi_i(s_i^*(s_j), s_j) = 0, \quad (3.19)$$

- (ii) *increasing in s_j , and its derivative with respect to s_j is*

$$\frac{\partial s_i^*(s_j)}{\partial s_j} = - \frac{\partial \phi_i(s_i^*(s_j), s_j)/\partial s_j}{\partial \phi_i(s_i^*(s_j), s_j)/\partial s_i}, \quad (3.20)$$

- (iii) *increasing in θ , p_i , b_i , and decreasing in h_i .*

Theorem 3.2 (i) states that under condition (3.18), $\Pi_i(s_i, s_j)$ has a unique maximum, guaranteeing the uniqueness of the best response. Condition (3.18) is very mild and is easily satisfied. From (3.15), the first term in $\partial \phi_i(s_i, s_j)/\partial s_i$ involving $G_i''(s_i)$ is negative by (3.6). If the demand density $f(w)$ is non-increasing, as is the case with the exponential distribution, then the second term is non-positive, and the condition is met.

A more careful look at (3.15) reveals that in order for the second term to be non-positive, $f(s_i)$ does not have to be decreasing for all $s_i > 0$ but only for $s_i \in (s_i^m, s_i^M)$, because we know from (3.17) that $s_i^*(s_j) \in (s_i^m, s_i^M)$. For example, if $f(w)$ is unimodal with mode ν (i.e., ν is the maximizer of $f(w)$, above which $f(w)$ is

decreasing) and $s_i^m \geq \nu$, then $f'(s_i) \leq 0, s_i \geq s_i^m$. The condition $s_i^m \geq \nu$ holds for many unimodal distributions, for reasonable values of the newsvendor critical fractile $b_i/(h_i + b_i)$. Indicatively, if $b_i/(h_i + b_i) \geq 0.5$, then $s_i^m \geq \mu_{1/2}$, where $\mu_{1/2}$ denotes the median of $f(w)$. In this case, if $\mu_{1/2} \geq \nu$, then $s_i^m \geq \nu$. The inequality $\mu_{1/2} \geq \nu$ is satisfied for many common distributions with non-negative skewness, such as the Normal, Lognormal, Weibull, Gamma, and other distributions. For instance, if $w \sim \text{Normal}(\theta, \sigma^2)$, then $\mu_{1/2} = \nu = \theta$. If $w \sim \text{Lognormal}(\mu, \sigma^2)$, then $\mu_{1/2} = \exp(\mu) > \exp(\mu - \sigma^2) = \nu$. If $w \sim \text{Weibull}(\lambda, m)$, with $m < 1/(1 - \ln 2) \approx 3.2589$, then $\mu_{1/2} = \lambda(\ln 2)^{1/m} > \lambda[(m - 1)/m]^{1/m} = \nu$. If $w \sim \text{Gamma}(m, \xi)$, with $m \geq 1$, then $\mu_{1/2} \in ((m - 1/3)\eta, m\xi) > (m - 1)\xi = \nu$ Chen and Rubin (1986).

Even if $f(w)$ is increasing in all or parts of the interval (s_i^m, s_i^M) , condition (3.18) will still hold if

$$f'(s_i) < -(\bar{F}(s_j) + \bar{F}(s_i)) \frac{G_i''(s_i)}{G_i(s_i)}, \quad s_i \in (s_i^m, s_i^M), \quad s_j \in [0, \infty), \quad (3.21)$$

where the right-hand side of the above inequality is positive.

Finally, note that (3.18) is a sufficient and not a necessary condition, that is, $s_i^*(s_j)$ may be unique even if (3.18) does not hold. More specifically, (3.18) implies that $\phi_i(s_i, s_j)$ is strictly decreasing in (s_i^m, s_i^M) , which guarantees that it will cross zero at exactly one point. It is possible, however, that $\phi_i(s_i, s_j)$ crosses zero at exactly one point without being decreasing everywhere in (s_i^m, s_i^M) . In this case, $s_i^*(s_j)$ will still be unique.

If the first-order condition $\partial \Pi_i(s_i, s_j)/\partial s_i = 0$ does not have a unique solution, then each solution $s_i^*(s_j)$ is either a local extremum or an inflection point. In this case, one can always evaluate $\Pi_i(s_i, s_j)$ at each solution to determine the maximizer of the payoff.

Theorem 3.2 (ii) and (iii) provide important monotonicity properties of $s_i^*(s_j)$. Property (iii) is the same as the respective property in the model of a single multi-period newsvendor with backorders. Property (ii) states that $s_i^*(s_j)$ is increasing in

s_j and can be expressed as

$$s_i^*(s_j) = s_i^*(0) + \int_0^{s_j} \frac{\partial s_i^*(y)}{\partial y} dy, \quad (3.22)$$

where $s_i^*(0)$ is the solution of equation (3.14), for $s_j = 0$, and $\partial s_i^*(y)/\partial y$ is given by (3.20) for $s_j = y$. This implies that if supplier j increases his active basestock level, supplier i will follow suit to mitigate his loss of demand share.

3.4.2 Nash equilibrium

From the previous discussion, competition pushes both suppliers to move away from their myopic basestock levels in an escalating inventory contest, benefiting the buyer. Does this rivalry ever settle? The following theorem suggests that it does.

Theorem 3.3. *If condition (3.18) holds for $i = 1, 2$, then*

- (i) *There exists at least one pure-strategy Nash equilibrium (s_i^e, s_j^e) satisfying (3.19) for $i = 1, 2$.*
- (ii) *Each Nash equilibrium (s_i^e, s_j^e) is increasing in $\theta, p_i, p_j, b_i, b_j$ and decreasing in h_i, h_j .*
- (iii) *If the following condition holds:*

$$\frac{\partial s_i^*(s_j)}{\partial s_j} < 1, \quad s_j \in (s_j^m, s_j^M), \quad i = 1, 2, \quad (3.23)$$

the Nash equilibrium is unique.

Figure 3.2 shows indicative graphs of the best response functions of the two suppliers under condition (3.18) which guarantees their uniqueness. As both graphs show, the two functions cross each other at least at one point because both are increasing and bounded from above and below. The crossing points are the Nash equilibria that always belong in region B , i.e., $s_i^e \in (s_i^m, s_i^M)$, $i = 1, 2$. In graph (b), condition (3.23) holds, and therefore the Nash equilibrium is unique. Theorem 3.3 (ii) implies




Figure 3.2: Best response functions of the two suppliers and Nash equilibrium (s_i^e, s_j^e) .

that if one of supplier j 's parameter values changes, not only does he respond by changing his active basestock level, but supplier i also changes his active basestock level in response to j 's response. Condition (3.23) is necessary and sufficient for the best response to being a contraction mapping, implying the uniqueness of the Nash equilibrium.

The fact that the suppliers' active basestock levels at equilibrium are higher than their myopic levels implies that the buyer's carrot-and-stick behavior is successful in raising the fill rate that she enjoys under supplier competition. The reduction in the frequency of stockouts resulting from the increase in the basestock levels limits the role of the backorder cost rates b_i for the suppliers.

If the suppliers are symmetric (identical), i.e., $p_i = p$, $h_i = h$, $b_i = b$, for $i = 1, 2$, implying that $G_i(y) = G(y)$, for $i = 1, 2$, the Nash equilibrium is unique and is given by the following proposition.

Proposition 3.2. *If the suppliers are symmetric and condition (3.18) holds, then*

- (i) *There exists at least one symmetric pure-strategy Nash equilibrium (s^e, s^e) , where s^e satisfies*

$$\frac{f(s^e)}{F(s^e)} = -\frac{2G'(s^e)}{G(s^e)}. \quad (3.24)$$

The resulting payoff of each supplier i is

$$\Pi_i(s^e, s^e) = \frac{G(s^e)}{2}, \quad i = 1, 2. \quad (3.25)$$

(ii) *There exist no asymmetric pure-strategy Nash equilibria.*

(iii) *If the following condition holds:*

$$\frac{\partial \hat{\phi}(s)}{\partial s} < 0, \quad s \in (s^m, s^M), \quad (3.26)$$

where $\hat{\phi}(s) = \phi(s, s)$, then the symmetric equilibrium is unique.

Note that if the suppliers are symmetric, condition (3.23) is not needed for the uniqueness of the Nash equilibrium, because the first-order conditions for the two suppliers reduce to one equation (because of symmetry) which has a unique solution, under (3.18).

3.5 Supplier cooperation

In the previous section, we saw that the non-cooperative game of the suppliers forces them to increase their active basestock levels above their myopic levels, compromising their profits. Now, suppose that the suppliers decide to team up to reduce their total inventory costs, perhaps as a result of consolidation (merger or acquisition) or an agreement to split the benefits from this reduction. What is the optimal joint inventory policy and gain of the suppliers in this case, and what is the adverse impact of their cooperation on the buyer's fill rate? Moreover, what can the buyer do to recover the fill rate that she enjoyed under competition? In this section, we address these questions.

3.5.1 The suppliers' gain

If the suppliers team up, the optimal ordering policy of each supplier in the team has the same structure as that under competition, given by Theorem 3.1. Moreover, the team's payoff is the sum of the individual payoffs of the suppliers. The problem for the team is to choose an active basestock level pair (s_i, s_j) that maximizes the team's payoff, denoted by $\Pi(s_i, s_j)$, by carefully balancing the expected period profits

and the long-term average demand share of each supplier defined in (3.10). More specifically, $\Pi(s_i, s_j)$ and its first partial derivative with respect to s_i are

$$\Pi(s_i, s_j) = \Pi_i(s_i, s_j) + \Pi_j(s_i, s_j) = \pi_i(s_i, s_j)G_i(s_i) + \pi_j(s_i, s_j)G_j(s_j). \quad (3.27)$$

$$\frac{\partial \Pi(s_i, s_j)}{\partial s_i} = \frac{\bar{F}(s_j)}{(\bar{F}(s_j) + \bar{F}(s_i))^2} \psi_i(s_i, s_j), \quad (3.28)$$

where $\psi_i(s_i, s_j)$ and its first partial derivatives are given by

$$\psi_i(s_i, s_j) = \phi_i(s_i, s_j) - f(s_i)G_j(s_j), \quad (3.29)$$

$$\frac{\partial \psi_i(s_i, s_j)}{\partial s_i} = \frac{\partial \phi_i(s_i, s_j)}{\partial s_i} - f'(s_i)G_j(s_j), \quad (3.30)$$

$$\frac{\partial \psi_i(s_i, s_j)}{\partial s_j} = -f(s_j)G'_i(s_i) - f(s_i)G'_j(s_j), \quad (3.31)$$

and where $\phi_i(s_i, s_j)$, $\partial \phi_i(s_i, s_j)/\partial s_i$, and $\partial \phi_i(s_i, s_j)/\partial s_j$ are given by (3.14)–(3.16).

Under competition, we saw that the Nash equilibrium (s_i^e, s_j^e) always resides in the region B of Figure 3.2, that is, the active basestock levels at equilibrium are larger than the myopic basestock levels. Does this also hold for the optimal active basestock level pair under cooperation, denoted by (s_i^c, s_j^c) ? The following theorem answers this question.

Theorem 3.4. *The optimal active basestock level pair (s_i^c, s_j^c) and the resulting maximum team payoff $\Pi(s_i^c, s_j^c)$ satisfy*

(i) *If $G_i(s_i^m) = G_j(s_j^m)$, then*

$$s_i^c = s_i^m, \quad i = 1, 2, \quad (3.32)$$

$$\Pi(s_i^c, s_j^c) = \Pi(s_i^m, s_j^m) = G_i(s_i^m) = G_j(s_j^m). \quad (3.33)$$

(ii) *If $G_i(s_i^m) < G_j(s_j^m)$, then*

$$s_i^c \in [0, s_i^m) \text{ and } s_j^c \in (s_j^m, s_j^M), \quad (3.34)$$

$$G_i(s_i^m) < \Pi(s_i^m, s_j^m) < \Pi(s_i^c, s_j^c) < G_j(s_j^c). \quad (3.35)$$

In both cases:

$$\Pi(s_i^c, s_j^c) > \Pi(s_i^e, s_j^e), \quad (3.36)$$

where $\Pi(s_i^e, s_j^e) = \Pi_i(s_i^e, s_j^e) + \Pi_j(s_i^e, s_j^e)$ is the sum of the payoffs of the two suppliers at equilibrium under competition.

Expressions (3.32) and (3.34) state that under cooperation, the optimal active basestock level of one supplier is at or below his myopic basestock level, and therefore below his Nash equilibrium, whereas the optimal active basestock level of the other supplier is at or above his myopic basestock level. Expression (3.36) states that in both cases, the team payoff under cooperation is greater than the sum of the payoffs at equilibrium under competition.

If $G_i(s_i^m) = G_j(s_j^m)$, both suppliers use their myopic basestock levels, reaping the maximum possible profits for the team. A special case is when the suppliers are symmetric, given by the following corollary.

Corollary 3.1. *If the suppliers are symmetric, the optimal active basestock levels and the resulting team payoff under cooperation are:*

$$s_i^c = s^m, \quad i = 1, 2, \quad (3.37)$$

$$\Pi(s^m, s^m) = G(s^m). \quad (3.38)$$

Corollary 3.1 implies that the two cooperating symmetric suppliers behave as one newsvendor in the buyer's eyes.

If $G_i(s_i^m) < G_j(s_j^m)$, supplier i has a smaller myopic profit than supplier j , so he uses an active basestock level that is below his myopic basestock level, ceding a part of his demand share to the more profitable supplier j , who uses an active basestock level which is above his myopic basestock level. Therefore, the suppliers use active basestock levels that reside in the region A of Figure 3.2. If $G_i(s_i^m) > G_j(s_j^m)$, the reverse is true, and the suppliers use active basestock levels in region D . In both cases, both suppliers sacrifice some of their myopic profits to optimally rebalance their demand shares by transferring some of the buyer's business from the less profitable to the more profitable supplier.

To visualize the behavior of (s_i^c, s_j^c) , suppose that $h_i = I_i c_i$ and $b_i = J_i c_i$, where I_i and J_i are some proportionality constants. Then, from (3.4), $G_i(s_i)$ is decreasing in c_i and so is the ratio $G_i(s_i^m)/G_j(s_j^m)$, while from (3.7), s_i^m is constant in c_i because $b_i/(h_i + b_i) = J_i/(I_i + J_i)$ is independent of c_i . Figure 3.3 shows a curve tracing the position of (s_i^c, s_j^c) as a function of c_i . The red part of the curve represents the set




Figure 3.3: Position of the optimal active basestock level pair when the two suppliers cooperate.

of points (s_i^c, s_j^c) that corresponds to large values of c_i such that $G_i(s_i^m) < G_j(s_j^m)$. The blue part represents the set of points that correspond to small values of c_i such that $G_i(s_i^m) > G_j(s_j^m)$. If $G_i(s_i^m)$ is too small or too big, then s_i^c or s_j^c becomes zero, respectively.

A question that arises naturally is, when does the less profitable supplier set his active basestock level at zero, thereby ceding almost all his demand share to the more profitable supplier? We say “almost,” because even if the less profitable supplier sets his active basestock level at zero, the buyer will still return to him occasionally for a supply run of just one period whenever the more profitable supplier fails her. A follow-up question is, how is the optimal active basestock level of the more profitable supplier compared to his active basestock level at equilibrium under competition? The following theorem answers these questions under conditions that guarantee the uniqueness of the optimal active basestock level pair.

Theorem 3.5. *Assuming without loss of generality that $G_i(s_i^m) < G_j(s_j^m)$, if the*

following conditions hold:

$$\frac{\partial \psi_k(s_i, s_j)}{\partial s_k} < 0, \quad k = i, j, \quad s_i \in [0, s_i^m], \quad s_j \in (s_j^m, s_j^M), \quad (3.39)$$

the optimal active basestock level pair (s_i^c, s_j^c) is a global maximizer of the team payoff $\Pi(s_i, s_j)$ satisfying

- (i) If $\psi_i(0, s_j^c) > 0$, then $s_i^c > 0$ and the pair (s_i^c, s_j^c) uniquely satisfies $\partial \Pi(s_i, s_j) / \partial s_k = 0$, $k = i, j$, which reduces to

$$\psi_k(s_i^c, s_j^c) = 0, \quad k = i, j, \quad (3.40)$$

implying that

$$f_i(s_i^c)G_j'(s_j^c) + f(s_j^c)G_i'(s_i^c) = 0. \quad (3.41)$$

Otherwise, $s_i^c = 0$ and s_j^c uniquely satisfies $\partial \Pi(0, s_j) / \partial s_j = 0$, which reduces to

$$\psi_j(0, s_j^c) = 0. \quad (3.42)$$

- (ii) If condition (3.18) holds for both suppliers, then $s_j^c < s_j^e$.

Theorem 3.5 (i) provides a condition under which the active basestock level of the less profitable supplier i is strictly positive. If this condition holds, (3.41) implies that the relative values of s_i^c and s_j^c depend only on h_i , b_i , and f , even though from (3.40) their individual values depend on all problem parameters. If this condition does not hold, supplier i sets his active basestock level at zero, ceding almost all his business to the more profitable supplier j .

A question that arises, in this case, is, how high does s_j^c become to curtail the frequency of the buyer's occasional visits to supplier i . Does it ever increase above the Nash equilibrium s_j^e ? Theorem 3.5 (ii) implies that it does not if the conditions ensuring the existence of a Nash equilibrium hold. Therefore, under these conditions, the optimal active basestock levels of both suppliers under cooperation are smaller than their respective active basestock levels at Nash equilibrium.

Finally, as mentioned earlier, from (3.36), the team payoff under cooperation is greater than the sum of the payoffs at equilibrium. The loss in efficiency for the suppliers, if they compete instead of cooperating, can be measured by the ratio of their optimal team payoff under cooperation to their worst total payoff at equilibrium under competition, known as the *price of anarchy* (PoA), i.e.,

$$\text{PoA} = \frac{\Pi(s_i^c, s_j^c)}{\min_{(s_i^e, s_j^e)} \{\Pi(s_i^e, s_j^e)\}}. \quad (3.43)$$

3.5.2 The buyer's perspective

We note that PoA defined in (3.43) usually refers to the degradation of social welfare due to the selfish behavior of agents, whereas in our case, it is the suppliers' profits that are at stake. Moreover, the loss in efficiency for the suppliers is a gain in service quality for the buyer. The buyer's carrot-and-stick behavior is precisely meant to raise the fill rate that she enjoys by stimulating competition, and it does so successfully. If the tables are turned and the suppliers decide to cooperate instead of competing, the buyer loses the high-fill rate advantage that her behavior incites. What counteroffensive action can she take in this case to gain back that advantage? One plausible countermeasure for the buyer is to charge the suppliers an extra backorder penalty rate—different from the *regular* backorder cost rate b_i that we have been using thus far—to force them to increase their active basestock levels. The question then is, what should the value of this penalty cost be to make the suppliers raise their active basestock levels to their equilibrium values? We call this value the *adjustment backorder penalty* rate. To compute this penalty rate, we assume that the suppliers are symmetric for mathematical simplification. Moreover, we assume that their common regular backorder cost rate b is zero. As mentioned earlier, under supplier competition, the role of b is weakened anyway because the increase in the suppliers' active basestock levels reduces the frequency of stockouts. As the suppliers are symmetric, the buyer charges them a common backorder penalty rate denoted by b^c .

If the suppliers compete, from Proposition (3.2), there exists a unique symmetric

Nash equilibrium s^e satisfying (3.24). After substituting $G(s^e)$ and $G'(s^e)$ from (3.4) and (3.5), respectively, dividing them by h , and setting $b = 0$, expression (3.24) becomes

$$f(s^e) \left((p/h)\theta - E[(s^e - w)^+] \right) = 2F(s^e)\bar{F}(s^e). \quad (3.44)$$

From the above expression, s^e depends on the buyer's distribution and the ratio p/h . If $h = Ic$, where I is the interest rate, then $p/h = [(r - c)/c]/I$, i.e., p/h is the ratio of the margin rate to the interest rate. If the suppliers decide to cooperate, then from (3.37), they set their symmetric active basestock levels at s^m , where s^m is given by (3.7) after replacing b with b^c . So, if the buyer wants to recover the fill rate that she can enjoy under supplier competition, she must set b^c so that $s^m = s^e$, where s^e satisfies (3.44), i.e.,

$$b^c = h \frac{F(s^e)}{\bar{F}(s^e)}. \quad (3.45)$$

3.6 Exponentially distributed demand

To better comprehend the results developed in the previous sections and their implications, we apply them to the case where the buyer's demand is exponentially distributed. The exponential distribution has been used in many newsvendor model applications over the years Mahajan and van Ryzin (2001); Liyanage and Shanthikumar (2005); Rossi, Prestwich, Tarim and Hnich (2014); Ülkü and Gürler (2018); Siegel and Wagner (2021). Besides, the mathematical tractability that it offers, it has been recognized to effectively describe highly variable demand Lau (1997); Gallego, Katircioglu and Ramachandran (2007).

3.6.1 Analytical results

If the buyer's demand is exponentially distributed with rate λ , then

$$f(w) = \lambda e^{-\lambda w}, \quad \bar{F}(w) = e^{-\lambda w}, \quad w \geq 0, \lambda > 0, \text{ and } \theta = 1/\lambda. \quad (3.46)$$

To facilitate our analysis, we use the Lambert W function, defined as the inverse

function of we^w , i.e., $W(z) = w \Leftrightarrow z = we^w$. The Lambert W function is often used to solve equations in which the unknown appears both outside and inside an exponential function or a logarithm. It has the following properties which are useful for our analysis:

- (i) $W(z)$ is increasing and concave in $(-1/e, \infty)$ and positive in $(0, \infty)$.
- (ii) $W(we^w) = w$.
- (iii) $W'(z) = \frac{W(z)}{z[1 + W(z)]}$, for $z \notin \{0, -1/e\}$.

To simplify notation, we also define the following ratios:

$$\rho_i = \frac{p_i}{h_i}, \quad (3.47)$$

$$\beta_i = \frac{b_i + h_i}{h_i}. \quad (3.48)$$

As mentioned in the previous section, if $h_i = I_i c_i$, where I_i is the interest rate used by supplier i , then ρ_i represents the margin-to-interest rate ratio of supplier i . Using this notation, if $f(w)$ is given by (3.46), we can obtain exact expressions for $G_i(s_i)$, $G'_i(s_i)$, s_i^m , s_i^M , $G_i(s_i^m)$, $\Pi_i(s_i, s_j)$, $\phi_i(s_i, s_j)$, and $\psi_i(s_i, s_j)$. These expressions are given in Appendix B.

As mentioned in the discussion following Theorem 3.2, if $f(w)$ is non-increasing, condition (3.18) holds, guarantying the uniqueness and monotonicity of the best response. In the case of the exponential distribution, the best response is given in closed form in the following proposition.

Proposition 3.3. *If $f(w)$ is given by (3.46), the best response $s_i^*(s_j)$ is*

(i) *unique and given by*

$$s_i^*(s_j) = \frac{\rho_i + \beta_i e^{-\lambda s_j} - W\left(e^{\rho_i + \beta_i e^{-\lambda s_j} - \lambda s_j}\right)}{\lambda}, \quad s_j \in [0, \infty), \quad (3.49)$$

(ii) bounded as follows:

$$s_i^m < s_i^*(0) < s_i^*(s_j) < \lim_{s_j \rightarrow \infty} s_i^*(s_j) < s_i^M, \quad (3.50)$$

where $s_i^*(0) = [\rho_i + \beta_i - W(e^{\rho_i + \beta_i})]/\lambda$, $\lim_{s_j \rightarrow \infty} s_i^*(s_j) = \rho_i/\lambda$, and s_i^m and s_i^M are given by (B.6) and (B.10), respectively, in Appendix B.

Expression (3.49) implies that $s_i^*(s_j)$ is increasing in ρ_i and β_i which also verifies Theorem 3.2 (iii). Moreover, (3.50) provides tighter lower and upper bounds than s_i^m and s_i^M , respectively.

From Theorem (3.3) (i), condition (3.18) also implies the existence of at least one pure-strategy Nash equilibrium. The following proposition states that the Nash equilibrium is unique and provides the equations to compute it.

Proposition 3.4. *If $f(w)$ is given by (3.46), there exists a unique pure-strategy Nash equilibrium (s_i^e, s_j^e) satisfying*

$$s_j^e = \frac{1}{\lambda} \ln \left(\frac{e^{\lambda s_i^e} - \beta_i}{\rho_i - \lambda s_i^e} \right), \quad i = 1, 2. \quad (3.51)$$

The proof is based on showing that condition (3.23) holds for the exponential case. From (3.51), the active basestock level of supplier j at equilibrium is increasing and concave in s_i^e and decreasing in ρ_i, β_i . The system of equations given by (3.51) cannot be solved analytically. However, for the symmetric case, we can obtain a closed-form solution which is given by the following corollary.

Corollary 3.2. *If the suppliers are symmetric and $f(w)$ is given by (3.46), there exists a unique pure-strategy Nash equilibrium (s_i^e, s_j^e) which is symmetric, i.e., $s_i^e = s_j^e = s^e$, where s^e is given by:*

$$s^e = \frac{\rho - 1 + W(\beta e^{1-\rho})}{\lambda}. \quad (3.52)$$

The resulting payoff of each supplier i is

$$\Pi_i(s^e, s^e) = \frac{h(1 - W(\beta e^{1-\rho}))}{\lambda}, \quad i = 1, 2. \quad (3.53)$$

The proof follows from Proposition 3.2. Equation (3.52) can also be derived from (3.51) after dropping the supplier indexes and solving for s^e .

The following result regards the cooperation of the suppliers.

Proposition 3.5. *If $f(w)$ is given by (3.46) and assuming without loss of generality that*

$$\Delta p > h_j \ln(\beta_j) - h_i \ln(\beta_i), \quad (3.54)$$

where $\Delta p = p_j - p_i = h_j \rho_j - h_i \rho_i$, then the optimal active basestock level pair (s_i^c, s_j^c) satisfies:

If $\Delta p < h_j \ln(K)$, where $K = \beta_j + (\beta_i - 1)h_i/h_j$, then $s_i^c \in (0, s_i^m)$, $s_j^c \in (s_i^m, s_j^M)$ and the pair (s_i^c, s_j^c) uniquely satisfies

$$h_j s_j^c - h_i s_i^c = \frac{\Delta p}{\lambda}, \quad (3.55)$$

$$h_i e^{\lambda s_i^c} + h_j e^{\lambda s_j^c} = h_i \beta_i + h_j \beta_j. \quad (3.56)$$

Otherwise:

$$s_i^c = 0, \quad \text{and} \quad s_j^c = \frac{\mu - W(e^\mu)}{\lambda}, \quad (3.57)$$

where $\mu = \Delta p/h_j - K$.

Proposition (3.5) states that if the difference in the margins of the suppliers, $\Delta p = p_j - p_i$, is larger than the difference $h_j \ln(\beta_j) - h_i \ln(\beta_i)$, then $G_i(s_i^m) < G_j(s_j^m)$. This means that supplier j is more profitable than supplier i . So, when the suppliers team up, supplier j uses a basestock level that is above his myopic basestock level, while supplier i uses a level that is below his myopic basestock level, ceding a part of his demand share to supplier j . Therefore, the suppliers use active basestock levels in region A of the (s_i, s_j) space shown in Figure 3.3. If, in addition, $\Delta p \geq h_j \ln(K)$, then supplier i sets his active basestock level at zero, ceding almost all his demand share to supplier j .

Note that Δp does not have to be positive for supplier j to be more profitable than supplier i . That is, supplier j can have a smaller margin than supplier i , i.e., $\Delta p < 0$, and still be more profitable than i , if $\Delta p > h_j \ln(\beta_j) - h_i \ln(\beta_i)$. In this case, however, Δp will certainly be smaller than $h_j \ln(K)$, which means that supplier i will not set his active basestock level at zero.

If $\Delta p < h_j \ln(K)$, the optimal active basestock levels of the two suppliers uniquely solve equations (3.55) and (3.56). If we substitute λs_j^c from the first equation into the second, we obtain an equation of the form

$$a_1 x^{a_2} + a_2 x = a_3,$$

where $a_1 = e^{\Delta p/h_j}$, $a_2 = h_i/h_j$, $a_3 = (h_i \beta_i + h_j \beta_j)/h_j$, and $x = e^{\lambda s_i^c}$. This equation is increasing in x and has a unique solution. In general, however, we cannot obtain a closed form for it, except for special cases, e.g., when $a_2 = 1, 2$, etc. An interesting result is given by the following corollary.

Corollary 3.3. *If $\Delta p = 0$ and $h_i = h_j = h$, then $s_i^c = s_j^c = s^c$, where*

$$s^c = \frac{\ln(\bar{\beta})}{\lambda}, \quad (3.58)$$

where $\bar{\beta} = (\bar{b} + h)/h$ and $\bar{b} = (b_i + b_j)/2$.

The proof follows from (3.55) and (3.56). The intuition behind Corollary 3.3 is that if the suppliers have the same margins and inventory cost rates, they bring in the same profits to the team and incur the same inventory costs, so there is no reason for them not to split the demand by setting their active basestock levels equal to each other. If the suppliers have different backorder cost rates, say $b_i > b_j$, then $s_i^c < s_i^m$ and $s_j^c > s_j^m$, but the important fact remains that $s_i^c = s_j^c$. To see why the difference in the backorder cost rates makes no difference, consider the following. Every time supplier i fails to deliver on-demand, the team pays b_i , and every time supplier j fails, the team pays b_j . Because the supplier switches from one supplier to the other, the team pays $b_i + b_j$ in every full cycle with two switches. Although, this cost matters, how it is divided among the suppliers is not important for the team's profit, because

it will be paid by the team as a sum. Based on this argument, we conjecture that the main result of the corollary, i.e., that $s_i^c = s_j^c = s^c$, holds for any demand distribution. Expression (3.58) is special to the exponential distribution.

Finally, if the suppliers are symmetric, then by Corollary 3.1, their maximum team payoff is $G(s^m)$, which for the exponential case is given by (B.11), after dropping the supplier index. On the other hand, the payoff of each supplier under the unique Nash equilibrium is given by (3.53). Therefore, for the symmetric case, the price of anarchy defined in (3.43) becomes:

$$\text{PoA} = \frac{\rho - \ln(\beta)}{2(1 - W(\beta e^{1-\rho}))}. \quad (3.59)$$

If the symmetric suppliers decide to cooperate, the adjustment backorder penalty rate b^c that the buyer must charge them to recover the fill rate that she can enjoy under supplier competition (assuming that $b = 0$) is found, after the analysis in Section 3.5.2, by setting $s^m = s^e$, where s^m is given by (B.6) in Appendix B with $\beta = (b^c + h)/h$ and s^e is given by (3.52) with $\beta = 1$ (since $b = 0$), and solving for b^c . The solution is

$$b^c = h \left(e^{\rho-1+W(e^{1-\rho})} - 1 \right). \quad (3.60)$$

From (3.60), b^c is h times a factor that is approximately exponentially increasing in ρ since $W(e^{1-\rho}) \in (0, 1)$ for $\rho > 0$. This is expected, because as ρ increases, the margin p becomes increasingly more important than the inventory holding cost rate h , pushing the active basestock levels at equilibrium increasingly higher. Therefore, if the suppliers cooperate, the buyer needs to charge them an increasingly larger penalty rate b^c to make them raise their active basestock levels to the equilibrium values.

3.6.2 Numerical example

We illustrate the analytical results developed in the previous subsection with a numerical example, also investigating the effect of the problem parameters on the optimal active basestock levels and the resulting performance measures. In the example, we assume that the buyer's demand is exponentially distributed with rate $\lambda = 1$ and that

the inventory cost rate of each supplier k ($k = i, j$) is given by $h_k = I_k c_k$, where I_k is the interest rate used by the supplier. Initially, we consider a nominal instance in which the suppliers are symmetric with parameter values $c_k = 1$, $I_k = 0.4$, $r_k = r = 3$, and $b_k = b = 0.7$, for $k = i, j$. Then, we vary the values of certain parameters one at a time within a certain range.

As we vary each parameter value, we calculate the myopic basestock level pair (s_i^m, s_j^m) from (B.6), the active basestock level pair at equilibrium (s_i^e, s_j^e) by solving (3.51), and the optimal active basestock level pair under cooperation (s_i^c, s_j^c) by solving either (3.55) and (3.56) or (3.57). We also calculate the resulting payoffs (Π_i^e, Π_j^e) and (Π_i^c, Π_j^c) from (B.7) in Appendix B, the demand shares (π_i^e, π_j^e) and (π_i^c, π_j^c) from (3.11), and the fill rates q_i^e and q^c from (3.12). Finally, we compute the price of anarchy PoA from (3.43), where $\Pi(s_i^c, s_j^c) = \Pi_i^c + \Pi_j^c$ and $\Pi(s_i^e, s_j^e) = \Pi_i^e + \Pi_j^e$, and the adjustment backorder penalty rate b^c from (3.60).

Figures 3.4, 3.5, and 3.6 show plots of the above-calculated values as we vary c_j , I_j , and r_j ¹, respectively. The difference between c_j and I_j is that while c_j






Figure 3.4: Optimal basestock levels and performance measures vs. c_j .

affects the margin $p_j = r_j - c_j$ and the inventory cost $h_j = I_j c_j$, I_j affects only h_j . Increasing either parameter, however, reduces supplier j ' payoff. On the other hand, increasing r_j , raises his payoff. For this reason, the plots in figures 3.4 and 3.5 have the same structure, whereas the plots in Figure 3.6 have a symmetric structure. We will therefore briefly discuss only the plots in Figure 3.4.

From the first three plots, we observe that supplier j 's active basestock level,

¹We vary r_j for the sake of completeness because, as the suppliers compete solely on availability, it is natural to assume that $r_i \approx r_j$

Figure 3.5: Optimal basestock levels and performance measures vs. I_j .Figure 3.6: Optimal basestock levels and performance measures vs. r_j

payoff, and demand share at equilibrium and under cooperation are decreasing in c_j , reflecting the resulting drop in p_j and rise in h_j . Supplier i 's active basestock level at equilibrium is also decreasing in c_j , although at a smaller rate, echoing the drop in s_j^e . His active basestock level under cooperation, however, as well as his payoff and demand share at equilibrium and under cooperation are increasing in c_j , reflecting the decrease in supplier j 's profitability. The active basestock levels of both suppliers under competition are higher than their myopic levels, confirming Proposition 3.1.

When $c_j = c_i = 1$, the suppliers are symmetric and have the same active basestock levels, payoffs, and demand shares. More specifically, their active basestock levels and payoffs under cooperation are equal to the corresponding myopic values, confirming Corollary 3.1. When $c_j > 1$, supplier j becomes less profitable than supplier i , so s_j^c drops below s_j^m , whereas s_i^c rises above s_i^m , confirming Theorem 3.4 (ii). When $c_j \gtrapprox 1.6$, $s_j^c = 0$, confirming Theorem 3.5 (i), whereas s_i^c keeps increasing in c_j but remains below s_i^e , confirming Theorem 3.5 (ii).

The fill rate that the buyer enjoys under competition is generally very high and

is not significantly affected by c_j . The fill rate under cooperation, on the other hand, is significantly lower, dropping to approximately 65% when $c_j \approx 1.6$. The team payoff under cooperation is more than double the sum of the suppliers' payoffs under competition, as indicated by the PoA plot.

Figure 3.7 shows plots of the optimal basestock levels and performance measures as we vary b_j . As expected, s_j^c is increasing in b_j while Π_i^c is decreasing, although




Figure 3.7: Optimal basestock levels and performance measures vs. b_j .

these changes are very subtle. We also observe that $s_i^c = s_j^c$ and $d_i^c = d_j^c$ for all values of b_j , confirming Corollary 3.3. From Figure 3.7, it appears that b_j has no effect on the active basestock levels and performance measures at equilibrium. A close-up of the first three plots in Figure 3.7, shown in Figure 3.8, reveals that b_j affects these quantities but the effect is negligible. The reason for this is that the exponential




Figure 3.8: Close-up of optimal active basestock levels and performance measures vs. b_j .

terms in equations (3.51) are much larger than the other terms, and as a result, the solution (s_i^e, s_j^e) is sensitive to the multiplicative terms ρ_i but insensitive to the

additive terms β_i . Hence it is insensitive to the backorder cost rates. The intuition behind this observation is that the main concern of the suppliers under competition is to maintain the buyer's loyalty because losing it as a result of a stockout means relinquishing profits for many periods following the stockout. This concern drives the suppliers to significantly increase their active basestock levels above their myopic levels as can be seen in Figure 3.7. Avoiding paying the backorder cost is therefore of minimal concern. Because the active basestock levels at equilibrium are so much larger than the myopic levels, the frequency of stockouts is significantly reduced, further limiting the impact of backorder costs, as mentioned earlier.

Figures 3.9 and 3.10 show plots of the optimal active basestock levels and performance measures as we vary r and b .




Figure 3.9: Optimal active basestock levels and performance measures vs. r .




Figure 3.10: Optimal active basestock levels and performance measures vs. b .

From these two figures, we observe that the optimal symmetric active basestock levels and performance measures under cooperation are sensitive in b and insensitive in r , whereas, under competition, they are sensitive in r and insensitive in b . The reason for this is that from (3.37), $s^c = s^m$ where from (B.6), s^m depends only on

the tradeoff between h and b , expressed by β , and is independent of r . On the other hand, from (3.52), s^e depends mainly on the tradeoff between h and $p = r - c$, expressed by ρ . It also depends on β through the Lambert W function, but this dependence is negligible for the same reason that the asymmetric active basestock levels are insensitive to b_j , discussed earlier.

Finally, Figure 3.11 shows plots of the adjustment backorder penalty rate b^c as we vary c , I , and r . From these figures, we observe that the effect of these parameters on




Figure 3.11: Adjustment backorder penalty cost b^c vs. c , I , and r , for $b = 0$.

b^c is dramatic, verifying our remark following (3.60) that b^c is exponentially increasing in ρ . For instance, for a 40% margin rate and a 20% interest rate, $\rho = [(r - c)/c]/I = 0.4/0.2 = 2$. For $\rho = 2$, $b^c = 2.6h$ from (3.60). For $\rho = 4$ and $\rho = 6$, $b^c = 20.1h$ and $148.4h$, respectively. In practice, it does not make sense for the buyer to raise b^c above a fraction of the selling price r . So, if the suppliers decide to cooperate and ρ is large, the buyer will not be able to match the fill rate that she can enjoy under supplier competition.

3.7 Extension to multiple sourcing

As we wrote in the Introduction, dual sourcing predominates multiple sourcing. Nevertheless, the analysis of the two-supplier model can be straightforwardly extended to $n > 2$ suppliers under the following setting. The buyer arranges the n suppliers in a rank order list based on the last service she received. At the beginning of period t ,

each supplier i orders a non-negative quantity ahead of demand, based on his inventory level, $x_{i,t} \in \mathbb{R}$, and his placement in the buyer's list (ranking), $\alpha_{i,t} \in \{1, 2, \dots, n\}$, where 1 indicates the top of the list (highest ranking) and n indicates the bottom (lowest ranking). The order arrives before the end of the period, raising the supplier's inventory level to $y_{i,t} \geq x_{i,t}$.

At the end of the period, the buyer selects the highest-ranking supplier at the top of the list and demands from him a random quantity w_t . If the supplier meets all the demand at once, he is kept at the top of the list and carries any leftover inventory to the next period. If he fails to meet all the demand at once, the buyer backorders the unmet demand with him and moves him to the bottom of the list, thereby bringing all the other suppliers one position closer to the top of the list. This way, the buyer selects all the suppliers in a round-robin fashion, switching suppliers after every stockout.

Round-robin is a common process for fair resource allocation. It is a popular method for scheduling processes in computer and communication systems, traffic and transportation systems, production systems—most notably in the context of the *stochastic economic lot scheduling problem (SELSP)*—and other polling systems Boon, van der Mei and Winands (2011), scheduling games in sports tournaments Rasmussen and Trick (2008), and other applications. It is simple, easy to implement, and starvation-and-envy-free. It is also one of the choices for supplier allocation in many ERP systems. For example, in SYSPRO's Preferred Supplier feature, it is one of the sourcing options SYSPRO (2022). In SAP's Allocation Quota Arrangement feature, each procurement lot is assigned to a source of supply based on its quota rating SAP (2022c). If the quota of all sources are set equal (which is the default value), the lots are assigned on a round-robin basis. What we propose in this paper is for the buyer to use round-robin as a fair and starvation-free supplier selection scheme, allowing each supplier to keep his “preferred supplier” status as long as he can afford to before giving his turn to the next supplier.

Under this setting, Theorem 3.1 immediately extends to n suppliers. Namely, the optimal ordering policy of supplier i is a ranking-dependent basestock policy, denoted by $y_i^*(\alpha_i)$, given by: $y_i^*(\alpha_i) = 0, \alpha_i \neq 1$ and $y_i^*(1) = s_i \geq 0$. Assuming, without loss

of generality, that the suppliers are numbered according to their initial position in the buyer's list, the ranking vector of the suppliers, $\alpha = (\alpha_1, \dots, \alpha_n)$, under the optimal policy is a discrete-time Markov chain with n states, $(1, 2, \dots, n-1, n)$, $(n, 1, \dots, n-2, n-1)$, $(n-1, n, \dots, n-3, n-2)$, $\dots$, $(2, 3, \dots, n, 1)$, arranged clockwise in a circle, and transition probabilities from the state where $\alpha_i = 1$ to the state where $\alpha_{i+1 \bmod n} = 1$, equal to $\bar{F}(s_i)$, $i = 1, \dots, n$, (see Figure 3.12, for $n = 4$).




Figure 3.12: Markov chain transition diagram of the ranking vector for $n = 4$.

It is trivial to show that the steady-state probability of the state where $\alpha_i = 1$, representing the expected average demand share of supplier i as a function of the vector of active basestock levels $\mathbf{s} = (s_1, \dots, s_n)$, denoted $\pi_i(\mathbf{s})$, is given by

$$\pi_i(\mathbf{s}) = \frac{\prod_{k \neq i} \bar{F}(s_k)}{\sum_l \prod_{k \neq l} \bar{F}(s_k)}. \quad (3.61)$$

From the above expression, the expected average profit (payoff) of supplier i , denoted by $\Pi_i(\mathbf{s})$, and its first partial derivative with respect to s_i are

$$\Pi_i(\mathbf{s}) = \pi_i(\mathbf{s}) G'_i(s_i), \quad (3.62)$$

$$\frac{\partial \Pi_i(\mathbf{s})}{\partial s_i} = \frac{\prod_{k \neq i} \bar{F}(s_k)}{\left(\sum_l \prod_{k \neq l} \bar{F}(s_k) \right)^2} \phi_i(\mathbf{s}), \quad (3.63)$$

where $\phi_i(\mathbf{s})$ is given by:

$$\phi_i(\mathbf{s}) = \left(\prod_{k \neq i} \bar{F}(s_k) \right) G'_i(s_i) + \left(\sum_{l \neq i} \prod_{k \neq l, i} \bar{F}(s_k) \right) (\bar{F}(s_i) G'_i(s_i) + f(s_i) G_i(s_i)), \quad (3.64)$$

where by convention, the products in the above expressions equal 1 if they contain no terms.

The following Theorem extends Proposition 3.1 and Theorem 3.2 to multiple suppliers, where $\mathbf{s}_{-i}$ denotes the vector of active basestock levels of all suppliers except i .

Theorem 3.6. *The best response $s_i^*(\mathbf{s}_{-i})$ is bounded as follows:*

$$0 < s_i^m < s_i^*(\mathbf{s}_{-i}) < s_i^M, \quad s_k \in [0, \infty), \quad k \neq i, \quad (3.65)$$

where s_i^m and s_i^M satisfy (3.7) and (3.8). Moreover, if the following condition holds:

$$\frac{\partial \phi_i(\mathbf{s})}{\partial s_i} < 0, \quad s_i \in (s_i^m, s_i^M), \quad s_k \in [0, \infty), \quad k \neq i, \quad (3.66)$$

the best response $s_i^*(\mathbf{s}_{-i})$ is

(i) a global maximizer of the payoff $\Pi_i(\mathbf{s})$, $s_k \in [0, \infty)$, $k \neq i$, uniquely satisfying the first-order condition $\partial \Pi_i(\mathbf{s}) / \partial s_i = 0$, which reduces to:

$$\phi_i(s_i^*(\mathbf{s}_{-i}), \mathbf{s}_{-i}) = 0, \quad (3.67)$$

(ii) increasing in s_j , and its derivative with respect to s_j is

$$\frac{\partial s_i^*(\mathbf{s}_{-i})}{\partial s_j} = - \frac{\partial \phi_i(s_i^*(\mathbf{s}_{-i}), \mathbf{s}_{-i}) / \partial s_j}{\partial \phi_i(s_i^*(\mathbf{s}_{-i}), \mathbf{s}_{-i}) / \partial s_i}. \quad (3.68)$$

(iii) increasing in θ , p_i , b_i , and decreasing in h_i .

As in the case of the dual-sourcing model, condition (3.66) is very mild and is easily satisfied. For example, if $f(w)$ is non-increasing, (3.66) is immediately met. The expression for $\partial \phi_i(\mathbf{s}) / \partial s_i$ is given by equation (B.2) in the proof. If we use that expression and rearrange terms, condition (3.66) reduces to:

$$f'(s_i) < - \left(\frac{\sum_l \Pi_{k \neq l} \bar{F}(s_k)}{\sum_{l \neq i} \Pi_{k \neq l, i} \bar{F}(s_k)} \right) \frac{G_i''(s_i)}{G_i'(s_i)}, \quad s_i \in (s_i^m, s_i^M), \quad s_k \in [0, \infty), \quad k \neq i, \quad (3.69)$$

where the right-hand side of the above inequality is positive. The above expression is the extension of (3.21) to n suppliers.

As in the dual-sourcing model, Theorem 3.6 (ii) implies that under condition (3.66) if supplier j increases his active basestock level, supplier i will follow suit to mitigate his loss of demand share. Therefore, competition drives all the suppliers to move away from their myopic basestock levels, s_i^m .

The issue of the existence and uniqueness of a Nash equilibrium for the general case of asymmetric suppliers is outside the scope of this paper. Here, it suffices to point out that the best response functions are increasing and that “increasing best response functions is the only major requirement for an equilibrium to exist” Cachon and Zhang (2006). The following proposition extends Proposition 3.2 to multiple symmetric suppliers.

Proposition 3.6. *If the suppliers are symmetric and condition (3.66) holds, then*

- (i) *There exists at least one symmetric pure-strategy Nash equilibrium $\mathbf{s}^e = (s^e, \dots, s^e)$, where s^e satisfies:*

$$\frac{f(s^e)}{\bar{F}(s^e)} = -\frac{n}{n-1} \frac{G'(s^e)}{G(s^e)}. \quad (3.70)$$

The resulting payoff of each supplier i is:

$$\Pi_i(\mathbf{s}^e) = \frac{G(s^e)}{n}, \quad i = 1, \dots, n. \quad (3.71)$$

- (ii) *There exists no asymmetric pure-strategy Nash equilibrium.*

- (iii) *If the following condition holds:*

$$\frac{\partial \hat{\phi}(s)}{\partial s} < 0, \quad s \in (s^m, s^M), \quad (3.72)$$

where $\hat{\phi}(s) = \phi(s, \dots, s)$, then the symmetric equilibrium is unique.

The proof is similar to that of Proposition 3.2 and is therefore omitted. As in the dual sourcing case, note that if the suppliers are symmetric, condition (3.72) instead of (3.66) is needed for the uniqueness of the Nash equilibrium, because the first-order

conditions for the two suppliers reduce to one equation (because of symmetry) which has a unique solution, under (3.72).

If the suppliers decide to cooperate, then the optimal ordering policy of each supplier in the team has the same structure as that under competition. Moreover, the payoff of the team, denoted by $\Pi(\mathbf{s})$, is the sum of the individual payoffs of the suppliers, i.e.:

$$\Pi(\mathbf{s}) = \sum_i \Pi_i(\mathbf{s}) = \sum_i \pi_i(\mathbf{s}) G_i(s_i). \quad (3.73)$$

Assuming without loss of generality that the suppliers are numbered from 1 to n so that $G_1(s_1^m) \leq G_2(s_2^m) \leq \dots \leq G_n(s_n^m)$, the following theorem extends Theorem 3.4 to multiple suppliers.

Theorem 3.7. *The optimal basestock level vector $\mathbf{s}^c$ and the resulting maximum team payoff $\Pi(\mathbf{s}^c)$ satisfy:*

(i) *If $G_1(s_1^m) = G_2(s_2^m) = \dots = G_n(s_n^m)$, then*

$$s_i^c = s_i^m, \quad i = 1, 2, \dots, n, \quad (3.74)$$

$$\Pi(\mathbf{s}^c) = \Pi(\mathbf{s}^m) = G_1(s_1^m) = G_2(s_2^m) = \dots = G_n(s_n^m). \quad (3.75)$$

(ii) *If $G_1(s_1^m) < G_2(s_2^m) < \dots < G_n(s_n^m)$, then there exists an index k such that:*

$$s_i^c \in [0, s_i^m), \quad i \leq k \quad \text{and} \quad s_j^c \in (s_j^m, s_j^M), \quad j > k. \quad (3.76)$$

$$G_1(s_1^m) < \Pi(\mathbf{s}^m) < \Pi(\mathbf{s}^c) < G_n(s_n^c). \quad (3.77)$$

In both cases:

$$\Pi(\mathbf{s}^c) > \Pi(\mathbf{s}^e), \quad (3.78)$$

where $\Pi(\mathbf{s}^e) = \sum_i \Pi_i(\mathbf{s}^e)$ is the sum of the payoffs of the n suppliers at equilibrium under competition.

The proof is similar to that of Theorem 3.4 and is therefore omitted. Theorem 3.7 implies that if the condition in (i) holds, all the suppliers have the same myopic profit.

In this case, any supplier moving away from his myopic basestock level hurts the team's profit. If the condition in (ii) holds, the customers are arranged in ascending order of myopic profit. The suppliers with the k smallest myopic profits use active basestock levels that are below their myopic basestock levels, ceding a part of their demand shares to the remaining $n - k$ suppliers who use active basestock levels that are above their myopic basestock levels. Therefore, all suppliers sacrifice some of their myopic profits to optimally balance their demand shares by transferring some of the buyers' business from the less profitable to the more profitable suppliers. The total expected average profit of the suppliers at equilibrium, if one exists, is higher than their team profit under cooperation, because under competition all suppliers use active basestock levels which are above and far from their myopic basestock levels, whereas under cooperation they use active basestock levels which, roughly speaking, are closer to their myopic basestock levels.

The round-robin switching policy of the buyer that we considered in this section is a particular policy that the buyer can use to stimulate competition on availability among multiple suppliers, but there can be other policies. One such policy, for example, is to let all the suppliers compete for the buyer's business on their active basestock levels without punishing them for failing to deliver on demand. Under such a policy, every supplier will try to overbid the other suppliers. As every supplier i incurs losses when his active basestock level is above s_i^M , the winning supplier will be the supplier with the highest value of s_i^M . He will set his active basestock level just above the second highest s_i^M value, and earn the supplier's loyalty, driving all the other suppliers out of the supplier's business. Such an outcome may not be desirable for the buyer.

To avoid this situation, an alternative policy is to let all the suppliers compete for the buyer's business on their active basestock levels, except for the supplier who failed most recently. Under this variant, the winning supplier will again be the supplier with the highest value of s_i^M unless he is the one who failed most recently. If this is the case, the winning supplier will be the supplier with the second-highest value of s_i^M . When that supplier fails, the supplier with the highest value of s_i^M will be eligible for selection again and will win the buyer's loyalty until he fails again. This cycle

will be repeated between the buyers with the two highest values of s_i^M , and all other suppliers will have been driven out of the buyer's business, leaving the buyer in the dual sourcing situation studied in the main part of this chapter.

The idea of banishing a supplier who fails for one supply run can be generalized to banishing him for k supply runs, where $k \in \{0, \dots, n\}$. In this case, the switching cycle will be repeated between the buyers with the k highest values of s_i^M , and all other suppliers will have been driven out of the buyer's business.

In the round-robin switching policy, $k = n$. Under this policy, when a supplier fails to deliver on demand, the buyer punishes him by sending him to the bottom of the list. This implies that the failed supplier will be selected again only after all the other suppliers fail, one after the other, i.e., his turn will come up again after $n - 1$ stockouts. Therefore, sending a failed supplier at the bottom of the list makes sense if the buyer wants to keep all the suppliers in business.

3.8 Discussion and future research

The behavior of an always-a-share buyer who plays her suppliers against each other by rewarding availability with loyalty and punishing stockouts with switching has significant implications for the suppliers' inventory policy and long-run average profit. The ordering decision of the supplier who enjoys the buyers' loyalty requires the careful balancing of his inventory and backorder costs against his future profit loss resulting from ceding the buyer's loyalty to his competitor(s). There are several possible directions for future work.

In our model, we assume that the buyer backorders any unmet demand with the supplier that she selects to meet the demand to ensure the uniformity and traceability of her order. A different possibility is to presume that if the selected supplier runs out of stock, the buyer tries to procure the missing items from the other supplier. In this case, it may be in the interest of the suppliers to hold some spare inventory even when they are inactive. This interest will be more intense if fully satisfying the residual demand results in gaining the buyer's loyalty in the next period.

A similar situation arises if the inactive supplier is not informed about the active

supplier's failure to serve the buyer. In this case, he must also hold some spare inventory to meet the buyer's demand if she calls on him without warning following a stockout by the other supplier.

Another intelligence-related issue concerns the information that the suppliers know about each other. In our model, we assume that each supplier knows the other supplier's cost and revenue parameters. In practice, these parameters can be unknown, in which case the supplier will have to estimate them through learning. The same holds for the buyer's demand distribution.

Finally, we assume that the buyer is loyal to one supplier as long as he serves her well but immediately switches to the other supplier after the first failure. If there is friction associated with switching, the buyer may think twice before switching at the first stockout incident. An alternative is to issue a warning to the supplier who fails the first time providing him with another chance to stay active, but switch after the second failure. In this case, the active supplier will use different active basestock levels depending on whether he has been issued a warning or not. Intuitively, the basestock level before the warning should be smaller than that after the warning, but it would be interesting to see how much smaller and also how both levels compare to the active basestock level when switching occurs after the first failure.

Chapter 4

Dynamic ordering and buyer selection policies in a newsvendor setting with service-dependent demand

4.1 Introduction

In this chapter, we study a newsvendor model of a firm that orders items for a group of repeat buyers. The buyers generate different revenues and have different average visit rates that depend on whether they are satisfied or dissatisfied with their last visit. In Section 4.2, we formulate the dynamic ordering and buyer selection problem of the firm. In Section 4.3, we determine the myopic policy and derive some important structural results on the optimal policy. In Section 4.4, we fully characterize the optimal policy for two buyers and compare it with other policies which form the basis of heuristic policies for larger problems. In Section 4.5, we probe into the optimal policy for more than two buyers by numerically solving and discussing a problem instance with three buyers. In Section 4.6, we set up the Lagrangian relaxation of the original problem. In Section 4.7, we develop three heuristic index policies for the buyer selection problem based on the relaxed problem. For the Lagrangian index policy, we

derive the “best” Lagrangian price in closed form as the solution to the Lagrangian dual problem. In Section 4.8, we explore and compare the performance of the three index policies by numerically solving a large number of problem instances with five and ten buyers. For the five-buyer instances, we also compare the index policies with the optimal policy. Finally, in Section 4.9, we summarize our findings and propose directions for future work. Supplemental material for this chapter, including proofs, can be found in Appendix C.

4.2 Model formulation

A firm supplies items to a finite set of buyers $\mathcal{B} = \{1, 2, \dots, n\}$ over consecutive time periods. At the beginning of each period t , the firm orders a quantity $y_t \in \mathcal{B}_0 = \{0, 1, \dots, n\}$, and at the end of the period, each buyer i (she) visits the firm with a probability that depends on whether she is satisfied or not with her previous visit. This probability is denoted by $q_i(\alpha_{i,t})$, where $\alpha_{i,t} \in \{0, 1\}$ is the *satisfaction state* of buyer i at the beginning of period t , with 0 meaning *dissatisfied* and 1 meaning *satisfied*. We assume that every buyer that visits the firm demands one item, and that satisfaction depends solely on item availability so that our results are not overshadowed by the complexity of other influencing factors. We refer to $q_i(\alpha_{i,t})$ and its complement $\bar{q}_i(\alpha_{i,t}) = 1 - q_i(\alpha_{i,t})$ as the average *visit rate* and *deferral rate* of buyer i , respectively, when in satisfaction state $\alpha_{i,t}$. The buyer’s demand is denoted by $d_i(\alpha_{i,t})$. For notational simplicity, henceforth, we omit the dependence of q_i and d_i on $\alpha_{i,t}$, wherever possible. The demand d_i is Bernoulli distributed with:

$$P(d_i = 1) = q_i = 1 - P(d_i = 0) = 1 - \bar{q}_i, \quad i \in \mathcal{B}.$$

We assume that the demands of different buyers are independent and that a satisfied buyer is more likely to demand service than a dissatisfied buyer is, i.e.:

$$q_i(1) \geq q_i(0) > 0, \quad i \in \mathcal{B}.$$

For convenience, we define:

$$\gamma_i = \frac{q_i(1) - q_i(0)}{q_i(1)}, \quad i \in \mathcal{B} \quad (4.1)$$

$$\bar{\gamma}_i = \frac{\bar{q}_i(0) - \bar{q}_i(1)}{\bar{q}_i(1)} = \frac{q_i(1) - q_i(0)}{\bar{q}_i(1)}, \quad i \in \mathcal{B}. \quad (4.2)$$

We refer to γ_i and $\bar{\gamma}_i$ as the *loss-of-visit-rate coefficient (LVC)* and *gain-of-deferral-rate coefficient (GDC)* of buyer i , respectively, when switching from the satisfied to the dissatisfied state.

The total demand of all buyers for any subset of buyers $\mathcal{A} \subseteq \mathcal{B}$ is denoted by $D_{\mathcal{A}}$, i.e., $D_{\mathcal{A}} = \sum_{j \in \mathcal{A}} d_j(\alpha_j)$. Random variable $D_{\mathcal{A}}$ is the sum of $n_{\mathcal{A}} = |\mathcal{A}|$ independent non-identical Bernoulli random variables, so it follows a Poisson binomial distribution Wang (1993). The p.m.f. and c.d.f. of $D_{\mathcal{A}}$ are denoted by $f_{\mathcal{A}}(k)$ and $F_{\mathcal{A}}(y)$, respectively, and are given by:

$$f_{\mathcal{A}}(k) = \sum_{\mathcal{X} \in [\mathcal{A}]^k} \prod_{j \in \mathcal{X}} q_j(\alpha_j) \prod_{j \in \mathcal{A} \setminus \mathcal{X}} \bar{q}_j(\alpha_j), \quad k = 0, \dots, n_{\mathcal{A}}, \quad (4.3)$$

$$F_{\mathcal{A}}(y) = \sum_{k=0}^y f_{\mathcal{A}}(k), \quad y = 0, \dots, n_{\mathcal{A}}, \quad (4.4)$$

where $[\mathcal{A}]^k$ is the set of k -combinations of $\mathcal{A}$. Computing $f_{\mathcal{A}}(k)$ and $F_{\mathcal{A}}(y)$ is computationally demanding even for modest values of $n_{\mathcal{A}}$, because $|[\mathcal{A}]^k| = C_k^{n_{\mathcal{A}}}$ can be very large, particularly for k close to $n_{\mathcal{A}}/2$, where $C_k^{n_{\mathcal{A}}}$ denotes the binomial coefficient.

In addition to $D_{\mathcal{A}}$, we also define the total demand of all buyers in $\mathcal{A}$ when they are all satisfied, $D_{\mathcal{A}}^1 = \sum_{j \in \mathcal{A}} d_j(1)$. The p.m.f. and c.d.f. of $D_{\mathcal{A}}^1$ are denoted by $f_{\mathcal{A}}^1(k)$ and $F_{\mathcal{A}}^1(y)$, respectively, and are given by (4.3) and (4.4) for $\alpha_j = 1$.

In our analysis, we will focus on two special subsets of buyers. The first is denoted by $\mathcal{A}_{(i)}$ and contains the first i buyers after all buyers have been reordered in some way so that (i) denotes the index of the i^{th} buyer in the reordered set, i.e., $\mathcal{A}_{(i)} = \{(1), (2), \dots, (i)\}$, $i \in \mathcal{B}$. For this subset, $n_{\mathcal{A}_{(i)}} = i$ and $D_{\mathcal{A}_{(i)}} = \sum_{j=1}^i d_{(j)}(\alpha_{(j)})$. The p.d.f. and c.d.f. of $D_{\mathcal{A}_{(i)}}$ are given by (4.3) and (4.4) for $\mathcal{A} = \mathcal{A}_{(i)}$. For notational

simplicity, $D_{\mathcal{A}_{(i)}}$, $f_{\mathcal{A}_{(i)}}(k)$, and $F_{\mathcal{A}_{(i)}}(y)$ are henceforth denoted by $D_{(i)}$, $f_{(i)}(k)$, and $F_{(i)}(y)$, respectively. By convention, we define $D_{(0)} = 0$, $f_{(0)}(0) = 1$, and $F_{(0)}(\cdot) = 0$. Note that $D_{(n)}$ is the total demand of all buyers.

The second subset is denoted by $\mathcal{A}_{-i}$ and contains all buyers except i , i.e., $\mathcal{A}_{-i} = \mathcal{B} \setminus \{i\}$, $i \in \mathcal{B}$. For this subset, $n_{\mathcal{A}_{-i}} = n - 1$ and $D_{\mathcal{A}_{-i}} = \sum_{j \in \mathcal{B} \setminus \{i\}} d_j(\alpha_j)$. The p.d.f. and c.d.f. of $D_{\mathcal{A}_{-i}}$ are given by (4.3) and (4.4) for $\mathcal{A} = \mathcal{A}_{-i}$. For notational simplicity, $D_{\mathcal{A}_{-i}}$, $f_{\mathcal{A}_{-i}}(k)$, and $F_{\mathcal{A}_{-i}}(y)$ are henceforth denoted by D_{-i} , $f_{-i}(k)$, and $F_{-i}(y)$, respectively.

Next, we define a component-wise ordering of satisfaction state vectors.

Definition 4.1. For two satisfaction state vectors $\alpha' = (\alpha'_1, \dots, \alpha'_n)$ and $\alpha = (\alpha_1, \dots, \alpha_n)$, if $\alpha'_i \geq \alpha_i$, $\forall i \in \mathcal{B}$, we say that α' is greater than (or equal) to α and write $\alpha' \geq \alpha$. If $\alpha' \geq \alpha$, the total demand in α' is stochastically larger than the total demand in α , denoted $D_{(n)}(\alpha') \geq_{st} D_{(n)}(\alpha)$.

After the demand is realized, the firm must select which buyers to serve. We refer to buyers who visit the firm demanding an item as *active* and to those who do not visit as *inactive*. We denote the buyer selection decision by $u_{i,t} \in \{0, 1\}$, where 0 means *do not serve* and 1 means *serve* buyer i . The vector of selection decisions is denoted by $\mathbf{u}_t = (u_{1,t}, u_{2,t}, \dots, u_{n,t})$. For notational simplicity, henceforth, we drop the time index t , wherever possible. Given order quantity $y \in \mathcal{B}_0$ and demand realization $\mathbf{d} \in \{0, 1\}^n$, the action space $\mathcal{U}(y, \mathbf{d})$, representing the possible values of $\mathbf{u}$, is defined as:

$$\mathcal{U}(y, \mathbf{d}) = \left\{ \mathbf{u} \in \{0, 1\}^n : u_i \leq d_i, i \in \mathcal{B}, \sum_{i \in \mathcal{B}} u_i \leq y \right\}. \quad (4.5)$$

The first inequality ensures that inactive buyers are not served. The second inequality is a capacity constraint which states that the number of buyers served cannot exceed the order quantity y .

Active buyers become satisfied if served and dissatisfied if not served. Inactive buyers remain in their previous satisfaction state. Mathematically, this is expressed as:

$$\alpha_{i,t+1} = \varphi(\alpha_{i,t}, d_{i,t}, u_{i,t}), \quad i \in \mathcal{B}, \quad (4.6)$$

where φ is the satisfaction state transition function defined as:

$$\varphi(\alpha_i, d_i, u_i) = u_i + (1 - d_i)\alpha_i, \quad i \in \mathcal{B}, \quad (4.7)$$

and in vector form as:

$$\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u}) = \mathbf{u} + (1 - \mathbf{d}) \circ \boldsymbol{\alpha}, \quad (4.8)$$

where “ $\circ$ ” denotes the Hadamard product (component-wise multiplication).

The short-memory behavior of buyers that we assume is adopted in several models that link demand to past service (e.g., Hall and Porteus (2000); Liu et al. (2007)). In the B2B setting that we consider, it is supported by the finding in Dion and Banting (1995) that multiple stockouts seem not to have serious consequences for buyer loyalty, beyond that of the initial occurrence. This behavior is also consistent with the *peak-end* rule, which suggests that the remembered utility from an experience largely depends on its peak and its end Fredrickson and Kahneman (1993). It is also compatible with the related concept of the *binary bias*, which is the persistent tendency that people have to dichotomize evidence leading to binary perceptions—in our case, satisfied vs. satisfied—Fisher and Keil (2018); Fisher, Newman and Dhar (2018).

The firm pays an acquisition cost c per item ordered at the beginning of each period and receives a revenue r_i from each buyer i that it serves at the end of the period, where $r_i > c$, $i \in \mathcal{B}$. Any unsold items have zero salvage value. Therefore, the profit per period of the firm, denoted by $g(y, \mathbf{u})$, is given by:

$$g(y, \mathbf{u}) = \sum_{i \in \mathcal{B}} r_i u_i - cy. \quad (4.9)$$

The model described above can also be used to represent a firm providing service instead of goods to a group of buyers, e.g., a technical support company that each day must decide the number of technicians on call to respond to requests for in-situ technical support from its clients.

The decision problem of the firm in each period is to select $y \in \mathcal{B}_0$ and $\mathbf{u} \in \mathcal{U}(y, \mathbf{d})$ to maximize its long-run average expected profit, denoted by $\Pi^{y, \mathbf{u}}$. The optimality

equation for this problem can be written as follows:

$$\Pi + V(\boldsymbol{\alpha}) = \max_{y \in \mathcal{B}_0} \left\{ \mathbb{E} \left[\max_{\mathbf{d}} \left[\max_{\mathbf{u} \in \mathcal{U}(y, \mathbf{d})} \{g(y, \mathbf{u}) + V(\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u}))\} \right] \right] \right\}, \quad \forall \boldsymbol{\alpha}, \quad (4.10)$$

where Π is the maximum average expected profit and $V(\boldsymbol{\alpha})$ is the optimal differential profit function starting in $\boldsymbol{\alpha}$. Solving (4.10) in one go is practically impossible because the optimal decisions y^* and $\mathbf{u}^*$ are sequential and interdependent, with y^* depending on $\boldsymbol{\alpha}$ given $\mathbf{u}^*$, and $\mathbf{u}^*$ depending on $\boldsymbol{\alpha}$, y^* , and the realization of $\mathbf{d}$. To unravel the self-reference in (4.10), we can decompose it into the following two subproblems.

Subproblem A: Given ordering policy $y = y(\boldsymbol{\alpha}) \in \mathcal{B}_0$, find the optimal buyer selection policy $\mathbf{u}^{y,*} = \mathbf{u}^*(\boldsymbol{\alpha}, y, \mathbf{d})$ by solving the following Dynamic Programming (DP) problem:

$$\Pi^{y,*} + V^{y,*}(\boldsymbol{\alpha}) = \mathbb{E} \left[\max_{\mathbf{d}} \left[\max_{\mathbf{u} \in \mathcal{U}(y(\boldsymbol{\alpha}), \mathbf{d})} \{g(y(\boldsymbol{\alpha}), \mathbf{u}) + V^{y,*}(\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u}))\} \right] \right], \quad \forall \boldsymbol{\alpha}, \quad (4.11)$$

where superscript “ $y, *$ ” indicates operation under the optimal buyer selection policy for given ordering policy $y = y(\boldsymbol{\alpha})$.

Subproblem B: Given buyer selection policy $\mathbf{u} = \mathbf{u}(\boldsymbol{\alpha}, y, \mathbf{d}) \in \mathcal{U}(y, \mathbf{d})$, find the optimal ordering policy $y^{*,\mathbf{u}} = y^*(\boldsymbol{\alpha}|\mathbf{u})$ by solving the following DP:

$$\Pi^{*,\mathbf{u}} + V^{*,\mathbf{u}}(\boldsymbol{\alpha}) = \max_{y \in \mathcal{B}_0} \left\{ \mathbb{E} [g(y, \mathbf{u}(\boldsymbol{\alpha}, y, \mathbf{d})) + V^{*,\mathbf{u}}(\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u}(\boldsymbol{\alpha}, y, \mathbf{d})))] \right\}, \quad (4.12)$$

$\forall \boldsymbol{\alpha}$, where superscript “ $*, \mathbf{u}$ ” indicates operation under the optimal ordering policy for given buyer selection policy $\mathbf{u} = \mathbf{u}(\boldsymbol{\alpha}, y, \mathbf{d})$.

The optimal decisions y^* and $\mathbf{u}^*$ that solve (4.10) can be obtained by simultaneously solving Subproblems A and B, i.e., y^* and $\mathbf{u}^*$ satisfy $\mathbf{u}^* = \mathbf{u}^{y^*,*} = \mathbf{u}^*(\boldsymbol{\alpha}, y^*, \mathbf{d})$ and $y^* = y^{*,\mathbf{u}^*} = y^*(\boldsymbol{\alpha}|\mathbf{u}^*)$. Solving either of the two subproblems exactly, however, is generally infeasible; therefore, resorting to numerical methods, such as value iteration, is the only viable option. Even so, numerically solving (4.11) becomes computationally intractable as the number of buyers increases, due to the curse of dimensionality and because the second constraint in (4.5) couples the selection decisions across buyers.

The number of computations that must be performed in each value iteration as a function of n , denoted by $N(n)$, is:

$$N(n) = 2^n \sum_{k=0}^n \left[C_k^n \sum_{y=0}^n C_{\min(y,k)}^k \right],$$

where 2^n is the number of satisfaction state vectors, C_k^n is the number of demand vectors for which the total demand is equal to k , and $C_{\min(y,k)}^k$ is the number of possible buyer selection decision vectors when the total demand is k and the order quantity is y . Indicatively, the above formula for different values of n yields: $N(2) = 52$, $N(3) = 312$, $N(5) = 10,336$, $N(7) = 337,280$, $N(10) = 65.71 \times 10^6$, $N(15) = 478.24 \times 10^9$, and $N(25) = 28.44 \times 10^{18}$.

A special class of easy-to-implement buyer selection policies that will play an important role in our analysis is the class of *index* policies. An index policy, denoted by $\mathbf{u}^x$, is a feasible buyer selection policy that assigns to each buyer i an index, denoted by x_i , which is, in general, a function of $\boldsymbol{\alpha}$ and y . The active buyers are served in descending order of their indices until either the order quantity is exhausted or there are no more active buyers to serve. A formal definition of $\mathbf{u}^x$ follows.

Definition 4.2. Under index policy $\mathbf{u}^x$, the buyer selection decision, denoted by $u_{(i)}^x(\boldsymbol{\alpha}, y, \mathbf{d})$, $i \in \mathcal{B}$, is given by:

$$u_{(i)}^x(\boldsymbol{\alpha}, y, \mathbf{d}) = d_{(i)} 1_{\{D_{(i-1)} \leq (y-1)^+\}}, \quad \boldsymbol{\alpha} \in \{0, 1\}^n, \quad y \in \mathcal{B}_0, \quad i \in \mathcal{B},$$

where (i) indexes the buyer with the i^{th} highest index, i.e.,

$$x_{(1)} \geq x_{(2)} \geq \cdots \geq x_{(n)}.$$

As was mentioned earlier, the optimal buyer selection policy $\mathbf{u}^*$ depends in general on $\boldsymbol{\alpha}$, y , and the realization of $\mathbf{d}$. This means that for the same $\boldsymbol{\alpha}$ and y but different realizations of $\mathbf{d}$, the order in which buyers are served may differ, as has been verified by our numerical experiments. Under an index policy $\mathbf{u}^x$, for the same $\boldsymbol{\alpha}$ and y , buyers are always served in the same order for any realization of $\mathbf{d}$. In other words, the

indices are computed before the demand (ex-ante) but are applied after the demand (ex-post). Thus, index policies are suboptimal in general.

Indices can have varying degrees of sophistication. An index that depends only on the parameters and/or satisfaction state of the buyer to which it pertains is referred to as *uncoupled*. An index that also depends on y is referred to as *weakly coupled*, whereas an index that depends on the parameters and/or satisfaction states of all buyers is called *strongly coupled*. The simplest uncoupled index is the revenue rate r_i , which leads to a *revenue-greedy* (or margin-greedy since the order cost c is uniform) buyer selection policy.

4.3 On the optimal policy

This section characterizes the optimal policy for the single-period problem (myopic policy) and provides some properties and conjectures on the optimal policy for the infinite-horizon problem.

4.3.1 Myopic policy

The single-period problem is a newsvendor problem where the newsvendor sells items to n heterogeneous buyers with independent, non-identical Bernoulli demands and different revenue rates. The optimal ordering and buyer selection policy for this problem is henceforth referred to as myopic policy, and the resulting myopic expected profit function is given by the following theorem.

Theorem 4.1. *The myopic buyer selection policy for any order quantity $y \in \mathcal{B}_0$ is an index policy $\mathbf{u}^r$ with index r_i , $i \in \mathcal{B}$. The resulting myopic expected profit, denoted by $G(y)$, is:*

$$G(y) = \sum_{i=1}^y q_{(i)} r_{(i)} + \sum_{i=y+1}^n F_{(i-1)}(y-1) q_{(i)} r_{(i)} - cy, \quad y \in \mathcal{B}_0, \quad (4.13)$$

where (i) indexes the buyer with the i^{th} highest revenue rate, i.e.:

$$r_{(1)} \geq r_{(2)} \geq \cdots \geq r_{(n)}. \quad (4.14)$$

Function $G(\cdot)$ is concave in y , and the myopic ordering policy, denoted by y^m , is given by:

$$y^m = \arg \min_{y \in \mathcal{B}_0 \setminus \{n\}} \left\{ \sum_{i=y+1}^n f_{(i-1)}(y) q_{(i)} r_{(i)} \leq c \right\}. \quad (4.15)$$

If there is no y satisfying (4.15) then $y^m = n$.

The proof is in Appendix A. Theorem 4.1 states that the myopic buyer selection policy is revenue-greedy and the myopic ordering policy is newsvendor-type. As was already mentioned in Section 4.2, computing $f_{(i)}(y)$ is not straightforward. Approximation methods such as the Poisson and normal approximations have been used in the literature, and Hong (2013) has derived an exact formula with a closed-form expression for the c.d.f. of the Poisson binomial distribution.

Corollary 4.1. *The quantity y^m is bounded as follows:*

$$F_{(n)}^{-1} \left(\frac{r_{(n)} - c}{c} \right) \leq y^m \leq F_{(n)}^{-1} \left(\frac{r_{(1)} - c}{c} \right). \quad (4.16)$$

The proof is similar to that of Proposition 2 in Sen and Zhang (1999) and hence is omitted. In all the above expressions, we have suppressed the dependence on α for notational simplicity. In fact, y^m depends on α because $q_{(i)}$ and $f_{(i)}$ in (4.15) are functions of $\alpha_{(i)}$ and $(\alpha_{(1)}, \dots, \alpha_{(i-1)})$, respectively. The following proposition states an important property of $y^m(\alpha)$.

Proposition 4.1. *If $\alpha' \geq \alpha$, then $y^m(\alpha') \geq y^m(\alpha)$.*

The proof is in Appendix A. Proposition 4.1 states that the greater the satisfaction state vector, the larger the myopic order quantity. Proposition 4.1 and constraint (4.16) are useful in reducing the search space of the optimal order quantity, which can be computationally demanding for large n , given the exponential growth of the state space (2^n) and the computational complexity of evaluating f and F , as was

mentioned earlier. We caution that the concavity of $G(y)$ stated in Proposition 4.1 holds only under the optimal buyer selection policy. That is if the firm does not use the revenue-greedy policy to select buyers, $G(y)$ may not be concave, and $y^m(\alpha)$ may not be non-decreasing in α .

4.3.2 Properties of the optimal policy

The myopic policy is appealing because it is simple and focuses on short-term revenue and hence profit. However, it is suboptimal and can be arbitrarily bad for the infinite-horizon problem because it ignores the effect of decisions on buyer satisfaction and future demand. Nevertheless, the monotonicity property of the myopic order quantity stated in Proposition 4.1 is fundamental and should hold beyond the confines of the single-period problem. This intuition is supported by the following proposition.

Proposition 4.2. *If $\alpha' \geq \alpha$, then $V(\alpha') \geq V(\alpha)$.*

The proof is in Appendix A. It is based on the sample-path argument that the firm will perform better if it starts from state α' and follows the optimal policy starting from α , than if it starts from α . Proposition 4.2 suggests the following analog to Proposition 4.1 for the infinite-horizon problem.

Conjecture 1. *If $\alpha' \geq \alpha$, then $y^*(\alpha') \geq y^*(\alpha)$.*

Conjecture 1 is based on the fact that $\alpha' \geq \alpha$ implies $D_{(n)}(\alpha') \geq_{st} D_{(n)}(\alpha)$ by Definition 4.1. The claim is that by ordering at least as many items in state α' as in state α , the firm can reap a higher revenue and avoid dissatisfying buyers who are more likely to be active in α' than they are in α . In the case of homogeneous buyers, where selection is not an issue, Deng et al. (2014) proves this monotonicity property by approximating the value function with a linear function in the number of satisfied buyers, since this approximation is exact for infinite buyers. If the buyers are heterogeneous, the problem is more complicated, due to the buyer selection decision. In this case, the monotonicity property makes sense only under the optimal selection policy, as was also the case with Proposition 4.1. Conjecture 1 is verified in all our numerical examples with three or more buyers and is proved for two buyers in Section

4.4.1. In Section 4.4.2, we provide a numerical counterexample with two buyers where the firm does not use the optimal selection policy and Conjecture 1 does not hold.

When considering the long-term problem, the firm must balance the current revenue from the satisfied buyers against the loss in future demand from the dissatisfied customers. The fact that the buyers have short memory suggests that the optimal selection should depend more on how decisions affect the near future than on how they affect the distant future. Intuitively, between two policies leading to different satisfaction state vectors, the policy that leads to the state where the total demand is stochastically larger should be preferable as far as the long-term average expected profit of the firm is concerned.

The following proposition states that for any two active buyers i and j competing for one item, if i has higher LVC and GDC than j , meaning that i is more reactive than j to quality-of-service changes, prioritizing i over j leads to a satisfaction state vector in which the total demand of both buyers is stochastically larger than the respective demand in the state vector led to by prioritizing j over i .

Proposition 4.3. *For any two active buyers $i, j \in \mathcal{B}$, $i \neq j$, if $\gamma_i \geq \gamma_j$ and $\bar{\gamma}_i \geq \bar{\gamma}_j$, then $d_i(1) + d_j(0) \geq_{st} d_i(0) + d_j(1)$.*

The proof is in Appendix A. It is trivial to show that $\gamma_i \geq \gamma_j$ and $q_i(1) \geq q_j(1)$ imply $\bar{\gamma}_i \geq \bar{\gamma}_j$; that is, if buyer i has a higher LVC and visits the firm more frequently when she is satisfied than j does, then i also has a higher GDC than j . Therefore, prioritizing i over j leads to a greater satisfaction state vector. If, in addition to being more reactive to service changes, i also has a higher revenue rate than j , then intuitively the firm is better off prioritizing i over j . This intuition is expressed in the following conjecture.

Conjecture 2. *For any two active buyers $i, j \in \mathcal{B}$, $i \neq j$, if $r_i \geq r_j$, $\gamma_i \geq \gamma_j$, and $\bar{\gamma}_i \geq \bar{\gamma}_j$, then $u_i^* \geq u_j^*$.*

Conjecture 2 is a useful rule of thumb for businesses where the higher the revenue rate of a buyer, the more important the buyer, the higher her service expectations, and therefore the higher her LVC and GDC. For cases where the more frequent the

buyer, the lower the price (revenue rate) she pays to the firm, conditions of Conjecture 2 do not hold, and buyer selection becomes more obscure.

In general, the optimal selection policy is not an index policy and cannot be obtained in closed form. We can only track it down for special cases, hoping to gain some intuition that can lead us to develop good heuristics. The following proposition characterizes the optimal selection policy for the special case where the firm uses a *fixed order quantity (FOQ)* policy with FOQ equal to $n - 1$.

Proposition 4.4. *If $y(\alpha) = n - 1$, $\forall \alpha$, the optimal buyer selection policy $\mathbf{u}^*$ is an index policy $\mathbf{u}^z$ with index for buyer i given by:*

$$z_i = \frac{r_i}{1 - \gamma_i F_{-i}^1(n - 2)}, \quad i \in \mathcal{B}, \quad (4.17)$$

The resulting maximum average expected profit, denoted by $\Pi^{n-1,*}$, is:

$$\Pi^{n-1,*} = R - R_j - (n - 1)c, \quad (4.18)$$

where R and R_j are given below:

$$R = \sum_{k \in \mathcal{B}} q_k(1)r_k, \quad (4.19)$$

$$R_j = \prod_{k \in \mathcal{B}} q_k(1)z_j, \quad \text{where } j = \arg \min_{k \in \mathcal{B}}(z_k). \quad (4.20)$$

The proof is in Appendix A. It is based on recognizing that if $y(\alpha) = n - 1$, $\forall \alpha$, buyer selection matters only when all buyers are active. In this case, the question is not who should be selected, but who should be assigned the lowest priority and be left out. If buyer j has the lowest priority, she will become dissatisfied, and all other buyers will be satisfied. Thereafter, j will not contribute to the firm's revenue until she becomes satisfied. For this to happen, the total demand of the other $n - 1$ buyers must not exceed $n - 2$. The term $R - R_j$ in (4.18) expresses the difference between the total expected revenue of the firm if all buyers are always served, and the lost revenue from buyer j every time she is not served—although active—because

all other buyers are also active and j has the lowest priority.

From (4.17), z_i is an increasing function of three terms: r_i , γ_i , and $F_{-i}^1(n-2)$. The first term is the *revenue* from buyer i , if buyer i is served. The second is the *loss in the future demand* of buyer i , if buyer i is not served. Both terms refer to buyer i 's parameters. The last term is the *type-I service level of all other buyers* when they are satisfied, if buyer i is served, i.e., it is the probability that the total demand of the other $n-1$ buyers is at most $n-2$. From (4.3) and (4.4), this probability is equal to $1 - \prod_{k \in \mathcal{B} \setminus \{i\}} q_k(1)$. This term couples the demand of buyer i to the demand of all other buyers. Note that if $q_i(1) > q_j(1)$, then $F_{-i}^1(n-2) > F_{-j}^1(n-2)$, for $i \neq j$. This means that by favoring a buyer with a higher $q_i(1)$, the firm reduces the probability of stockout for the other buyers, thus increasing the chance of the low-priority buyer being satisfied and contributing to profit. Proposition 4.4 leads to the following property.

Corollary 4.2. *If $r_i \geq r_j$, $\gamma_i \geq \gamma_j$, and $\bar{\gamma}_i \geq \bar{\gamma}_j$, then $z_i \geq z_j$, $i \neq j$.*

Corollary 4.2 confirms Conjecture 2 when $y(\alpha) = n-1$, $\forall \alpha$. Note that the three conditions that it sets are sufficient. This means that it is possible that $z_i \geq z_j$ even if not all these conditions are met. From (4.17), the necessary condition for $z_i \geq z_j$ is $r_i[1 - \gamma_j(1 - q_i(1)Q)] > r_j[1 - \gamma_i(1 - q_j(1)Q)]$, where $Q = \prod_{k \in \mathcal{B} \setminus \{i,j\}} q_k(1)$.

Index z_i depends on the parameters of buyer i , the order quantity $n-1$, and the visit rates of all other buyers. It is therefore strongly coupled. As $n \rightarrow \infty$, $F_{-i}^1(n-2) \rightarrow 0$, so the limit of z_i as $n \rightarrow \infty$, denoted by s_i , becomes:

$$s_i = \lim_{n \rightarrow \infty} z_i = \frac{r_i}{1 - \gamma_i}, \quad i \in \mathcal{B}. \quad (4.21)$$

Index s_i augments the revenue rate r_i by a factor of $1/(1 - \gamma_i)$. Unlike z_i which depends on the visit rates of all buyers, s_i depends on the parameters of buyer i only, so it is uncoupled, like r_i . Therefore, increasing the number of buyers decreases the coupling between them. Unlike r_i which is myopic, s_i is far-sighted because it accounts for the loss in future demand. We refer to s_i as the *augmented revenue rate* and to the index policy that results from using s_i as the *augmented-revenue-greedy* policy $\mathbf{u}^s$.

4.4 The two-buyer problem

In this section, we characterize the optimal ordering and buyer selection policy for two buyers, and we analyze and compare the performance of any arbitrary index policy to that of the optimal selection policy.

4.4.1 Optimal policy

From Proposition 4.4, the optimal selection policy for two buyers, under the FOQ policy $y(\alpha) = 1, \forall \alpha$, is $\mathbf{u}^z$, where z_i given by (4.17) for $n = 2$. The following theorem states that the optimal ordering policy for two buyers is an FOQ policy, so when the optimal FOQ is 1, the optimal selection policy is $\mathbf{u}^z$; otherwise (when the optimal FOQ is 0 or 2), buyer selection is not an issue.

Theorem 4.2. *For $n = 2$ ($\mathcal{B} = \{1, 2\}$):*

- (a) *The optimal buyer selection policy $\mathbf{u}^*$ is an index policy $\mathbf{u}^z$, with index for buyer i given by:*

$$z_i = \frac{r_i}{1 - \gamma_i \bar{q}_j(1)} = \frac{r_i}{1 - \beta_i \bar{q}_i(1) \bar{q}_j(1)}, \quad i \in \mathcal{B}, \quad (4.22)$$

where

$$\beta_i = \gamma_i + \bar{\gamma}_i = \frac{q_i(1) - q_i(0)}{q_i(1) \bar{q}_i(1)}, \quad i \in \mathcal{B}. \quad (4.23)$$

- (b) *The optimal ordering policy y^* is an FOQ policy. Assuming without loss generality that $z_i > z_j, i \neq j$, the optimal FOQ, denoted by y^z , and the resulting maximum average expected profit Π are given in Table 4.1, where R and R_j are defined in (4.19) and (4.20).*

Table 4.1: Optimal FOQ and resulting average expected profit for $z_i > z_j, i \neq j$, and $n = 2$.

Region	Condition	y^z	Π
Y0	$R_j > R - c$	0	0
Y1 _i	$R_j < \min(R - c, c)$	1	$R - R_j - c$
Y2	$R_j > c$	2	$R - 2c$

The proof is in Appendix A. The first expression for z_i in (4.22) is the same as the expression in (4.17) for $n = 2$, after noting that $\bar{q}_j(1) = F_{-i}^1(0)$; therefore, it has the same interpretation as that expression. The second expression allows for a different interpretation. It implies that z_i is increasing in r_i and β_i , where β_i is the ratio of the drop of supplier i 's visit rate if she is active but not served, $q_i(1) - q_i(0)$, to the variance of her demand when she is satisfied, $q_i(1)\bar{q}_i(1)$ (note that $q_i(1)\bar{q}_i(1)$ is maximized at $q_i(1) = 0.5$). Therefore, prioritizing the buyers based on β_i implies favoring a buyer who is *more reactive to quality-of-service changes* and has a *more predictable visit behavior when satisfied*.

Theorem 4.2 (b) confirms Conjecture 1. The conditions in Table 4.1 under which each FOQ value is optimal when $z_i > z_j$, have the form of inequalities involving the expected revenues R and R_j defined in (4.19) and (4.20), respectively. These conditions partition the (z_i, z_j, R, R_j) space into three regions, denoted Y0, Y1_{*i*}, and Y2, where $y^z = 0, 1$, and 2, respectively. The conditions imply that region Y1_{*i*} borders with Y0 and Y2, but Y0 and Y2 do not share a border. The interpretation of the optimal FOQ policy is simple and intuitive. If $\Pi^{0,\mathbf{u}^z} > \Pi^{1,\mathbf{u}^z}$, then $y^z = 0$ (region Y0). If $\Pi^{2,\mathbf{u}^z} > \Pi^{1,\mathbf{u}^z}$, then $y^z = 2$ (region Y2). Otherwise, $y^z = 1$ (region Y1_{*i*}). Figure 4.1 displays graphs of two indicative problem instances with different sets of visit and deferral rates, showing the regions, projected onto the (r_1, r_2) space, where each FOQ value is optimal, under $\mathbf{u}^z$. In both instances, $c = 1$. Note that in region




Figure 4.1: Optimal FOQ policy under the optimal index policy $\mathbf{u}^z$ for $n = 2$ ($c = 1$).

$Y1_1$, $z_1 > z_2$, whereas in region $Y1_2$, $z_2 > z_1$. In the left graph, $\beta_1 > \beta_2$, so region $Y1_1$ covers a part of the space between the lines $z_1 = z_2$ and $r_1 = r_2$, where buyer 1 has priority over 2 (because $z_1 > z_2$) even though $r_1 < r_2$. The reverse is true in the right graph, where region $Y1_2$ covers a part of the space where $r_1 > r_2$.

4.4.2 Effect of the index on the optimal FOQ and the average expected profit

In the previous section, we saw that for two buyers, the optimal buyer selection policy is index policy $\mathbf{u}^z$ and the optimal ordering policy is FOQ policy y^z . In this section, we examine how the performance of an arbitrary index policy $\mathbf{u}^x$ compares to that of $\mathbf{u}^z$ and what the resulting optimal FOQ, y^x , is. This investigation is of interest because the result for $n = 2$ can be an indication of the outcome for $n > 2$, where the optimal policy is not tractable and different heuristic index policies may be considered. For two buyers, when $y^z = 0$ or 2, buyer priority is irrelevant; therefore,

Table 4.2: Optimal FOQ and resulting average expected profit under $\mathbf{u}^x$; difference in average expected profit under $\mathbf{u}^z$ and $\mathbf{u}^x$ when $y^z = 1$, for $x_i > x_j$, $z_i < z_j$, $i \neq j$, and $n = 2$.

Region	Condition	y^x	Area	$\Pi^{*,\mathbf{u}^x}$	$\Pi - \Pi^{*,\mathbf{u}^x}$
Y0	$R_i < R - c < \min(R_j, c)$	0	Blue	0	$R - R_i - c$
$Y1_i$	$R_i < R_j < \min(R - c, c)$	1	Gray	$R - R_j - c$	$R_j - R_i$
Y2	$R_i < c < R_j$ and $R > 2c$	2	Red	$R - 2c$	$c - R_i$

$\mathbf{u}^x$ has the same optimal FOQ and average expected profit as $\mathbf{u}^z$, i.e., $y^x = y^z$ and $\Pi^{*,\mathbf{u}^x} = \Pi$. Policies $\mathbf{u}^x$ and $\mathbf{u}^z$ differ only when $y^z = 1$ and the two policies prioritize buyers in the reverse order. In this case, y^x may be either 0, 1, or 2, and $\Pi^{*,\mathbf{u}^x} < \Pi$. Table 4.2 displays y^x , $\Pi^{*,\mathbf{u}^x}$, and $\Pi - \Pi^{*,\mathbf{u}^x}$ when $y^z = 1$ and $\mathbf{u}^x$ and $\mathbf{u}^z$ prioritize buyers in the reverse order.

Two index policies that are of particular interest are $\mathbf{u}^r$ and $\mathbf{u}^s$ because they capture key buyer attributes (profitability and reactivity) and arise in Lagrangian relaxation approximations of the original problem, as we will see in Section 4.7.

Figures 4.2 and 4.3 display graphs for the same problem instances as those in



Figure 4.2: Optimal FOQ policy under index policy $\mathbf{u}^r$ for $n = 2$ ($c = 1$).Figure 4.3: Optimal FOQ policy under index policy $\mathbf{u}^s$ for $n = 2$ ($c = 1$).

Figure 4.1, showing the regions, projected onto the (r_1, r_2) space, where different FOQ values are optimal under $\mathbf{u}^x$, for $x = r$ and $x = s$, respectively. Note that in the right graph of both figures, regions $Y0$ and $Y2$ share a border, whereas under $\mathbf{u}^z$ in Figure 4.1, these regions do not communicate, as was pointed out earlier.

In both figures, the areas where $\mathbf{u}^z$ outperforms $\mathbf{u}^x$ are shown in color, as defined in Table 4.2. These areas cover a strip between the lines $z_1 = z_2$ and $x_1 = x_2$ ($r_1 = r_2$ in Figure 4.2 and $s_1 = s_2$ in Figure 4.3) where $\mathbf{u}^x$ and $\mathbf{u}^z$ prioritize buyers in the reverse order and $y^z = 1$. In both figures, the areas where $\mathbf{u}^z$ outperforms $\mathbf{u}^x$ are shown in color, as defined in Table 4.2. These areas cover a strip between the lines $z_1 = z_2$ and $x_1 = x_2$ ($r_1 = r_2$ in Figure 4.2 and $s_1 = s_2$ in Figure 4.3) where $\mathbf{u}^x$ and $\mathbf{u}^z$ prioritize

In the above analysis, we restricted our search for the optimal ordering policy to FOQ policies. In the proof of Theorem 4.2, we show that the invariance of the optimal order quantity arises from the monotonicity property of $y(\boldsymbol{\alpha})$ which holds only under the optimal index policy $\mathbf{u}^z$. That is, if a suboptimal index policy $\mathbf{u}^x \neq \mathbf{u}^z$ is used, then $\boldsymbol{\alpha}' \geq \boldsymbol{\alpha}$ does not imply that $y^x(\boldsymbol{\alpha}') \geq y^x(\boldsymbol{\alpha})$, and therefore $y^x(\boldsymbol{\alpha})$ is not necessarily fixed. We explain this counterintuitive behavior with an example. Consider a problem instance with $c = 1$, $r_1 = 1.1$, $r_2 = 1.05$, $q_1(1) = 0.2$, $q_1(0) = 0.1$, $q_2(1) = 0.98$, and $q_2(0) = 0.8$. Note that $r_1 > r_2$, $\gamma_1 = 0.5 > 0.184 = \gamma_2$, and $\bar{\gamma}_1 = 0.125 < 9 = \bar{\gamma}_2$. Suppose the firm uses the revenue-greedy policy $\mathbf{u}^r$, which means that it prioritizes buyer 1. Numerically solving problem (4.12) under $\mathbf{u}^r$, yields the optimal ordering policy $y^r(\boldsymbol{\alpha})$ shown for each state $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$ in the transition state diagram in Figure 4.4. The components $q_i(\alpha_i)$ of each transition probability appear in green or red color depending on whether active buyer i is selected or is left out of service, respectively, in the corresponding transition. The components $\bar{q}_i(\alpha_i)$ of inactive buyers appear in black. Self-transitions are omitted.




Figure 4.4: State transition diagram (omitting self-transitions) corresponding to the optimal ordering policy $y^r(\alpha)$ under the revenue-greedy policy $\mathbf{u}^r$, for a two-buyer example with $r_1 > r_2$, $\gamma_1 > \gamma_2$, and $\bar{\gamma}_1 < \bar{\gamma}_2$.

4.5 Current revenue vs. loss in future demand

In the previous section, we presented the optimal ordering and buyer selection policy for two buyers. Generalizing this analysis to more buyers, however, is impossible. To gain insight into the optimal policy for more than two buyers, in this section, we analyze the optimal policy for an example with three buyers, which we compute by numerically solving the optimality equation (4.10).

The parameters of the example are: $\mathcal{B} = \{1, 2, 3\}$, $c = 1$, $r_1 = 1.3$, $r_2 = 1.25$, $r_3 = 1.2$, $q_1(1) = q_2(1) = 0.66$, $q_1(0) = q_2(0) = 0.33$, $q_3(1) = 0.93$, and $q_3(0) = 0.46$. Note that $r_1 > r_2 > r_3$ and, from (4.1) and (4.2), $\gamma_1 = \gamma_2 = \gamma_3 = 0.5$ and $\bar{\gamma}_1 = \bar{\gamma}_2 = 1 < 7 = \bar{\gamma}_3$. Figure 4.5 shows the state transition diagram corresponding to the optimal policy. The values of the optimal ordering policy $y^*(\alpha)$ and the maximum average expected profit, denoted by $G^*(\alpha)$, computed as $G^*(\alpha) = \sum_{i \in \{1,2,3\}} r_i \mathbb{E}_{\mathbf{d}}[u_i^*(\alpha, y^*(\alpha), \mathbf{d})] - cy(\alpha)$, are shown next to each state $\alpha = (\alpha_1, \alpha_2, \alpha_3)$. The optimal buyer selection policy $\mathbf{u}^*(\alpha, y^*(\alpha), \mathbf{d})$ is implied by the transition probabilities. As in Figure 4.4, each component q_i of each transition probability appears in green or red color depending on whether active buyer i is selected or is left out of service in the respective transition. For notational simplicity, the dependence of y^* and G^* on α and of q_i on α_i is omitted. From Figure 4.5, we observe that $y^*(\alpha) = 1$, $\alpha \neq (1, 1, 1)$, and $y^*(1, 1, 1) = 2$, confirming Conjecture 1. We also observe that buyer 1 is always prioritized over buyer 2, because both buyers have the same visit rates, and hence the same LVC and GDC, but $r_1 > r_2$. In general, however, $\mathbf{u}^*$ depends on α , y , and the realization of $\mathbf{d}$; therefore, it is not an index policy. More specifically, if the ending state after $\mathbf{d}$ is realized is one where two buyers are satisfied and one buyer is dissatisfied, henceforth referred to as a *good state*, the optimal selection policy is revenue-greedy and hence myopic. This means that the firm opts for state $(1, 1, 0)$ then $(1, 0, 1)$ and lastly $(0, 1, 1)$, despite the fact that $G^*(0, 1, 1) > G^*(1, 0, 1) > G^*(1, 1, 0)$. The transitions that lead to states $(1, 1, 0)$ and $(1, 0, 1)$ are indicated with blue arrows. On the other hand, if the ending state is one where one buyer is satisfied and two buyers are dissatisfied, henceforth referred to as a *bad state*, the optimal selection policy is $\bar{\gamma}_i$ -greedy and hence far-sighted. This means that the firm opts for state $(0, 0, 1)$ and




Figure 4.5: State transition diagram (omitting self-transitions) corresponding to the optimal ordering and buyer selection policy, y^* and u^* , for a three-buyer example with $r_1 > r_2 > r_3$, $\gamma_1 = \gamma_2 = \gamma_3$ and $\bar{\gamma}_1 = \bar{\gamma}_2 < \bar{\gamma}_3$.

then states $(1, 0, 0)$ and $(0, 1, 0)$. The transitions into state $(0, 0, 1)$ are indicated with orange arrows. If the only option is between the last two states, then the firm opts for $(1, 0, 0)$ because it always prioritizes buyer 1 over 2, as was mentioned earlier. The sole transition into state $(1, 0, 0)$ is indicated with a green arrow. States $(0, 1, 0)$ and $(0, 0, 0)$ are transient and hence omitted. To better understand how the optimal selection policy works, consider the situation where the initial state is $(1, 0, 1)$. If buyers 2 and 3 are active and 1 is inactive, the firm will end up in a good state, no matter which buyer it selects. Under this demand scenario, the firm uses a revenue-greedy policy seeking to maximize the current revenue, so it selects buyer 2 over 3 and ends up in state $(1, 1, 0)$. If, on the other hand, buyers 1 and 3 are active and buyer 2 is inactive, the firm will end up in a bad state, no matter which buyer it selects. In this scenario, the firm uses a $\bar{\gamma}_i$ -greedy policy seeking to maximize future demand, so it selects buyer 3 over 1 and ends up in state $(0, 0, 1)$.

4.6 Lagrangian relaxation

As was mentioned at the end of Section 4.2, to solve the optimality equation (4.10), we can decompose it into two subproblems: Subproblem A given by (4.11) and Subproblem B given by (4.12). Solving Subproblem A exactly is intractable because of capacity constraint (4.5) which couples the selection decisions across buyers. If it were not for this constraint, each buyer could be analyzed independently. As a result, Subproblem A fits the definition of a *weakly coupled* DP problem Adelman and Mersereau (2013); Bertsimas and Mišić (2016) and is amenable to decomposition via relaxation.

In this section, we consider a Lagrangian relaxation of (4.11) that is obtained by relaxing the coupling constraint (4.5) and adding the penalty term $\lambda (y(\boldsymbol{\alpha}) - \sum_{i \in \mathcal{B}} u_i)$ to the objective function, where $\lambda \geq 0$ is the Lagrange multiplier or penalty price for violating (4.5). After rearranging terms, the relaxed problem becomes:

$$\hat{\Pi}^{y,\lambda} + \hat{V}^\lambda(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{d}} \left[\max_{\mathbf{u} \in \mathcal{U}(\mathbf{d})} \left\{ \sum_{i \in \mathcal{B}} (r_i - \lambda) u_i + \hat{V}^\lambda(\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u})) \right\} \right] + (\lambda - c) y(\boldsymbol{\alpha}), \quad (4.24)$$

$$\mathcal{U}(\mathbf{d}) = \{\mathbf{u} \in \{0, 1\}^n : u_i \leq d_i, i \in \mathcal{B}\}.$$

The term $\lambda (y(\boldsymbol{\alpha}) - \sum_{i \in \mathcal{B}} u_i)$ that has been added to (4.11) is non-negative for policies satisfying (4.5). This implies that $\Pi^{y,*} \leq \hat{\Pi}^{y,\lambda}$, and hence $\hat{\Pi}^{y,\lambda}$ is an upper bound for $\Pi^{y,*}$. Note that $\hat{\Pi}^{y,\lambda}$ depends on $y(\boldsymbol{\alpha})$ because of the last term in the r.h.s. of (4.24). To remove this dependence, we define:

$$\hat{\Pi}^\lambda = \hat{\Pi}^{y,\lambda} - (\lambda - c) y(\boldsymbol{\alpha}), \quad (4.25)$$

so that (4.24) can be written as:

$$\hat{\Pi}^\lambda + \hat{V}^\lambda(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{d}} \left[\max_{\mathbf{u} \in \mathcal{U}(\mathbf{d})} \left\{ \sum_{i \in \mathcal{B}} (r_i - \lambda) u_i + \hat{V}^\lambda(\Phi(\boldsymbol{\alpha}, \mathbf{d}, \mathbf{u})) \right\} \right], \quad \forall \boldsymbol{\alpha}. \quad (4.26)$$

The advantage of Lagrangian relaxation is that for any fixed λ , DP (4.26) can

be decomposed into the sum of n buyer-specific DP problems that can be solved independently, as stated in the following Proposition.

Proposition 4.5. *For any $y(\alpha) > 0$, $\alpha \in \{0, 1\}^n$, and $\lambda \geq 0$:*

$$\hat{V}^\lambda(\alpha) = \sum_{i \in \mathcal{B}} \hat{V}_i^\lambda(\alpha_i), \quad (4.27)$$

$$\hat{\Pi}^\lambda = \sum_{i \in \mathcal{B}} \hat{\Pi}_i^\lambda, \quad (4.28)$$

where $\hat{V}_i^\lambda(\alpha_i)$ and $\hat{\Pi}_i^\lambda$ solve the following buyer-specific DP:

$$\hat{\Pi}_i^\lambda + \hat{V}_i^\lambda(\alpha_i) = \mathbb{E} \left[\max_{d_i} \left\{ (r_i - \lambda)u_i + \hat{V}_i^\lambda(u_i + (1 - d_i)\alpha_i) \right\} \right], \quad i \in \mathcal{B}, \quad (4.29)$$

$$\mathcal{U}_i(d_i) = \{u_i \in \{0, 1\} : u_i \leq d_i\}, \quad i \in \mathcal{B}.$$

The proof is in Appendix A. The first term in the maximization of (4.29) is the revenue generated by buyer i when she is active. It is non-negative if $r_i \geq \lambda$ and negative if $r_i < \lambda$, leading to a simple solution provided by the following proposition.

Proposition 4.6. *The solution of (4.29), denoted by $u_i^\lambda(d_i)$, $i \in \mathcal{B}$, is given by:*

$$u_i^\lambda(d_i) = d_i 1_{\{r_i \geq \lambda\}}, \quad i \in \mathcal{B}. \quad (4.30)$$

$$\hat{V}_i^\lambda(1) = (r_i - \lambda)^+ \frac{\gamma_i}{1 - \gamma_i}, \quad \hat{V}_i^\lambda(0) = 0, \quad i \in \mathcal{B}, \quad (4.31)$$

$$\hat{\Pi}_i^\lambda = (r_i - \lambda)^+ q_i(1), \quad i \in \mathcal{B}. \quad (4.32)$$

The proof is in Appendix A. The optimal buyer selection policy given by (4.30) depends on the choice of the Lagrange multiplier λ . As was mentioned earlier, $\hat{\Pi}^{y, \lambda}$ is an upper bound for $\Pi^{y, *}$. To obtain the tightest bound, we must solve the Lagrangian dual problem Topaloglu (2009); Brown and Smith (2020):

$$\min_{\lambda \geq 0} \hat{\Pi}^{y, \lambda}. \quad (4.33)$$

Proposition 4.7. *The solution of (4.33) is:*

$$\lambda^* = r_{(i^*)}, \text{ where } i^* = \arg \max_{i=1, \dots, n+1} \left\{ y(\alpha) \geq \sum_{k=1}^{i-1} q_{(k)}(1) \right\}, \quad (4.34)$$

where (i) indicates the index of the buyer with the i^{th} highest revenue rate, and by convention $r_{(n+1)} = 0$.

The proof is in Appendix A. Expression (4.34) implies that λ^* is the $(i^*)^{\text{th}}$ highest revenue rate, where i^* is the first buyer whose demand is not covered by $y(\alpha)$ if the firm prioritizes buyers in descending order of their revenue rates and replaces their demands by their expected value in the satisfied state. Figure 4.6 displays a graph of λ^* vs. $y(\alpha)$. If $y(\alpha) = \sum_{k=1}^{i-1} q_{(k)}(1)$, for some i , then $\lambda^* \in [r_{(i)}, r_{(i-1)}]$. In this case, to break the tie, we set $\lambda^* = r_{(i)}$.




Figure 4.6: Optimal Lagrange multiplier λ^* vs. $y(\alpha)$.

4.7 Index policies

The buyer selection policy given by (4.30) is not feasible because it violates the linking constraint (4.5) when $\sum_{i \in \mathcal{B}} d_i 1_{\{r_i \geq \lambda\}} > y(\alpha)$. In this section, we consider three heuristic index policies that respect constraint (4.5). Following Brown and Smith (2020), the index in each policy approximates the value added to the firm by

selecting buyer i when she is active, given by:

$$v_i = r_i + W_i(1) - W_i(0), \quad i \in \mathcal{B}, \quad (4.35)$$

where $W_i(\alpha_i)$ is some buyer-specific approximation of the value function, e.g., $\hat{V}_i^\lambda(\alpha_i)$.

4.7.1 Whittle index policy

The first index policy that we consider is the Whittle index policy, where $W_i(\alpha_i)$ is the value function $\hat{V}_i^\lambda(\alpha_i)$ in (4.29), for a given buyer-specific Lagrange multiplier $\lambda = w_i$. The Whittle index is the value of w_i which, if given to the firm as a subsidy, makes it indifferent between selecting vs. not selecting buyer i Whittle (1988). From (4.35), w_i satisfies:

$$w_i = r_i + \hat{V}_i^{w_i}(1) - \hat{V}_i^{w_i}(0), \quad i \in \mathcal{B}.$$

Substituting $\hat{V}_i^{w_i}(1)$ and $\hat{V}_i^{w_i}(0)$ from (4.31) into the above expression yields:

$$w_i = r_i + (r_i - w_i)^+ \frac{\gamma_i}{1 - \gamma_i}, \quad i \in \mathcal{B}.$$

The solution of the above equation is:

$$w_i = r_i, \quad i \in \mathcal{B}. \quad (4.36)$$

Therefore, the Whittle index of buyer i is r_i , and the Whittle index policy coincides with the revenue-greedy policy, as is also shown in Adelman and Mersereau (2013) in a similar setting.

4.7.2 Lagrangian index policy

The second index policy that we consider is the Lagrangian index policy, where $W_i(\alpha_i)$ is the value function $\hat{V}_i^\lambda(\alpha_i)$ in (4.29) for a given Lagrange multiplier λ which is common for all buyers Brown and Smith (2020). From (4.35), the Lagrangian index,

denoted by l_i satisfies:

$$l_i = r_i + \hat{V}_i^\lambda(1) - \hat{V}_i^\lambda(0), \quad i \in \mathcal{B}.$$

Substituting $\hat{V}_i^\lambda(1)$ and $\hat{V}_i^\lambda(0)$ from (4.31) yields:

$$l_i = r_i + (r_i - \lambda)^+ \frac{\gamma_i}{1 - \gamma_i}, \quad i \in \mathcal{B}. \quad (4.37)$$

The above index can also be viewed as a greedy index w.r.t. $\hat{V}_i^\lambda(\alpha_i)$ that is derived by solving the following optimization problem Adelman and Mersereau (2013):

$$\mathbb{E}_d \left[\max_{\mathbf{u} \in \mathcal{U}(y(\boldsymbol{\alpha}), \mathbf{d})} \left\{ \sum_{i \in \mathcal{B}} r_i u_i + \hat{V}_i^\lambda(u_i + (1 - d_i)\alpha_i) \right\} \right].$$

The above problem is equivalent to a 0-1 knapsack problem that is solved by selecting buyers in descending order of indices l_i given by (4.37). Although we can use any Lagrange multiplier λ in (4.37), we expect that values leading to tighter performance bounds result in better approximate value functions and generate better heuristics. For this reason, we use the optimal multiplier λ^* from (4.34) to obtain the optimal Lagrangian index:

$$l_{(i)}^* = r_{(i)} + (r_{(i)} - r_{(i^*)})^+ \frac{\gamma_{(i)}}{1 - \gamma_{(i)}}, \quad i \in \mathcal{B}, \quad (4.38)$$

where (i) indicates the index of the buyer with the i^{th} highest revenue rate. Based on (4.38), the buyers are divided into two groups: those with the $i^* - 1$ highest revenue rates and those with the $n - i^* + 1$ lowest revenue rates. The index for each buyer in the first group depends on her goodwill and is constructed by augmenting her revenue rate $r_{(i)}$ by a term that is proportional to $r_{(i)} - r_{(i^*)}$ and $\gamma_{(i)}/(1 - \gamma_{(i)})$. The index for each buyer in the second group is $r_{(i)}$. Note that $l_{(i)}^*$ depends on $y(\boldsymbol{\alpha})$, because from (4.34), i^* depends on $y(\boldsymbol{\alpha})$; therefore, it is weakly coupled. Figure 4.7 shows a plot of $l_{(i)}^*$ in (4.38) vs. (i) .

A question that comes to mind is, how does the optimal Lagrangian index policy




Figure 4.7: Optimal Lagrangian index $l_{(i)}^*$ vs. buyer index (i) using the optimal Lagrange multiplier λ^* .

compare to the optimal buyer selection policy? To address this question, we compute $l_{(i)}^*$ for the FOQ policy $y(\alpha) = n - 1$, $\forall \alpha$, for which we know from Proposition 4.4 that the optimal selection policy is index policy u^z where the optimal index $z_{(i)}$ is given by (4.17) and, in the case of two buyers, (4.22).

Corollary 4.3. *If $y(\alpha) = n - 1$, $\forall \alpha$, the optimal Lagrangian index for buyer i is given by:*

$$l_{(i)}^* = \begin{cases} s_{(i)}, & \text{if } \sum_{k=1}^n q_{(k)}(1) \leq n - 1, \\ r_{(i)} + (r_{(i)} - r_{(n)}) \frac{\gamma_{(i)}}{1 - \gamma_{(i)}}, & \text{otherwise,} \end{cases} \quad i \in \mathcal{B}, \quad (4.39)$$

where (i) indicates the index of the buyer with the i^{th} highest revenue rate. For $n = 2$ ($\mathcal{B} = \{1, 2\}$), (4.39) is equivalent to:

$$l_{(i)}^* = \begin{cases} s_{(i)}, & \text{if } q_1(1) + q_2(1) \leq 1, \\ r_{(i)}, & \text{otherwise,} \end{cases} \quad i \in \mathcal{B}. \quad (4.40)$$

The proof is in Appendix A. By comparing $l_{(i)}^*$ in (4.39) and $z_{(i)}$ in (4.17) it is obvious that the two indices differ, although both are increasing in $r_{(i)}$ and $\gamma_{(i)}$. It can be shown that $l_{(i)}^* \leq z_{(i)}$ if $\sum_{k=1}^n q_{(k)}(1) > n - 1$ and $r_{(i)}/r_{(n)} \leq \gamma_{(i)} + (1 - \gamma_{(i)})/\prod_{k \in \mathcal{B} \setminus \{i\}} q_k(1)$; otherwise, $l_{(i)}^* > z_{(i)}$. For $n = 2$, expression (4.40) states that if the total demand is relatively low, $l_{(i)}^* = s_{(i)}$; otherwise, $l_{(i)}^* = r_{(i)}$. This behavior echoes the observed behavior of the optimal policy in Section 4.5, that when the ending state is *bad* (i.e., with low expected demand), the optimal selection focuses on

maximizing future demand; otherwise, it focuses on maximizing the current revenue. Moreover, as was mentioned earlier, $l_{(i)}^*$ depends on $y(\alpha)$, so it is weakly-coupled. This implies that the Lagrangian index policy has some of the properties of the optimal policy, which the Whittle index policy does not have. We, therefore, expect the former policy to outperform the latter.

4.7.3 Active-constraint index policy

The Whittle and Lagrangian indices are derived from (4.35), where the buyer-specific approximation of the value function $W_i(\alpha_i)$ is $\hat{V}_i^\lambda(\alpha_i)$ in (4.29). For the Whittle index, the penalty price w_i for violating capacity constraint (4.5) is discriminatory (buyer-specific) and hence ignores capacity. For the Lagrangian index, the penalty price λ^* is uniform (common for all buyers). In both cases, this price is applied whenever the firm selects buyer i ($u_i = 1$), even if constraint (4.5) is not active, i.e., even if there is enough capacity to serve buyer i without depriving another buyer of service. This makes capacity more expensive than it really is and introduces a bias in the approximation. To remedy this, we consider an alternative index policy, which we refer to as the active-constraint index policy, where a discriminatory penalty price is applied only when the capacity constraint is active, i.e., when $D_{-i} \geq y(\alpha)$. The active-constraint index, denoted by θ_i , is defined as:

$$\theta_i = r_i + \tilde{V}_i^{\theta_i}(1) - \tilde{V}_i^{\theta_i}(0), \quad i \in \mathcal{B}, \quad (4.41)$$

where the buyer-specific value function $\tilde{V}_i^{\theta_i}(\alpha_i)$ solves the following the DP:

$$\tilde{\Pi}_i^{\theta_i} + \tilde{V}_i^{\theta_i}(\alpha_i) = \mathbb{E} \left[\max_{\mathbf{d}} \left[\max_{u_i \in \mathcal{U}_i(d_i)} \left\{ (r_i - \theta_i 1_{\{D_{-i} \geq y(\alpha)\}}) u_i + \tilde{V}_i^{\theta_i}(u_i + (1 - d_i)\alpha_i) \right\} \right] \right], \quad (4.42)$$

for $i \in \mathcal{B}$. DP (4.42) is the outcome of a stronger relaxation compared to DP (4.29), at the expense of requiring more computations, as its solution depends on the distribution of D_{-i} , $F_{-i}(y)$.

Proposition 4.8. *The solution of DP (4.42) is:*

$$\tilde{V}_i^{\theta_i}(1) = r_i \frac{\gamma_i F_{-i}(y(\boldsymbol{\alpha}) - 1)}{1 - \gamma_i F_{-i}(y(\boldsymbol{\alpha}) - 1)}, \quad \tilde{V}_i^{\theta_i}(0) = 0, \quad i \in \mathcal{B}, \quad (4.43)$$

$$\tilde{\Pi}_i^{\theta_i} = r_i q_i(1) \frac{(1 - \gamma_i) F_{-i}(y(\boldsymbol{\alpha}) - 1)}{1 - \gamma_i F_{-i}(y(\boldsymbol{\alpha}) - 1)}, \quad i \in \mathcal{B}. \quad (4.44)$$

The resulting active-constraint index defined in (4.41) is:

$$\theta_i(\boldsymbol{\alpha}) = \frac{r_i}{1 - \gamma_i F_{-i}(y(\boldsymbol{\alpha}) - 1)}, \quad i \in \mathcal{B}. \quad (4.45)$$

The proof is in Appendix A. Note that $\theta_i(\boldsymbol{\alpha})$ is strongly coupled, since it depends on both $y(\boldsymbol{\alpha})$ and the vector of satisfaction states of all buyers except i (recall that D_{-i} is a function of α_k , $k \in \mathcal{B} \setminus \{i\}$). From (4.45), $\theta_i(\boldsymbol{\alpha})$ has a striking resemblance to z_i in (4.17). The buyer-specific terms r_i and γ_i are identical to those in z_i , and the term $F_{-i}(y(\boldsymbol{\alpha}) - 1)$ is similar to the term $F_{-i}^1(n - 2)$. It represents the *type-I service level of all other buyers*, if buyer i is served, i.e., it is the probability that the total demand of the other $n - 1$ buyers is at most $y(\boldsymbol{\alpha}) - 1$. Note that if $q_i(\alpha_i) > q_j(\alpha_j)$, then $F_{-i}(y(\boldsymbol{\alpha}) - 1) > F_{-j}(y(\boldsymbol{\alpha}) - 1)$, for $i \neq j$. This means that by favoring a buyer with a higher $q_i(\alpha)$, the firm reduces the probability of stockout for the other buyers, thus increasing the chance of ending up in a greater satisfaction state vector and leading to more well-balanced satisfaction and service levels among the buyers. This term gives the active-constraint index policy a significant advantage over policies that use uncoupled indices, such as the revenue-greedy policy $\mathbf{u}^r$ with index r_i and the augmented-revenue-greedy policy $\mathbf{u}^s$ with index $s_i = r_i/(1 - \gamma_i)$. These policies rely on a strict prioritization of the buyers which is independent of their satisfaction states and can lead to very unbalanced satisfaction and service levels that are biased towards the high-priority buyers.

Note that the term in the numerator of (4.45) refers to the current revenue of the firm, while the term in the denominator refers to the drop in future demand. The lower the vector of satisfaction states α_k , $k \in \mathcal{B} \setminus \{i\}$, the higher $F_{-i}(y(\boldsymbol{\alpha}) - 1)$, and therefore the higher the impact of the loss in future demand compared to that of the current

revenue in the computation of $\theta_i(\alpha)$. As in the case of the Lagrangian index policy, this behavior reflects the observed behavior of the optimal policy in Section 4.5 that when the ending state has low expected demand, the optimal selection focuses on the loss in future demand; otherwise, it focuses on the current revenue. An important difference, however, is that in the active-constraint index policy, the emphasis on the current revenue or future demand changes dynamically based on the satisfaction state vector, whereas, in the Lagrangian index policy, it is static. As was noted in the previous paragraph, this constitutes a significant advantage of the active-constraint index policy.

Finally, the following property further reinforces our intuition that the active-constraint index policy is expected to outperform the Whittle and Lagrangian index policies.

Corollary 4.4. *If $y(\alpha) = n - 1$, $\forall \alpha$, the active-constraint index policy is identical to the optimal buyer selection policy $\mathbf{u}^z$ given by Proposition 4.4.*

The proof is in Appendix A. Corollary 4.4 implies that if $y(\alpha) = n - 1$, $\forall \alpha$, the active-constraint index policy is optimal. This is a very attractive property that none of the other two Lagrangian-relaxation-based index policies have.

4.8 Numerical results

We complement our analytical results with a computational study in which we numerically solve equation (4.12) for the three index policies considered in Section 4.7, for a large number of problem instances with five and ten buyers. Our aim is to explore and compare the performance of these policies, reinforce some of our earlier insights, and make new observations. For the five-buyer instances, we also evaluate the index and the FOQ policies against the optimal policy.

4.8.1 Evaluation of index policies

To investigate the performance of the optimal policy and compare it to that of the three Lagrangian relaxation-based index policies, we numerically solve 250 instances

of a problem with five buyers ($n = 5$, $\mathcal{B} = \{1, 2, 3, 4, 5\}$) under all policies. In each instance, we set $c = 1$ and randomly generate the visit rates within the following ranges: $q_i(0) \in (0.005, 0.77)$ and $q_i(1) \in (q_i(0), 0.96)$, $i \in \mathcal{B}$. We also generate five revenue rate values in the interval $(1.15, 1.25)$, sort them in decreasing order, and assign them to the buyers so that $r_i > r_{i+1}$, $i \in \{1, 2, 3, 4\}$, to facilitate the presentation of the buyer-specific results.

For each instance, we numerically solve equation (4.10) for the optimal policy, and equation (4.12) for the three index policies, using value iteration. The solution of (4.10) yields the optimal buyer selection policy $\mathbf{u}^* = \mathbf{u}^{y^*,*} = \mathbf{u}^*(\boldsymbol{\alpha}, y^*, \mathbf{d})$, the optimal ordering policy $y^* = y^{*,\mathbf{u}^*} = y^*(\boldsymbol{\alpha}|\mathbf{u}^*)$, and the corresponding maximum average expected profit $\Pi = \Pi^{*,\mathbf{u}^*}$. The solution of (4.12), for each index policy $\mathbf{u}^x$, $x = r, l^*, \theta$, yields the optimal ordering policy $y^{*,\mathbf{u}^x} = y^*(\boldsymbol{\alpha}|\mathbf{u}^x)$ and the corresponding average expected profit $\Pi^{*,\mathbf{u}^x}$.

For each instance and for each policy $\mathbf{u}$, we also determine the average expected optimal order quantity, denoted by $\bar{y}^{*,\mathbf{u}}$, and the average expected demand and service rates of buyer i , denoted by $\bar{d}_i^{*,\mathbf{u}}$ and $\bar{u}_i^{*,\mathbf{u}}$, $i \in \mathcal{B}$, respectively. To compute these measures, we run, on the side of the main value iteration, additional value iterations of the following DP equations, where in each iteration we use the decisions $y(\boldsymbol{\alpha})$ and $\mathbf{u}(\boldsymbol{\alpha}, y(\boldsymbol{\alpha}), \mathbf{d})$ that result from the maximization step in the main iteration:

$$\bar{y}^{*,\mathbf{u}} + \bar{V}^y(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{d}} [y(\boldsymbol{\alpha}) + \bar{V}^y(\Phi(\boldsymbol{\alpha}, \mathbf{u}(\boldsymbol{\alpha}, y(\boldsymbol{\alpha}), \mathbf{d}), \mathbf{d}))], \quad \forall \boldsymbol{\alpha},$$

$$\bar{d}_i^{*,\mathbf{u}} + \bar{V}^{d_i}(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{d}} [d_i + \bar{V}_i^{d_i}(\Phi(\boldsymbol{\alpha}, \mathbf{u}(\boldsymbol{\alpha}, y(\boldsymbol{\alpha}), \mathbf{d}), \mathbf{d}))], \quad \forall \boldsymbol{\alpha}, \quad i \in \mathcal{B},$$

$$\bar{u}_i^{*,\mathbf{u}} + \bar{V}_i^{u_i}(\boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{d}} [u_i(\boldsymbol{\alpha}, y(\boldsymbol{\alpha}), \mathbf{d}) + \bar{V}_i^{u_i}(\Phi(\boldsymbol{\alpha}, \mathbf{u}(\boldsymbol{\alpha}, y(\boldsymbol{\alpha}), \mathbf{d}), \mathbf{d}))], \quad \forall \boldsymbol{\alpha}, \quad i \in \mathcal{B}.$$

Finally, we calculate the average expected fill rate for each buyer i , denoted by $S_i^{*,\mathbf{u}}$, defined as the average expected probability that buyer i is served given that she demands service, as follows:

$$S_i^{*,\mathbf{u}} = \frac{\bar{u}_i^{*,\mathbf{u}}}{\bar{d}_i^{*,\mathbf{u}}}. \quad (4.46)$$

Figures 4.8 and 4.9 show plots of the average expected profit, order quantity, and fill rate for buyers 1, 3, and 5, under the four considered policies, for the first 100

instances. In each plot, the instances are sorted in ascending order of the values of the optimal policy, for ease of exposition.




Figure 4.8: Average expected profit and order quantity under different policies for 100 instances of a problem with five buyers.




Figure 4.9: Average expected fill rate for buyers 1, 3, and 5 under different policies for 100 instances of a problem with five buyers.

From Figure 4.8, we observe that the average expected profit and order quantity values of the active-constraint index policy are extremely close to the respective values of the optimal policy. For the Lagrangian and Whittle index policies, these performance measures deviate visibly from their optimal policy counterparts.

Figure 4.9 shows that in general, $S_1^{*,u} > S_3^{*,u} > S_5^{*,u}$. This is expected, since in all the instances $r_1 > r_3 > r_5$, whereas the visit rates are generated similarly for all

buyers. In more than half, about half, and less than half of the instances, S_1^{*,u^*} , S_3^{*,u^*} , and S_5^{*,u^*} are equal to one. The average expected fill rate of the active-constraint index policy is equal to or quite close to the respective value of the optimal policy, whereas, in the Lagrangian and Whittle index policies, it deviates substantially from the optimal. This deviation is mostly positive for buyer 1, both positive and negative for buyer 3, and negative for buyer 5. This is expected, because the Whittle index policy is revenue-greedy, assigning the highest priority to buyer 1 and the lowest to buyer 5, and the Lagrangian index policy is often identical to or close to a revenue-greedy policy.

Table 4.3 shows the sample mean and standard deviation, over all 250 instances, of the average expected profit, order quantity, and fill rates of the optimal policy (column 2). It also shows the sample mean and standard deviation of the percent difference of the average expected profit, order quantity, and fill rates of the three index policies from the respective values of the optimal policy (columns 3–5).

Table 4.3: Average performance (sample mean and standard deviation) of all policies, for 250 instances of a five-buyer problem.

Performance measure	Opt $u = u^*$	Active-constraint index $u = u^\theta$ (% diff from Opt)	Lagrangian index $u = u^{l^*}$ (% diff from Opt)	Whittle index $u = u^r$ (% diff from Opt)
$\Pi^{*,u}$	(0.33, 0.11)	(-0.28, 0.51)	(-14.02, 14.75)	(-15.36, 16.66)
$\bar{y}^{*,u}$	(2.81, 0.62)	(-0.34, 1.46)	(-7.18, 19.16)	(-9.66, 19.37)
$S_1^{*,u}$	(0.81, 0.29)	(0.76, 18.00)	(59.22, 117.61)	(60.86, 120.52)
$S_2^{*,u}$	(0.75, 0.31)	(-0.50, 14.65)	(58.20, 110.41)	(58.80, 114.24)
$S_3^{*,u}$	(0.76, 0.32)	(1.09, 18.66)	(23.53, 84.75)	(21.36, 84.28)
$S_4^{*,u}$	(0.68, 0.34)	(1.00, 17.46)	(-11.76, 55.07)	(-15.63, 56.05)
$S_5^{*,u}$	(0.58, 0.35)	(5.35, 22.19)	(-28.81, 52.84)	(-30.59, 53.63)

We observe that the mean average expected profit and order quantity of the optimal policy is 0.33 and 2.81, respectively, while the mean average expected fill rate ranges from 0.58 for buyer 5 to 0.81 for buyer 1. This suggests that the optimal selection policy tries to keep the fill rates relatively balanced. The average expected profit and order quantity of the active-constraint index policy decrease by only 0.28% and 0.34% from the respective values in the optimal policy, indicating that the active-constraint index policy is near-optimal. The average expected fill rate increases for

all buyers, except buyer 2, for whom it slightly decreases. The increase in the fill rates for almost all buyers under a suboptimal policy seems counter-intuitive at first. However, it can be explained by the fact that the suboptimal selection can cause the average demand of the buyers to drop more than their average service rate does, i.e., the denominator of S_i in (4.46) can decrease more than the numerator for all buyers.

The mean percent differences of the performance measures of the Lagrangian and Whittle index policies from the respective values of the optimal policy are significantly larger, with the Whittle index policy having the worst performance. Its mean average expected profit and order quantity are 15.36% and 9.66% lower than the respective values of the optimal policy, and its mean average expected fill rate ranges from 60.86% higher for buyer 1 to 30.59% lower for buyer 5 than the respective values of the optimal policy. The fact that $S_1^{*,u^r} < 100\%$, even though buyer 1 always has top priority, is because, in some satisfaction states, the order quantity is zero, as was discussed in the two-buyer example in Section 4.4.2.

The mean value of $\bar{y}^{*,u}$ of all the index policies is lower than the respective value of the optimal policy (this can also be seen from Figure 4.8). This suggests that the firm, by selecting buyers inefficiently, loses demand and is forced to cut its orders to adapt to the lower demand.

4.8.2 Evaluation of FOQ policy

In the previous section, we computed the optimal ordering policy for the optimal and the three index buyer selection policies considered in Section 4.7. A question that arises is, how well does the best FOQ policy perform compared to the optimal ordering policy? This question is of interest in situations where the firm must allocate a fixed capacity instead of a variable order quantity Adelman and Mersereau (2013); Klein and Kolb (2015); so, designing this capacity is a concern.

To address this question, we devise a procedure in which we fix $y(\alpha) = y_F$, $\forall \alpha$, and numerically solve equation (4.11) for the optimal buyer selection policy, and equation (4.12) without the maximization step for the three index policies, using value iteration. We run this procedure for each $y_F \in \mathcal{B}_0$ and for each of the 250

Table 4.4: Percent difference (sample mean and standard deviation) of the average performance of the optimal FOQ policy (under all selection policies) from the optimal policy, for 250 instances of a five-buyer problem.

Performance measure	Opt $\mathbf{u} = \mathbf{u}^*$	Active-constraint index $\mathbf{u} = \mathbf{u}^\theta$	Lagrangian index $\mathbf{u} = \mathbf{u}^{l^*}$	Whittle index $\mathbf{u} = \mathbf{u}^r$
$100 \times (\Pi^{y_F^*, \mathbf{u}} - \Pi^{*, \mathbf{u}^*}) / \Pi^{*, \mathbf{u}^*}$	(-0.47, 0.90)	(-0.72, 1.14)	(-15.32, 16.30)	(-16.91, 18.54)
$100 \times (y_F^{*, \mathbf{u}} - \bar{y}^{*, \mathbf{u}^*}) / \bar{y}^{*, \mathbf{u}^*}$	(0.42, 7.75)	(0.55, 8.0)	(-7.36, 21.09)	(-10.12, 21.40)

instances considered earlier. For each instance and each buyer selection policy $\mathbf{u}$, we select the optimal value of y_F that yields the highest average expected profit, denoted by $y_F^{*, \mathbf{u}}$.

Table 4.4 shows the sample mean and standard deviation, over all 250 instances, of the percent difference of the average expected profit and optimal FOQ of the optimal and the three index selection policies from the average expected profit and the average order quantity of the optimal policy.

We observe that the mean drop in the average expected profit of the optimal FOQ policy under the optimal selection policy $\mathbf{u}^*$ is only 0.47% and that $y_F^{*, \mathbf{u}^*}$ is only 0.42% higher than $\bar{y}^{*, \mathbf{u}^*}$ on average. The mean percent differences of the average expected profit and order quantity of the FOQ policy under any index policy from the respective values of the optimal policy also change very modestly compared to the corresponding differences in Table 4.3. For example, compare the 15.32% profit loss and 7.36% drop in the average order quantity of the Lagrange index policy using the best FOQ to the 14.02% profit loss and 7.18% drop in the average order quantity of the same index policy using optimal ordering.

These results suggest that the FOQ policy can be fairly efficient if y_F is chosen optimally. If the wrong value of y_F is used, however, the average expected profit can drop significantly. Figure 4.10 shows plots of the percent drop in the average expected profit of the optimal selection policy using an FOQ policy $y(\boldsymbol{\alpha}) = y_F$, $\forall \boldsymbol{\alpha}$ from the respective value of the optimal policy, for different y_F values, for seven representative instances.

The plots show that if the wrong y_F is used, the drop in average expected profit can be extremely high. For the particular set of 250 instances considered, $y_F^{*, \mathbf{u}^*} = 0, 1, 2, 3, 4, 5$ in 0% 2%, 26%, 61%, 11%, and 0% of the instances, respectively. In




Figure 4.10: Percent profit loss of the FOQ policy under the optimal selection, for different FOQ values, for seven representative instances of a problem with five buyers.

almost all instances, the value of y_F that yields the smallest average expected profit is 5.

4.8.3 Effect of number of buyers and revenue rates

To investigate the effect of the number of buyers on the performance of the different policies considered, we numerically solve 150 instances of a problem with ten buyers ($n = 10$, $\mathcal{B} = \{1, \dots, 10\}$). In each instance, we set $c = 1$ and randomly generate the rest of the parameters within the following ranges: $q_i(0) \in (0.005, 0.89)$, $q_i(1) \in (q_i(0), 0.99)$, and $r_i \in (1.15, 1.25)$, $i \in \mathcal{B}$. Numerically finding the optimal policy for this problem is computationally intractable, because the number of computations that must be performed in each value iteration is 65.71 million, as was mentioned in Section ?? . Hence, we limit our study to the three Lagrangian relaxation-derived index policies.

Figure 4.11 shows plots of the average expected profit and order quantity under the three index policies, for the first 100 instances. In each plot, the instances are sorted in ascending order of the values of the active-constraint index policy, for ease of exposition. These plots show that the average expected profit and order quantity under the Lagrangian and Whittle index policies deviate visibly from the respective values under the active-constraint policy, as was the case in the five-buyer problem.

Table 4.5 shows the sample mean and standard deviation, over all 150 instances, of




Figure 4.11: Average expected profit and order quantity under different policies for 100 instances of a problem with ten buyers.

the average expected profit and order quantity of the active-constraint index policy in column 2. Columns 3–4 show the sample mean and standard deviation of the percent difference of the average expected profit and order quantity of the Lagrangian and Whittle index policies from the respective values of the active-constraint policy.

Table 4.5: Average performance (sample mean and standard deviation) of all policies, for 150 instances of a ten-buyer problem.

Performance measure	Active-constraint index $u = u^\theta$	Lagrangian index $u = u^{l^*}$ (% diff from Active-constraint index)	Whittle index $u = u^r$ (% diff from Active-constraint index)
$\Pi^{*,u}$	(1.30, 0.15)	(-7.23, 5.46)	(-8.21, 6.25)
$\bar{y}^{*,u}$	(7.45, 0.58)	(-4.56, 7.51)	(-5.96, 7.04)

We observe that the mean average expected profit and order quantity of the active-constraint index policy is 1.30 and 7.45, respectively. The mean decrease in performance of the Lagrangian and Whittle index policies compared to the active-constraint policy is sizable but not as high as the respective drop in the five-buyer problem. As in the five-buyer problem, the Whittle index policy has the worst performance, with a mean average expected profit and order quantity which are 8.21% and 5.96% lower, respectively, than the respective values of the active-constraint policy.

Finally, to investigate the effect of the revenue rates on the performance of the

three index policies, we numerically solve 150 instances of a problem with ten buyers, where in each instance, we set $c = 1$, and we randomly generate the rest of the parameters in the following ranges: $r_i \in (1.5, 1.7)$, $\forall i \in \mathcal{B}$, $q_i(0) \in (0.1, 0.7)$ and $q_i(1) \in (q_i(0), 1)$, $i \in \mathcal{B}$. That is, the visit rates are slightly less differentiated between buyers than in the previous set of instances, while the revenue rates are more differentiated and higher than in the previous set of instances. The idea is to increase the weight of the current revenue relative to that of the loss in future demand.

For this set of instances, the mean average expected profit of the active-constraint index policy is 3.37, i.e., much higher than the respective value in the previous set of instances (1.30). Moreover, the mean percent difference of the average expected profit of the Lagrangian and Whittle index policies from the respective value of the active-constraint index policy is 2.24% and 5.07%, respectively, i.e., much smaller than the respective differences in the previous set of instances (7.23% and 8.21%). This is expected because, in the new set of instances, the revenue rates are much higher and more differentiated than those in the previous set. The higher and the more differentiated the revenue rates, the better the performance of the myopic revenue-greedy policy (i.e., the Whittle index policy) and the Lagrangian index policy, which, as was mentioned earlier, is often identical to or close to the revenue-greedy policy.

4.9 Discussion and future research

The reactions of heterogeneous buyers to stockouts give rise to a complicated set of trade-offs in inventory and buyer portfolio management. Firms often overlook these trade-offs and deal with the adverse effect of stockouts on buyer goodwill with a penalty cost or a service level constraint, ignoring the effect of goodwill changes on future demand. Moreover, they typically prioritize buyers based on their past sales or margins, ignoring the long-term importance of each buyer. Ordering and buyer selection decisions necessitate more sophisticated practical approaches that carefully balance the acquisition cost, the current revenue from the satisfied buyers, and the loss in future demand from the dissatisfied buyers.

Our model can be extended in several directions. One direction is to consider

markets where unused items are not perishable but are carried over to the next period at a cost, and/or unserved buyers incur a penalty cost in addition to the loss in future demand.

Another direction is to consider ways in which the firm can mitigate the stockout risk, for example by including an expensive backup supplier to cover some of the excess demand, or by incorporating personalized dynamic pricing to expedite the return of dissatisfied buyers.

Chapter 5

Thesis Summary

Designing inventory control policies that account for the adverse effect of stockouts on buyer (or customer) goodwill and future demand has long been a challenging issue for OR/OM researchers and practitioners. To address this issue, we develop and analyze three multi-period models of a supplier(s) selling items to the buyer(s) whose demand is driven by past service.

In **Chapter 2**, we develop a multiperiod model of a supplier selling items to a buyer who rates the supplier based on the history of her service, measured in terms of in-stock/out-of-stock incidents. We show that while the myopic policy is a basestock policy with rating-dependent basestock levels, the optimal policy for the infinite-horizon problem partitions the inventory space in several order-up-to and do-not-order intervals, for each rating. The optimal decision—order up to the next point or do not order—depends on whether ordering reduces the risk of downrating the supplier—lowering her expected future profits—enough to offset the resulting increase in ordering and inventory holding costs. This tradeoff depends on the inventory level and the buyer demand density function. We derive and evaluate bounds on the optimal policy and found properties of this policy. We show that a basestock policy is optimal and analytically tractable for cases where the buyer has random demand but short memory of service (two ratings), and constant demand but long memory (more than two ratings). If the buyer has random demand and long memory of service, a basestock policy is optimal under a certain condition on the demand and

other parameters. We use the service-driven model to impute the fixed stockout cost reflecting the loss of goodwill due to a stockout in a newsvendor model that posits such a cost. Our results show that using the imputed stockout cost is quite efficient, thus providing a justification for the use of the newsvendor model. However, using an arbitrary, fixed stockout cost can significantly reduce the supplier's profits.

In **Chapter 3**, we develop a model of a repeat buyer (she) sharing her patronage among two heterogeneous newsvendor-type suppliers over an infinite horizon. To enjoy the best service advantage, the buyer plays one supplier (him) against the other by rewarding product availability with repurchase (loyalty) and punishing stockouts with switching (disloyalty) in the next period.

Our analytical and numerical results provide new insight into these decisions. They suggest that the main concern of the suppliers under competition is to maintain the buyer's loyalty because losing it as a result of a stockout means foregoing profits for many periods following the stockout. This concern forces each supplier to significantly increase his active basestock level above his myopic level, reducing the frequency of stockouts and the role of the backorder cost.

The benefits of supplier competition for the buyer are completely wiped out if the suppliers decide to cooperate. In this case, the supplier with the lower myopic profit lowers his active basestock level below his myopic level—possibly down to zero—ceding his demand share to the more profitable supplier who sets his active basestock level above his myopic level but still below his basestock level at equilibrium under competition. The buyer can recover the high fill rate that she can enjoy under competition if she charges the cooperating suppliers an adjustment penalty backorder rate when products are unavailable on demand. This rate can be excessively high to be of practical use if the suppliers' margin-to-interest rate is high.

Finally, most of the results for two suppliers extend to multiple suppliers if the buyer uses a round-robin policy where she switches from one supplier to the next on a circular basis after each stockout.

In **Chapter 4**, we develop a newsvendor model of a firm with a number of heterogeneous buyers that captures the effect that the joint ordering and buyer selection decisions have on the visit dynamics of the buyers and the long-term average profit

of the firm.

We show that for two buyers, the optimal selection policy is index-based where the index of each buyer is increasing in her revenue rate, the relative loss in her future demand if she is not served, and the type-I service level of the other buyer if she is served. This implies favoring buyers who are profitable, reactive to quality-of-service changes, and predictable in their visit behavior when satisfied.

We demonstrate that for more buyers, the optimal selection policy is not index-based but may depend on the demand realization. Our results suggest that if the ex-post satisfaction state is high, the firm should select buyers to maximize current revenues; otherwise, it should select buyers to maximize future demand.

Our analysis suggests that under optimal selection, the optimal order quantity is non-decreasing in the satisfaction state of the buyers. For two buyers, this implies that effectively an FOQ policy is optimal in steady state. We demonstrate that if buyers are not selected optimally, it may the firm may be better off ordering fewer items in a higher satisfaction state than in a lower state, to drive buyer satisfaction to more profitable states which the suboptimal selection policy fails to do.

For the Lagrangian index policy that we develop based on the relaxed problem, we manage to derive the “best” Lagrangian price in closed form as the solution to the Lagrangian dual problem, allowing us to obtain the tightest bound of the original problem.

The active-constraint index that we develop resonates very well with the observed optimal policy, leading to relatively well-balanced satisfaction states and service levels among the buyers. It augments the revenue rate of each buyer by a factor that neatly separates into two terms: the drop in the buyers’ visit rate if she is not served, and the type-I service level of the other buyers if she is served. The higher the last term, the smaller the expected ex-post satisfaction state of the buyers, and the higher the relative weight of the future demand over the current revenue. Our numerical results show that the active-constraint index policy is near-optimal.

Appendix A

Chapter 2 Supplemental Material

Intuition behind the optimal policy

To gain intuition behind the optimal policy given by (2.28), consider Figure A.1 which shows graphs of two different demand density functions $f(w)$ and four different initial inventory levels x_0 , $\bar{x}_0$, x_1 , and $\bar{x}_1$ in different areas of the demand space. The difference between the two graphs is that in (a), $f(w)$ increases sharply and decreases smoothly, whereas in (b), it increases and decreases smoothly. For each inventory level, the shaded area of width ϵ represents the reduction in the probability of a stockout—hence in the risk of downgrading the supplier—if the supplier orders a small quantity ϵ . For $x_0 \in R_\alpha^0$ in graph (a) and $x_1 \in R_\alpha^1$ in graph (b), this reduction is significant, because by ordering ϵ , the supplier would eliminate a large area under the part of $f(w)$ that corresponds to backorders, defined as $w > x_0$ and $w > x_1$, respectively. Thus, in these cases, it would be worth it for her to order at least ϵ , despite the resulting higher ordering and inventory holding costs. For $\bar{x}_0 \in \bar{R}_\alpha^0$ and $\bar{x}_1 \in \bar{R}_\alpha^1$ in graph (b), the reduction is minor. In these cases, it would not be worth ordering more items. These examples suggest that the number of order-up-to-points beyond S_α^0 , n , depends on the shape of $f(w)$, and should be bounded by the number of its local maxima.




Figure A.1: Reduction in the probability that the supplier stocks out, if she orders a small quantity ϵ , for different demand density functions and initial inventory levels.

Proof of Proposition 2.1. Let S_α^{my} be the global maximizer of $\Lambda_\alpha(y)$. From expressions (2.11)-(2.12), we get $S_\alpha^{my} = \arg \min_y \{L_\alpha(y)\}$, where

$$L_\alpha(y) = K_1 y + K_2 \left\{ [q_\alpha \int_y^\infty (w - y) dF(w)] 1_{\{y \geq 0\}} + [q_\alpha (\theta - y) - \bar{q}_\alpha y] 1_{\{y < 0\}} \right\}.$$

The first two derivatives of $L_\alpha(y)$ are $L'_\alpha(y) = K_1 - K_2(q_\alpha \bar{F}(y) 1_{\{y \geq 0\}} + 1_{\{y < 0\}})$ and $L''_\alpha(y) = K_2 q_\alpha f(y) 1_{\{y \geq 0\}}$. Clearly, $L''_\alpha(y) \geq 0$, implying that $L_\alpha(y)$ is convex, hence S_α^{my} is its unique minimizer. From (2.11), this further means that $\Lambda''_\alpha(y) \leq 0$, hence $\Lambda_\alpha(y)$ is concave and S_α^{my} is its unique maximizer. Therefore, the optimal inventory control policy is a basestock policy given by (2.16). If $y < 0$, then $L'_\alpha(y) < 0$, since $K_1 - K_2 < 0$ from (2.9). This implies that $S_\alpha^{my} \geq 0$. Moreover, from the first-order condition, $S_\alpha^{my} = \arg \min_y \{q_\alpha \bar{F}(y) \leq K_1/K_2\}$, which can be rewritten as (2.17). $\square$

Proof of Proposition 2.2. Proof of (2.18). The proof of (2.18) follows immediately from (2.10) for $\Pi_\alpha(x)$, and from the definition $V_\alpha(x)$, once we note that $y \geq x$ implies that as $x \rightarrow \infty$, $y \rightarrow \infty$; therefore, $\lim_{y \rightarrow \infty} -K_1 y = -\infty$.

Proof of (2.20)-(2.21). Consider a nominal and a perturbed sample path that start from a nominal and a perturbed initial state, (x_0, α_0) and (x'_0, α'_0) , respectively, where $x'_0 = x_0$ and $\alpha'_0 > \alpha_0$. From (2.5), $q_{\alpha'_0} \geq q_{\alpha_0}$. Suppose that the nominal path follows the optimal inventory control policy, while the perturbed path follows the nominal path by setting $y'_t = y_t$, where $y_t \geq x_t$. Clearly, the policy followed by the

perturbed path is suboptimal. We will show that in each period, the perturbed path: (i) can always set $y'_t = y_t$, and (ii) is always at least as profitable as the nominal sample path, implying that $\Pi_{\alpha'_0}(x'_0) = \Pi_{\alpha'_0}(x_0) \geq \Pi_{\alpha_0}(x_0)$.

To create the two paths, in each period t , generate a common random number $z_t \in [0, 1]$ and use it to generate the demand seen by the supplier in the two paths, denoted by d_t and d'_t , respectively. There are three cases to consider. **Case i:** If $0 \leq z_t \leq q_{\alpha_t}$, then $d_t = d'_t = w_t$; **Case ii:** if $q_{\alpha_t} < z_t \leq q_{\alpha'_t}$, then $d_t = 0$ and $d'_t = w_t$; **Case iii:** if $q_{\alpha'_t} < z_t \leq 1$, then $d_t = d'_t = 0$, where w_t is the buyer demand in period t .

Start by setting $y'_0 = y_0$. This is feasible since $x'_0 = x_0$. Then, generate w_0 from $f(\cdot)$ and z_0 and use it to generate d_0 and d'_0 , respectively, based on the rule described above.

Case i: If $z_0 \leq q_{\alpha_0}$, then $d'_0 = d_0 = w_0$, $x_1 = y_0 - d_0 = y_0 - w_0$, and $x'_1 = y'_0 - d'_0 = y_0 - w_0 = x_1$. Moreover, $y_1 \geq x_1$. Setting $y'_1 = y_1$ is feasible since $x'_1 = x_1$. There are two subcases. **Subcase i-a:** If $w_0 > y'_0 = y_0$, then $\alpha_1 = \alpha_0 - \delta_{\alpha_0}^-$ and $\alpha'_1 = \alpha'_0 - \delta_{\alpha'_0}^-$. From (2.4) and the assumption $\alpha'_0 > \alpha_0$, it follows that $\alpha'_1 \geq \alpha_1$. If $\alpha_0 = 1$ and $\alpha'_0 = 2$, then from (2.4), $\alpha'_1 = \alpha_1 = 1$. **Subcase i-b:** If $w_0 \leq y'_0 = y_0$, then $\alpha_1 = \alpha_0 + \delta_{\alpha_0}^+$ and $\alpha'_1 = \alpha'_0 + \delta_{\alpha'_0}^+$. From (2.4) and the assumption $\alpha'_0 > \alpha_0$, it follows that $\alpha'_1 \geq \alpha_1$. If $\alpha_0 = M - 1$ and $\alpha'_0 = M$, then $\alpha'_1 = \alpha_1 = M$.

Case iii: If $z_0 > q_{\alpha'_0}$, then $d'_0 = d_0 = 0$, $x_1 = y_0 - d_0 = y_0$, and $x'_1 = y'_0 - d'_0 = y_0 = x_1$. Since $d_0 = 0$, $y_1 = y_0$. Again, setting $y'_1 = y_1$ is feasible since $x'_1 = x_1$. Also, $\alpha_1 = \alpha_0$ and $\alpha'_1 = \alpha'_0$, where $\alpha'_1 > \alpha_1$, since $\alpha'_0 > \alpha_0$.

In both cases i and iii, the perturbed and nominal paths are identical. Thus, both paths give rise to the same profits for the supplier. Also, the ratings in both paths either maintain their order or become identical, in which case, they coincide from that point on.

Case ii: If $q_{\alpha_0} < z_0 \leq q_{\alpha'_0}$, then $d_0 = 0$, $d'_0 = w_0$, $x_1 = y_0 - d_0 = y_0 \geq x_0$, and $x'_1 = y'_0 - d'_0 = y_0 - w_0 < x_1$. Moreover, $\alpha_1 = \alpha_0$, $y_1 = x_1 = y_0$, and the supplier's profit in the nominal path is $-hy_0$ (loss). In the perturbed path, there are two subcases. **Subcase ii-a:** If $w_0 > y'_0 = y_0$, the supplier receives a revenue of ry_0 and ends up with inventory $x'_1 = y_0 - w_0 < 0$, incurring a backorder cost of $b(w_0 - y_0)$. Considering also the revenue $r(x'_1)^- = r(w_0 - y_0)$ for the backordered

demand that she will receive in the next period at a discount of β , her total current revenue is $ry_0 + \beta r(w_0 - y_0)$. She can set $y'_1 = y_1 = y_0$ if she orders $w_0 > 0$ at a cost of βcw_0 , when rolled back into the current period. Therefore, her profit is $ry_0 + \beta r(w_0 - y_0) - b(w_0 - y_0) - \beta cw_0 = [(1 - \beta)r + b]y_0 + (\beta p - b)w_0 > -hy_0$ (profit in the nominal path), since by assumption $\beta p - b > 0$. Moreover, $\alpha'_1 = \alpha'_0 - \delta_{\alpha'_0}^-$. From (2.4) and the assumption $\alpha'_0 > \alpha_0 = \alpha_1$, it follows that $\alpha'_1 \geq \alpha_1$. Note that if $\alpha'_0 = \alpha_0 + 1$, then from (2.4), $\alpha'_1 = \alpha_1 = \alpha_0$. **Subcase ii-b:** If $w_0 \leq y'_0$, the supplier receives a revenue of rw_0 , ends up with inventory $x'_1 = y'_0 - d'_0 = y_0 - w_0 \geq 0$, incurring an inventory holding cost of $h(y_0 - w_0)$. She can set $y'_1 = y_1 = y_0$ if she orders $w_0 > 0$ at a cost of βcw_0 , when rolled back into the current period. Thus, her profit is $rw_0 - \beta cw_0 - h(y_0 - w_0) = (r - \beta c + h)w_0 - hy_0 > -hy_0$ (profit in the nominal path). Moreover, $\alpha'_1 = \alpha'_0 + \delta_{\alpha'_0}^+$. From (2.4) and the assumption $\alpha'_0 > \alpha_0$, it follows that $\alpha'_1 \geq \alpha_1$. From the above analysis, in both subcases, $y'_1 = y_1$, $\alpha'_1 \geq \alpha_1$, and the perturbed sample path has a higher profit than the nominal path.

Repeating the same argument for the next period and all the periods thereafter, we can see that in each period t , $y'_t = y_t$ and $\alpha'_t \geq \alpha_t$, and the profit of the perturbed path is greater than or equal to that of the nominal path. Almost surely, at some point, the rating of the perturbed path will coincide with the rating of the nominal path and the two paths will be identical from then on. Therefore, the total profit of the perturbed sample path is greater than or equal to the total profit of the nominal path, which implies that $\Pi_{\alpha'_0}(x'_0) = \Pi_{\alpha'_0}(x_0) \geq \Pi_{\alpha_0}(x_0)$. Clearly, from (2.10), $V_{\alpha'_0}(x'_0) = V_{\alpha'_0}(x_0) \geq V_{\alpha_0}(x_0)$, too.

Proof of (2.22)-(2.23). A lower bound for $\Pi_\alpha(x)$ can be constructed by considering the myopic policy given by Proposition 2.1, or any other feasible policy. Here, we consider an order-up-to policy with rating-dependent order-up-to points S_α , where, similarly to S_α^{my} , S_α is non-decreasing in α . This property guarantees that if $x_0 \leq S_{\alpha_0}$, then $y_0 = S_{\alpha_0}$, and more generally, $y_t = S_{\alpha_t}$, $t > 0$. As a result, $\{\alpha_t, t \geq 0\}$ is a discrete-time Markov chain with state-space A and non-zero transition probabilities

$p_{\alpha,\alpha'}$ that are independent of x_t and are given as follows:

$$\begin{aligned} p_{\alpha,\alpha-1} &= q_\alpha \bar{F}(S_\alpha), 1 < \alpha \leq M, \quad p_{\alpha,\alpha+1} = q_\alpha F(S_\alpha), 1 \leq \alpha < M-1, \\ p_{\alpha,\alpha} &= \bar{q}_\alpha + q_\alpha \bar{F}(S_\alpha) 1_{\{\alpha=1\}} + q_\alpha F(S_\alpha) 1_{\{\alpha=M\}}, 1 \leq \alpha \leq M. \end{aligned} \quad (\text{A.1})$$

The resulting discounted expected profit, which is denoted by $\Pi_\alpha^l(x)$, satisfies $\Pi_\alpha^l(x) = c(x)^+ + p(x)^- + V_\alpha^L(x)$, where $V_\alpha^L(x)$ is the value function corresponding to the considered policy and is a lower bound for $V_\alpha(x)$. For $x \leq S_\alpha$, $V_\alpha^L(x) = V_\alpha^L(S_\alpha)$, where $V_\alpha^L(S_\alpha)$ is obtained by solving the following equation:

$$V_\alpha^L(S_\alpha) = \Lambda_\alpha(S_\alpha) + \beta \sum_{\alpha'} p_{\alpha,\alpha'} V_{\alpha'}^L(S_{\alpha'}), \alpha \in A. \quad (\text{A.2})$$

From the monotonicity of q_α and S_α , it follows that $\Lambda_\alpha(S_\alpha)$ and $V_\alpha^L(S_\alpha)$ are also increasing in α . Letting $\mathbf{V}^L$ and $\mathbf{\Lambda}$ denote the vectors $(V_1^L(S_1), \dots, V_M^L(S_M))$, and $(\Lambda_1(S_1), \dots, \Lambda_M(S_M))$, respectively, and letting $\mathbf{P}$ denote transition probability matrix of the discrete-time Markov chain defined above, A.2 can be written in matrix form as $\mathbf{V}^L = \mathbf{\Lambda} + \beta \mathbf{P} \mathbf{V}^L$. The solution of this equation is $\mathbf{V}^L = (\mathbf{I} - \beta \mathbf{P})^{-1} \mathbf{\Lambda}$, where $\mathbf{I}$ is the identity matrix. Note that $\mathbf{P}$ is a tridiagonal matrix and so is $\mathbf{I} - \beta \mathbf{P}$. The inverse of the latter matrix can be obtained using a formula developed by Usmani (1994) for computing the inverse of tridiagonal matrices based on their principal minors. Applying that formula to our problem, yields:

$$\begin{aligned} V_\alpha^L(S_\alpha) &= \sum_{j=1}^{\alpha-1} \left(\beta^{\alpha-j} \Lambda_j(S_j) \frac{\eta_{j-1} \varphi_{\alpha+1}}{\eta_M} \prod_{k=j+1}^i q_k \bar{F}(S_k) \right) + \Lambda_\alpha(S_\alpha) \frac{\eta_{\alpha-1} \varphi_{\alpha+1}}{\eta_M} \\ &\quad + \sum_{j=\alpha+1}^M \left(\beta^{j-\alpha} \Lambda_j(S_j) \frac{\eta_{\alpha-1} \varphi_{j+1}}{\eta_M} \prod_{k=i}^{j-1} q_k F(S_k) \right), \end{aligned} \quad (\text{A.3})$$

where $\eta_i, i = 0, \dots, M$, are the principal minors of $\mathbf{I} - \beta \mathbf{P}$ and satisfy the recurrence equations:

$$\begin{aligned} \eta_0 &= 1, \quad \eta_1 = 1 - \beta(1 - q_1 F(S_1)), \\ \eta_i &= (1 - \beta \bar{q}_i - \beta q_i F(S_i) 1_{\{i=M\}}) \eta_{i-1} - \beta^2 q_i \bar{F}(S_i) q_{i-1} F(S_{i-1}) \eta_{i-2}, \end{aligned} \quad (\text{A.4})$$

for $i = 2, \dots, M$. Similarly, $\varphi_i, i = 1, \dots, M + 1$, are given by:

$$\begin{aligned}\varphi_{M+1} &= 1, \quad \varphi_M = 1 - \beta (1 - q_M \bar{F}(S_M)), \\ \varphi_i &= (1 - \beta \bar{q}_i - \beta q_i \bar{F}(S_i) 1_{\{i=1\}}) \varphi_{i+1} - \beta^2 q_{i+1} \bar{F}(S_{i+1}) q_i F(S_i) \varphi_{i+2},\end{aligned}\tag{A.5}$$

for $i = M - 1, \dots, 1$. If $x_0 > S_{\alpha_0}$, setting $y_t = S_{\alpha_t}$ may not be feasible. In this case, the transition probabilities from one rating to another depend on the current inventory level x , making the computation of $V_\alpha^L(x)$ too demanding and outside the scope of this paper. To bypass this difficulty, we set aside the quantity $x - S_\alpha$ and hold it as unused inventory forever, at a holding cost of $h(x - S_\alpha) / (1 - \beta)$. Thus, for any initial inventory level x , the lower bound of $V_\alpha(x)$ is given by $V_\alpha^L(x) = -h(x - S_\alpha)^+ / (1 - \beta) + V_\alpha^L(S_\alpha)$, where $V_\alpha^L(S_\alpha)$ is given by (A.3).

Note that the lower bound developed in Lemma 2.1 in Robinson (2016) for a similar model is a special case of our bound where $S_\alpha = 0, \alpha \in A$. In this case, $V_\alpha^L(0)$ can be found from (A.3), after replacing $F(0) = 1 - \bar{F}(0) = 0$, $\Lambda_\alpha(0) = (p\beta - b)q_\alpha\theta, \eta_i = \prod_{k=1}^i (1 - \beta + \beta q_k 1_{\{k>1\}})$, and $\varphi_i = \prod_{k=i}^M (1 - \beta + \beta q_k 1_{\{k>1\}}), i \in A$, from (A.4) and (A.5), as follows:

$$V_\alpha^L(0) = (p\beta - b)\theta \sum_{j=1}^i \beta^{\alpha-j} \prod_{k=j}^{\alpha} q_k / (1 - \beta + q_k \beta 1_{\{k>1\}}), \alpha \in A.$$

This bound is positive, implying (2.19), but smaller than or equal to $V_\alpha^L(S_\alpha^{my})$, because when $x \leq S_\alpha^0$, the optimal policy is to order up to S_α^0 , where $S_\alpha^0 \geq S_\alpha^{my}$, as was mentioned earlier.

An upper bound for $\Pi_{\alpha_0}(x_0)$ can be constructed by considering the ideal scenario under which definitions (2.4) are replaced by $\delta_{\alpha_t}^+ = 1_{\{\alpha_t < M\}}$ and $\delta_{\alpha_t}^- = 0$. From (2.3), this implies that α_t remains unchanged with probability $\bar{q}_{\alpha_t}$, as before, and increases by one unit with probability q_{α_t} , irrespectively of whether the supplier satisfies the demand or not. Hence, under this scenario, $\{\alpha_t, t \geq 0\}$ is decoupled from the inventory control policy of the supplier, and the optimal inventory control policy is the myopic policy given by Proposition 2.1. The resulting discounted expected profit, which is denoted by $\Pi_\alpha^u(x)$, satisfies $\Pi_\alpha^u(x) = c(x)^+ + p(x)^- + V_\alpha^U(x)$, where $V_\alpha^U(x)$ is the value

function corresponding to the myopic policy under the ideal scenario and is an upper bound for $V_\alpha(x)$. For $x \leq S_\alpha^{my}$, $V_\alpha^U(x) = V_\alpha^U(S_\alpha^{my})$, where $V_\alpha^U(S_\alpha^{my})$ is obtained by solving a system of equations identical to that in (A.2) with $V_\alpha^U(\cdot)$ instead of $V_\alpha^L(\cdot)$. The solution of that system is given by an expression identical to (A.3) with $V_\alpha^U(S_i^{my})$ instead of $V_\alpha^L(S_i)$. After replacing $F(S_j) = F(\infty) = 1$, $\Lambda(S_j) = \Lambda(S_j^{my})$, $j \in A$, $\eta_i = \prod_{k=1}^i (1 - \beta + \beta q_k 1_{\{k < M\}})$, and $\varphi_i = \prod_{k=i}^M (1 - \beta + \beta q_k 1_{\{k < M\}})$, $i = 1, \dots, M$, from (A.4) and (A.5), respectively, $V_\alpha^U(S_\alpha^{my})$ is given as follows:

$$V_\alpha^U(S_\alpha^{my}) = \sum_{j=\alpha}^M \left(\beta^{j-\alpha} \frac{\Lambda_j(S_j^{my})}{q_j} \prod_{k=\alpha}^j \frac{q_k}{1 - \beta + q_k \beta 1_{\{k < M\}}} \right). \quad (\text{A.6})$$

If $x_0 > S_{\alpha_0}^{my}$, then $y_t \neq S_{\alpha_t}^{my}$, $t \geq 0$, and we face the same difficulty in computing $V_\alpha^U(x)$ as in the case of $V_\alpha^L(x)$, mentioned earlier. To circumvent this difficulty, we simply discard quantity $x - S_\alpha^{my}$ at no cost; therefore, the starting inventory effectively is S_α^{my} and hence $V_\alpha^U(x) = V_\alpha^U(S_\alpha^{my})$, as was the case for $x \leq S_\alpha^{my}$.

Note that Robinson (2016) in Lemma 2.1 developed an upper bound for a similar model under the assumption that constraint $y_t \geq x_t$ is relaxed and the supplier is allowed to order after the demand has been observed. Obviously, in this case, the optimal policy is to set $y_t = d_t$, $\forall t$, which is equivalent to a make-to-order policy with known demand. Moreover, the supplier does not incur any holding and backorder costs and only receives revenue for the items sold. Therefore, her period profit is $\Lambda_\alpha(d) = q_\alpha p \theta$, which is larger than the profit $\Lambda_\alpha(S_\alpha^{my})$ that we have considered. The resulting upper bound for the value function is looser than our bound, and from (A.6), the upper bound of Robinson (2016) can be written as $V_\alpha^U(x_0) = p \theta \sum_{j=\alpha}^M \left(\beta^{j-\alpha} \prod_{k=\alpha}^j q_k / (1 - \beta + q_k \beta 1_{\{k < M\}}) \right)$. $\square$

Proof of Lemma 2.1. The proof is similar to that of Lemma 2 in Robinson (2016). After reinstating t , expression (2.15), for $y_t < 0$, becomes:

$$H_{\alpha_t}(y_t) = -L_{\alpha_t}(y_t) + \beta \left\{ q_{\alpha_t} \left[\int_0^\infty V_{\alpha_t - \delta_{\alpha_t}^-}(y_t - w_t) dF(w_t) \right] + \bar{q}_{\alpha_t} V_{\alpha_t}(y_t) \right\}. \quad (\text{A.7})$$

Consider a restricted policy where $y_t \geq 0$ for $t \geq \tau$, for some period τ . Clearly, the

discounted expected profit over an infinite horizon under the restricted policy is less than or equal to the corresponding profit under the optimal policy, but as $\tau \rightarrow \infty$, the two profits become equal, since the discounted expected profit is bounded. Therefore, it suffices to show that the optimal policy y_t^* satisfies $y_t^* \geq 0$ for a finite τ and then let $\tau \rightarrow \infty$. Under the restricted policy, $y_t^* \geq 0$ for $t \geq \tau$, by definition. To extend the result to periods $t < \tau$, we use backward induction. Assume that $y_{t+s}^* \geq 0, s > 1$. If $y_t < 0$, then $x_{t+1} < 0$, from (2.2); therefore, $x_{t+1} < y_{t+1}^*$. Using (2.14), equation (A.7) becomes,

$$\begin{aligned} H_{\alpha_t}(y_t) = & -K_1 y_t - K_2 [q_{\alpha_t} \theta - y_t] + \beta \{ q_{\alpha_t} [\int_0^\infty [K_3 q_{\alpha_t - \delta_{\alpha_t}^-} \theta \\ & + H_{\alpha_t - \delta_{\alpha_t}^-}(y_{t+1}^*)] dF(w_t)] + \bar{q}_{\alpha_t} [K_3 q_{\alpha_t - \delta_{\alpha_t}^-} \theta + H_{\alpha_t - \delta_{\alpha_t}^-}(y_{t+1}^*)] \} \}. \end{aligned}$$

The derivative of the above expression is $H'_{\alpha}(y_t) = K_2 - K_1$. The r.h.s. of the above expression is positive and equal to $(1 - \beta)p + b$ from (2.7), yielding the desired result. $\square$

Proof of Proposition 2.3. First, we show that $S_\alpha^n \leq \bar{S}_\alpha^n$. From (2.31), $S_a^n = \arg \min_{y \geq (x)^+} \{H'_\alpha(y) \leq 0\}$. From (2.33) and (2.29), a sufficient condition for $H'_\alpha(y) \leq 0$ is $H'_\alpha(y) = -L'_\alpha(y) + \beta q_\alpha f(y) [V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)] \leq 0$, where $L'_\alpha(y) = K_1 - K_2 q_\alpha \bar{F}(y)$, from (2.27). Substituting $L'_\alpha(y)$, replacing $V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)$ with its upper bound Δ_α from (2.25), rearranging terms, and using the above expression for S_a^n , yields the upper bound of S_a^n given by expression (2.35).

Next, we show that $S_\alpha^0 \geq S_\alpha^{my}$. For $\alpha = M$, there are two cases to consider. In case 1, $S_M^0 \geq 0$ and $H'_M(S_M^0) = 0$, i.e., S_M^0 satisfies the first-order condition. In case 2, $S_M^0 = 0$ and $H'_M(0) < 0$. From (2.33) and (2.31), we have,

$$H'_M(S_M^0) = -L'_M(S_M^0) + \beta q_M f(S_M^0) [V_M(0) - V_{M-1}(0)] = 0$$

and,

$$(1 - \beta \bar{q}_M) H'_M(0) = -L'_M(0) + \beta q_M f(0) [V_M(0) - V_{M-1}(0)] < 0,$$

for the two cases, respectively. From (2.20), $\beta q_M f(\cdot) [V_M(0) - V_{M-1}(0)] \geq 0$. Therefore, the above two expressions imply that $L'_M(S_M^0) \geq 0$ and $L'_M(0) \geq 0$, respectively. Given that $L''_\alpha(y) = K_2 q_\alpha f(y) > 0$ from (2.27), this further implies that $S_M^{my} \leq S_M^0$. For $\alpha = M - 1$, there are three cases to consider. In case 1, $0 \leq S_{M-1}^0 \leq S_M^0$ and $H'_{M-1}(S_{M-1}^0) = 0$, i.e., S_{M-1}^0 satisfies the first-order condition. In case 2, $S_{M-1}^0 = 0$ and $H'_{M-1}(0) < 0$. For these cases, the proof that $S_{M-1}^{my} \leq S_{M-1}^0$ is the same as in the corresponding cases for $\alpha = M$. In case 3, $S_{M-1}^0 > S_M^0$. From (2.17) and (2.5), $S_{M-1}^{my} \leq S_M^{my}$. Given that $S_M^0 \geq S_M^{my}$, as we showed earlier for $\alpha = M$, it immediately follows that $S_{M-1}^0 \geq S_{M-1}^{my}$. The above arguments for $\alpha = M - 1$ hold similarly for all the remaining ratings $\alpha \in \{1, M - 2\}$. $\square$

Proof of Proposition 2.4. We use induction. For the one-period problem, $V_\alpha(x) = \max_{y \geq x} \{\Lambda_\alpha(y)\}$, by (2.14). As was shown in Proposition 2.1, for this problem, $\Lambda''_\alpha(y) \leq 0$, which implies that $V''_\alpha(x) \leq 0$. Assuming that the proposition has been proved for $t - 1$ periods, we will show that it holds for t periods. This is equivalent to showing that if $V''_\alpha(x) < 0, x \geq 0$, then $H''_\alpha(y) < 0, y \geq 0$, too. From (2.34), it suffices to show that $-L''_\alpha(y) + \beta q_\alpha f'(y) [V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)] \leq 0$, because all the other terms in (2.34) are negative. More specifically, the terms containing $V''_\alpha(\cdot)$ are negative because of the induction hypothesis, and the term $V'_{\alpha+\delta_\alpha^+}(0^+)$ is negative from (2.29). Substituting $L''_\alpha(y) = K_2 q_\alpha f(y) > 0$ from (2.27) into the above inequality and rearranging terms yields $f'(y)/f(y) \leq K_2/\beta [V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0)]$. If (2.37) holds, the last inequality holds immediately, because $V_{\alpha+\delta_\alpha^+}(0) - V_{\alpha-\delta_\alpha^-}(0) \leq \Delta_\alpha$ from (2.24). $\square$

Proof of Proposition 2.6. To prove (i), write (2.32), for $\alpha = 1, 2$, after replacing $L_\alpha(y)$ from (2.27) and rearranging terms the function $G_\alpha(y)$ is equal to,

$$\frac{H_\alpha(y) - \beta \bar{q}_\alpha V_\alpha(y)}{q_\alpha} = -\frac{yK_1}{q_\alpha} - K_2 B(y) + \beta \left[\int_0^y V_{\alpha+\delta_\alpha^+}(y-w) dF(w) + V_1(0) \bar{F}(y) \right].$$

From (2.30), the function $G_\alpha(y)$ is either equal to $(H_\alpha(y) - C_\alpha^1)/q_\alpha$ or $(H_\alpha(y)(1 - \beta \bar{q}_\alpha) - C_\alpha^2)/q_\alpha$, where C_α^1, C_α^2 are constants, for $\alpha = 1, 2$, respectively. $\partial^n G_\alpha(y)/\partial y^n$ is either equal to $\partial^n H_\alpha(y)/\partial y^n$ or $(1 - \beta \bar{q}_\alpha) \partial^n H_\alpha(y)/\partial y^n$, where $1 - \beta \bar{q}_\alpha > 0$. This

implies that $G_\alpha(y)$ has the same shape as $H_\alpha(y)$. Therefore, both functions attain all their extrema at the same values of y . This includes the global maximizer, S_α^0 , i.e., $S_\alpha^0 = \arg \max_y \{H_\alpha(y)\} = \arg \max_y \{G_\alpha(y)\}$. Therefore, to show that $S_2^0 \geq S_1^0$, it suffices to show that $\arg \max_y \{G_2(y)\} \geq \arg \max_y \{G_1(y)\}$. Define $\Delta V_2(y) = V_{2+\delta_2^+}(y) - V_{1+\delta_1^+}(y) = V_{2+\delta_2^+}(y) - V_2(y)$. Clearly, from (2.20), $\Delta V_2(y) \geq 0$ (note that if $M = 2$, $\Delta V_2(y) = 0$). $G_2(y) - G_1(y)$ can now be written as follows

$$G_2(y) - G_1(y) = [(q_2 - q_1) K_1 / q_1 q_2] y + \beta \int_0^y \Delta V_2(y - w) dF(w).$$

Clearly, this difference is a positive non-decreasing function in y , implying that $\arg \max_y \{G_2(y)\} \geq \arg \max_y \{G_1(y)\}$.

To prove (ii), write (2.14) for $\alpha = 1, 2$ and $M = 2$, after replacing $H_\alpha(y)$ from (2.32), as follows:

$$V_\alpha(0) = K_3 q_\alpha \theta - L_\alpha(S_\alpha^0) + \beta q_\alpha [V_2(0) F(S_\alpha^0) + V_1(0) \bar{F}(S_\alpha^0)] + \beta \bar{q}_\alpha V_\alpha(0).$$

After some manipulations, the above expression can be specialized for $\alpha = 1, 2$ as follows:

$$V_1(0) = K_3 q_1 \theta - L_1(S_1^0) + \beta q_1 F(S_1^0) [V_2(0) - V_1(0)] + \beta V_1(0),$$

and

$$V_2(0) = K_3 q_2 \theta - L_2(S_2^0) - \beta q_2 \bar{F}(S_2^0) [V_2(0) - V_1(0)] + \beta V_2(0).$$

Subtracting the first from the second equation and rearranging terms yields

$$V_2(0) - V_1(0) = (K_3(q_2 - q_1)\theta - [L_2(S_2^0) - L_1(S_1^0)])(1 - \beta + \beta[q_2 \bar{F}(S_2^0) + q_1 F(S_1^0)])$$

From (2.33), $H'_\alpha(S_\alpha^0) = -L'_\alpha(S_\alpha^0) + \beta q_\alpha f(S_\alpha^0) [V_2(0) - V_1(0)] = 0$, where $L'_\alpha(S_\alpha^0) = K_1 - K_2 q_\alpha \bar{F}(S_\alpha^0)$, from (2.27). Note that $L'_1(0) > L'_2(0)$. There are three cases to consider:

Case 1. $S_1^0 = S_2^0 = 0$ and S_2^0 does not satisfy first-order conditions, i.e.,

$$L'_2(0) > \beta q_2 f(0) [V_2(0) - V_1(0)]$$

which can be rewritten as

$$(K_1 + K_2 q_2) (1 - \beta \bar{q}_2) > \beta q_2 f(0) (K_3 - K_2) (q_2 - q_1) \theta$$

after replacing $L'_2(0) = K_1 - K_2 q_2$ and $V_2(0) - V_1(0)$ from the expression above for $S_1^0 = S_2^0 = 0$.

Case 2. S_2^0 satisfies first-order conditions, i.e.,

$$L'_2(S_2^0) = \beta q_2 f(S_2^0) [V_2(0) - V_1(0)]$$

and $S_1^0 = 0$ and does not satisfy first-order conditions, i.e.,

$$LL'_1(0) > \beta q_1 f(0) [V_2(0) - V_1(0)],$$

which can be rewritten $K_1 - K_2 q_1 > \beta q_1 f(0) [V_2(0) - V_1(0)]$, where

$$V_2(0) - V_1(0) = (K_3 (q_2 - q_1) \theta - [L_2(S_2^0) - K_2 q_1 \theta]) / (1 - \beta + \beta q_2 \bar{F}(S_2^0))$$

Case 3. S_1^0 and S_2^0 satisfy first-order conditions, i.e.,

$$L'_\alpha(S_\alpha^0) = \beta q_\alpha f(S_\alpha^0) [V_2(0) - V_1(0)], \alpha = 1, 2$$

Replacing $L'_\alpha(S_\alpha^0)$ and $V_2(0) - V_1(0)$ from above yields (2.41). $\square$

Proof of Theorem 2.1. First, we will derive expressions for the average expected profit under each of the three candidate policies, denoted by $\tilde{\Pi}_{P_1}$, $\tilde{\Pi}_{P_M}$, and $\tilde{\Pi}_{P_{\alpha-1,\alpha}}$, $\in \{2, \dots, M\}$, and then we will compare their values. Under policy P_1 , the supplier's rating is absorbed in the lowest value 1, where she orders up to $S_1^0 < \theta$. From (2.44), $\tilde{\Pi}_{P_1}$ can be written as $\tilde{\Pi}_{P_1} = \Lambda_1(S_1^0) = (p - b) q_1 \theta + [(b + h) q_1 - h] S_1^0$. Under P_M , the supplier's rating is absorbed in M , where she orders up to $S_M^0 = \theta$. From (2.44),

$\tilde{\Pi}_{P_M}$ can be written as:

$$\tilde{\Pi}_{P_M} = \Lambda_M(\theta) = [(p+h)q_M - h]\theta. \quad (\text{A.8})$$

Finally, under $P_{\alpha-1,\alpha}, \alpha \in \{2, \dots, M\}$, the supplier's rating is absorbed in the set $\{\alpha-1, \alpha\}$. When in α , she orders up to $S_\alpha^0 < \theta$, and when in $\alpha-1$, she orders up to $S_{\alpha-1}^0 = \theta$. In fact, $\alpha-1$ and α are the states of a two-dimensional Markov chain, with $q_{\alpha-1}$ and q_α being the transition probabilities from $\alpha-1$ to α and from α to $\alpha-1$, respectively. The steady-state probabilities of $\alpha-1$ and α are $q_\alpha/(q_{\alpha-1} + q_\alpha)$ and $q_{\alpha-1}/(q_{\alpha-1} + q_\alpha)$, respectively. From (2.44), $\tilde{\Pi}_{P_{\alpha-1,\alpha}}$ can be written as:

$$\begin{aligned} \tilde{\Pi}_{P_{\alpha-1,\alpha}} &= \Lambda_{\alpha-1}(\theta) \frac{q_\alpha}{q_{\alpha-1} + q_\alpha} + \Lambda_\alpha(S_\alpha^0) \frac{q_{\alpha-1}}{q_{\alpha-1} + q_\alpha} \\ &= \frac{[(p+h)q_{\alpha-1} - h]q_\alpha\theta + [(p-b)q_\alpha\theta + ([b+h]q_\alpha - h)S_\alpha^0]q_{\alpha-1}}{q_{\alpha-1} + q_\alpha} \end{aligned} \quad (\text{A.9})$$

Next, we will show that P_1 and $P_{\alpha-1,\alpha}, \alpha \in \{2, \dots, M\}$, can be candidate overall optimal policies only if $S_1^0 = 0$ and $S_\alpha^0 = 0$, respectively. First, we show this for P_1 . Note that $\partial\tilde{\Pi}_{P_1}/\partial S_1^0 = (b+h)q_1 - h$. The optimal value of S_1^0 , denoted by S_1^{0*} , depends on the sign of $\partial\tilde{\Pi}_{P_1}/\partial S_1^0$. If $q_1 < h/(b+h)$, then $S_1^{0*} = 0$; otherwise, $S_1^{0*} = \theta^-$. If we replace S_1^0 with θ^- , we get $\tilde{\Pi}_{P_1} \leq [(p+h)q_1 - h]\theta \leq \tilde{\Pi}_{P_M}$, since $q_M \geq q_1$. Therefore, the option that policy P_1 is overall optimal (in the sense that $\tilde{\Pi}_{P_1} > \tilde{\Pi}_{P_M}$) and $S_1^{0*} = \theta^-$ is not feasible. Policy P_1 may be overall optimal only if $q_1 < h/(b+h)$, which implies that $S_1^{0*} = 0$. $\tilde{\Pi}_{P_1}$ for $S_1^0 = 0$ becomes:

$$\tilde{\Pi}_{P_1} = (p-b)q_1\theta. \quad (\text{A.10})$$

The result for $P_{\alpha-1,\alpha}, \alpha \in \{2, \dots, M\}$, is shown similarly. Namely, the optimal value of S_α^0 , denoted by S_α^{0*} , depends on the sign of

$$\partial\tilde{\Pi}_{P_{\alpha-1,\alpha}}/\partial S_\alpha^0 = [(b+h)q_\alpha - h]q_{\alpha-1}/(q_{\alpha-1} + q_\alpha).$$

If $q_\alpha < h/(b+h)$, then $S_\alpha^{0*} = 0$; otherwise, $S_\alpha^{0*} = \theta^-$. If we replace S_α^0 with θ^- in

(A.9), after some manipulations, we get $\tilde{\Pi}_{P_{\alpha-1,\alpha}} \leq [(p+h) 2q_{\alpha-1}q_{\alpha}/(q_{\alpha-1}+q_{\alpha}) - h] \theta \leq [(p+h)q_{\alpha} - h] \theta$. Noting that $q_M \geq q_{\alpha}$, it follows that $\tilde{\Pi}_{P_{\alpha-1,\alpha}} \leq \tilde{\Pi}_{P_M}$. Therefore, the option that policy $P_{\alpha-1,\alpha}, \alpha \in \{2, \dots, M\}$, is overall optimal (in the sense that $\tilde{\Pi}_{P_{\alpha-1,\alpha}} > \tilde{\Pi}_{P_M}$) and $S_{\alpha}^{0*} = \theta^-$ is not feasible. Policy $P_{\alpha-1,\alpha}$ may be overall optimal only if $q_{\alpha} < h/(b+h)$, which implies that $S_{\alpha}^{0*} = 0$. In this case, if we replace S_{α}^0 with zero in (A.9), we get:

$$\tilde{\Pi}_{P_{\alpha-1,\alpha}} = \frac{[(2p-b+h)q_{\alpha-1} - h]q_{\alpha}\theta}{q_{\alpha-1} + q_{\alpha}}. \quad (\text{A.11})$$

Moreover, $P_{\alpha-1,\alpha}$ may be overall optimal if $\tilde{\Pi}_{P_1} < \tilde{\Pi}_{P_{\alpha-1,\alpha}}$, which from (A.10) and (A.11) implies that $(p-b)q_1 < [(2p-b+h)q_{\alpha-1} - h]q_{\alpha}/(q_{\alpha-1} + q_{\alpha})$. From the assumption $\beta p > b$, it follows that for $\beta = 1$, the l.h.s. of the above inequality is positive. Therefore, the r.h.s. must also be positive, which means that $\lambda q_{\alpha-1} > h$, where λ is defined as $\lambda = 2p - b + h$. Note that $\lambda q_{\alpha} > \lambda q_{\alpha-1} > h$ from (2.5).

Next, we compare policies $P_{\alpha-1,\alpha}$ and $P_{\alpha,\alpha+1}$, for $2 < \alpha < M-1$, by considering the difference $\tilde{\Pi}_{P_{\alpha,\alpha+1}} - \tilde{\Pi}_{P_{\alpha-1,\alpha}}$. From (A.11), this difference can be written as follows:

$$\tilde{\Pi}_{P_{\alpha,\alpha+1}} - \tilde{\Pi}_{P_{\alpha-1,\alpha}} = \frac{\{(\lambda q_{\alpha} - h)(q_{\alpha-1} + q_{\alpha})q_{\alpha+1} - (\lambda q_{\alpha-1} - h)(q_{\alpha} + q_{\alpha+1})q_{\alpha}\}\theta}{(q_{\alpha-1} + q_{\alpha})(q_{\alpha} + q_{\alpha+1})}$$

The denominator in the r.h.s. of the above expression is positive. Therefore, the sign of the difference $\tilde{\Pi}_{P_{\alpha,\alpha+1}} - \tilde{\Pi}_{P_{\alpha-1,\alpha}}$ depends on the sign of the expression in the braces multiplying θ in the numerator. For this expression, we have: $(\lambda q_{\alpha} - h)(q_{\alpha-1} + q_{\alpha})q_{\alpha+1} - (\lambda q_{\alpha-1} - h)(q_{\alpha} + q_{\alpha+1})q_{\alpha} = \lambda q_{\alpha}^2(q_{\alpha+1} - q_{\alpha-1}) - h(q_{\alpha-1}q_{\alpha+1} - q_{\alpha}^2) > h[q_{\alpha}(q_{\alpha+1} - q_{\alpha-1}) - (q_{\alpha-1}q_{\alpha+1} - q_{\alpha}^2)] = h[q_{\alpha+1}(q_{\alpha} - q_{\alpha-1})(q_{\alpha+1} + q_{\alpha})] > 0$, where the first inequality follows from $\lambda q_{\alpha} > h$ and the second inequality follows from (2.5). Therefore, $\tilde{\Pi}_{P_{\alpha,\alpha+1}} > \tilde{\Pi}_{P_{\alpha-1,\alpha}}$. This further implies that the only policy that may be overall optimal, besides policies P_1 and P_M , is policy $P_{M-1,M}$. The resulting average expected profit from (A.11) is:

$$\tilde{\Pi}_{P_{M-1,M}} = \frac{[(2p-b+h)q_{M-1} - h]q_M\theta}{q_{M-1} + q_M}. \quad (\text{A.12})$$

To summarize, policies P_1 and $P_{M-1,M}$ may be overall optimal if $q_1 < h/(b+h)$ and $q_M < h/(b+h)$, respectively. Policy P_M is overall optimal if either $q_M > h/(b+h)$ or $q_M < h/(b+h)$ but $\tilde{\Pi}_{P_M} > \tilde{\Pi}_{P_1}$ and $\tilde{\Pi}_{P_M} > \tilde{\Pi}_{P_{M-1,M}}$. From (A.8),(A.10), and (A.12), the last two conditions can be written as $q_M > h/(p+h) + q_1(p-b)/(p+h)$ and $q_M > (q_{M-1}/q_M) h/(p+h) + q_{M-1}(p-b)/(p+h)$. Multiplying both sides of the second condition with $(p+h)q_M$, this condition can be rewritten as $(p+h)q_M^2 - (p-b)q_{M-1}q_M - hq_{M-1} > 0$. The l.h.s. of this inequality is a quadratic function in q_M with a positive and a negative solution. Using the positive solution, we get $q_M > \left[(p-b)q_{M-1} + \sqrt{[(p-b)q_{M-1}]^2 + 4(p+h)hq_{M-1}} \right] / [2(p+h)]$. Similarly, P_1 is overall optimal if $q_1 < h/(b+h)$, $\tilde{\Pi}_{P_1} > \tilde{\Pi}_{P_M}$, and $\tilde{\Pi}_{P_1} > \tilde{\Pi}_{P_{M-1,M}}$. From (A.8),(A.10), and (A.12) these last two conditions can be written as

$$q_M < \frac{h}{p+h} + \frac{q_1(p-b)}{p+h}$$

and

$$q_M < \frac{(p-b)q_1q_{M-1}}{(p-b)(q_{M-1}-q_1) + (p+h)q_{M-1} - h}.$$

□

Proof of Proposition 2.7. Expression (2.47) can be written as

$$L(y) = (K_1y + K_2B(y) + \hat{b}q\bar{F}(y))1_{\{y \geq 0\}} + (K_1y + K_2[q(\theta - y) - \bar{q}y] + \hat{b}q)1_{\{y < 0\}}.$$

The first and second derivatives of the above expression are

$$L'(y) = (K_1 - K_2q\bar{F}(y) - \hat{b}qf(y))1_{\{y \geq 0\}} + (K_1 - K_2)1_{\{y < 0\}}.$$

and,

$$L''(y) = (K_2qf(y) - \hat{b}qf'(y))1_{\{y \geq 0\}}$$

Clearly, $L'(y) < 0, y < 0$, since $K_1 - K_2 < 0$ from (2.9). This implies that all the minima of $L(y)$ are non-negative. For $\Lambda(y), y \geq 0$, to be concave, we need $\Lambda''(y) \leq 0, y \geq 0$, or equivalently $L''(y) \geq 0, y \geq 0$, which can be rewritten as (2.48). If

(2.48) holds, the unique minimizer of $L(y)$ is given by $S^{my} = \arg \min_{y \geq 0} \{L'(y) \geq 0\}$, which can be written as (2.49). $\square$

Proof of Proposition 2.8. Consider the SD model under a basestock policy with a single basestock level S for all ratings, i.e., $S_\alpha = S, \alpha \in A$. Under this policy, $\{\alpha_t, t \geq 0\}$ is a discrete-time birth-death process with state-space A and non-zero transition probabilities $p_{\alpha, \alpha-1} = q_\alpha \bar{F}(S), 1 < \alpha \leq M, p_{\alpha, \alpha+1} = q_\alpha F(S), 1 \leq \alpha < M-1$, and $p_{\alpha, \alpha} = \bar{q}_\alpha + q_\alpha \bar{F}(S) 1_{\{\alpha=1\}} + q_\alpha F(S) 1_{\{\alpha=M\}}, 1 \leq \alpha \leq M$. Using standard Markov chain analysis, the steady-state probabilities for this process, $\pi_\alpha(S)$, are:

$$\pi_\alpha(S) = \frac{\pi_1(S) q_1 \Phi(S)^{\alpha-1}}{q_\alpha}, \quad \alpha = 2, \dots, M$$

and,

$$\pi_1(S) = \left[\frac{q_1 \sum_{\alpha \in A} \Phi(S)^{\alpha-1}}{q_\alpha} \right]^{-1}$$

where $\Phi(S) = F(S) / \bar{F}(S)$. Using the above expressions,

$$\tilde{q}(S) = \sum_{\alpha \in A} \pi_\alpha(S) q_\alpha = \pi_1(S) q_1 \sum_{\alpha \in A} \Phi(S)^{\alpha-1},$$

which leads to (2.53). Similarly, we have that

$$\tilde{\Pi}(S) = \sum_{\alpha \in A} \pi_\alpha(S) \Lambda_\alpha(S) = [K_3 \theta - K_2 B(S)] \sum_{\alpha \in A} \pi_\alpha(S) q_\alpha - K_1 S \sum_{\alpha \in A} \pi_\alpha(S),$$

which leads to (2.52). The global maximizer of $\tilde{\Pi}(S)$, S^* , is given by (2.50). Once S^* is found, $\tilde{q}(S^*)$ is given by (2.53).

To find the imputed $\hat{b}^*$ in the FS model, we must solve (2.49) for $\hat{b}$, after setting $S^{my} = S^*$ and $q = \tilde{q}(S^*)$. There are two cases to consider: $S^* > 0$ and $S^* = 0$. If $S^* > 0$, then $\hat{b}^*$ is obtained from (2.49) which leads to the top expression in (2.51). If $S^* = 0$, $F(0) = 1 - \bar{F}(0) = 0$, hence $\pi_1 = 1$, $\pi_\alpha = 0, \alpha > 1$, and $\tilde{q}(0) = q_1$. From (2.49), $\hat{b}$ belongs to the interval given by the bottom expression in (2.51), where the right end of that interval is necessarily positive, i.e., $K_1 - q_1 K_2 > 0$. To see this,

note that if $S^* = 0$, then $\tilde{\Pi}'(0) = \tilde{q}'(0)(K_3 - K_2)\theta + q_1K_2 - K_1$. The first term in this expression is positive because $K_3 - K_2 > 0$ from (2.9) and $\tilde{q}'(S) > 0$, because $\tilde{q}(S)$ is increasing in S . If $q_1K_2 - K_1 \geq 0$, then $\tilde{\Pi}'(0) > 0$. However, this cannot be true, because it would imply that $S^* > 0$. Therefore, $q_1K_2 - K_1 < 0$, or equivalently, $K_1 - q_1K_2 > 0$. $\square$

Appendix B

Chapter 3 Supplemental Material

Proof of Theorem 3.1. Suppose that supplier i is highly ranked but fails to fully meet the buyer's demand. Then, his inventory level becomes negative, and his ranking turns low. In this case, it is optimal for him to order just enough to satisfy the backordered demand and end up with zero inventory. Ordering less would result in him not fully satisfying the backordered demand and receiving the full margin for it. Ordering more would result in him holding costly inventory that would remain unused because the buyer never selects the low-ranking supplier. Once supplier i 's inventory level reaches zero, it is optimal to keep it at zero as long as his ranking remains low. Therefore, $y_i^*(2) = 0$.

Now suppose that supplier j is highly ranked but fails to fully meet the buyer's demand. Then, his inventory level becomes negative, and his ranking turns low, while the inventory level of supplier i remains at zero and his ranking turns high. In this case, it is optimal for supplier i to order some quantity $s_i \geq 0$, ending up with inventory s_i , in anticipation of the buyer's demand in the next period; s_i is, therefore, a target inventory level. If the demand is greater than s_i , supplier i will fail to fully meet the demand, his inventory level will become negative, and his ranking will turn low. If the demand is less than or equal to s_i , he will fully meet the demand, his inventory level will drop below s_i , and his ranking will remain high. In this case, it is optimal for him to order up to the target level s_i again. Therefore, $y_i^*(1) = s_i$.

Based on the above, the ranking vector (α_i, α_j) of the two suppliers is a discrete-time Markov chain with two states: $(1, 2)$ and $(2, 1)$. When $(\alpha_i, \alpha_j) = (1, 2)$, customer i orders up to s_i , and his expected profit is given by (3.4), because he receives the buyer's demand. Customer j orders up to zero, and his expected profit is zero, because he does not get any demand. When $(\alpha_i, \alpha_j) = (2, 1)$, the reverse is true. The probability of switching from $(1, 2)$ to $(2, 1)$ is $\bar{F}(s_i)$, while the probability of switching from $(2, 1)$ to $(1, 2)$ is $\bar{F}(s_j)$. It is trivial to show that the steady-state probabilities of states $(1, 2)$ and $(2, 1)$ are $\bar{F}(s_j)/[\bar{F}(s_j) + \bar{F}(s_i)]$ and $\bar{F}(s_i)/[\bar{F}(s_j) + \bar{F}(s_i)]$, respectively, which implies (3.10). $\square$

Proof of Proposition 3.1. From (3.4)–(3.8), $G_i(s_i) > 0$, $G'_i(s_i) \geq 0$, for $0 \leq s_i \leq s_i^m$, and $G_i(s_i) \leq 0$, $G'_i(s_i) < 0$, for $s_i \geq s_i^M$. From (3.13), this means that $\partial \Pi_i(s_i, s_j)/\partial s_i > 0$, for $0 \leq s_i \leq s_i^m$, and $\partial \Pi_i(s_i, s_j)/\partial s_i < 0$, for $s_i \geq s_i^M$, which implies (3.17). $\square$

Proof of Theorem 3.2. (i) The sign of $\phi_i(s_i, s_j)$ given by (3.14) determines the sign of $\partial \Pi_i(s_i, s_j)/\partial s_i$ given by (3.13), since the term $\bar{F}(s_j)/[\bar{F}(s_j) + \bar{F}(s_i)]^2$ in (3.13) is positive. As mentioned in the discussion following Proposition 3.1, $\partial \Pi_i(s_i, s_j)/\partial s_i > 0$, for $0 \leq s_i \leq s_i^m$, and $\partial \Pi_i(s_i, s_j)/\partial s_i < 0$, for $s_i \geq s_i^M$, implying that $\phi_i(s_i, s_j) > 0$, for $0 \leq s_i < s_i^m$, and $\phi_i(s_i, s_j) < 0$, for $s_i \geq s_i^M$. If condition (3.18) holds, then $\phi_i(s_i, s_j)$ is decreasing in s_i , for $s_i \in (s_i^m, s_i^M)$. As a result, the first-order condition $\partial \Pi_i(s_i, s_j)/\partial s_i = 0$, which reduces to $\phi_i(s_i, s_j) = 0$, has a unique solution, $s_i^*(s_j)$, satisfying (3.19).

(ii) The derivative of $s_i^*(s_j)$ with respect to s_j is given by (3.20) by using implicit differentiation. The denominator of the right-hand-side of (3.20) is given by (3.18) and is negative; therefore, the sign of $\partial s_i^*(s_j)/\partial s_j$ is determined by the sign of the numerator, which is given by (3.16). This quantity is positive for $s_i \in (s_i^m, s_i^M)$, because $G'_i(s_i) < 0$, for $s_i > s_i^m$; therefore, $\partial s_i^*(s_j)/\partial s_j > 0$.

(iii) The derivative of $s_i^*(s_j)$ with respect to any parameter q can be similarly computed as $\partial s_i^*(s_j)/\partial q = -[\partial \phi_i(s_i^*(s_j), s_j)/\partial q]/[\partial \phi_i(s_i^*(s_j), s_j)/\partial s_i]$. The denominator is again given by (3.18), so it is negative; therefore, the sign of $\partial s_i^*(s_j)/\partial q$ is determined by the sign of the numerator, which from (3.14) is $\partial \phi_i(s_i, s_j)/\partial q =$

$(\bar{F}(s_j) + \bar{F}(s_i)) \partial G'_i(s_i)/\partial q + f(s_i) \partial G_i(s_i)/\partial q$. If we substitute $f(s_i)$ from (3.19), this can be rewritten as $\partial \phi_i(s_i, s_j)/\partial q = [(\bar{F}(s_j) + \bar{F}(s_i)) / G_i(s_i)] [G_i(s_i) \partial G'_i(s_i)/\partial q - G'_i(s_i) \partial G_i(s_i)/\partial q]$. The sign of this expression is determined by the sign of the term in the second square bracket, where $G_i(s_i) > 0$ and $G'_i(s_i) < 0$, for $s_i \in (s_i^m, s_i^M)$, as mentioned earlier. From (3.4) and (3.5), we have the following: (a) For $q = \theta$, $\partial G_i(s_i)/\partial \theta = h_i + b_i > 0$ and $\partial G'_i(s_i)/\partial \theta = 0$; therefore, $\partial \phi_i(s_i, s_j)/\partial \theta > 0$. (b) For $q = p_i$, $\partial G_i(s_i)/\partial p_i = \theta > 0$ and $\partial G'_i(s_i)/\partial p_i = 0$; therefore, $\partial \phi_i(s_i, s_j)/\partial p_i > 0$. (c) For $q = h_i$, $\partial G_i(s_i)/\partial h_i = \theta - s_i - E[(w - s_i)^+] = -E[(s_i - w)^+] < 0$ and $\partial G'_i(s_i)/\partial h_i = -F(s_i) < 0$; therefore, $\partial \phi_i(s_i, s_j)/\partial h_i < 0$. (d) For $q = b_i$, $\partial G_i(s_i)/\partial b_i = -E[(w - s_i)^+] < 0$ and $\partial G'_i(s_i)/\partial b_i = \bar{F}(s_i) > 0$, so the term in the second square bracket of the numerator becomes $\bar{F}(s_i)[(h_i + p_i)\theta - h_i s_i] - h_i E[(w - s_i)^+]$. This expression equals $p_i \theta > 0$, for $s_i = 0$, goes to zero as $s_i \rightarrow \infty$, and its derivative with respect to s_i is $-f(s_i)[(h_i + p_i)\theta - h_i s_i] < 0$. Hence, it is positive and decreasing, implying that $\partial \phi_i(s_i, s_j)/\partial b_i > 0$. $\square$

Proof of Theorem 3.3. (i) By Theorem 3.2, under (3.18), $s_i^*(s_j)$ is increasing in s_j and is the unique solution of (3.14), for $i = 1, 2$. Moreover, from Proposition 3.1, $s_i^*(s_j)$ is bounded from above and below by s_i^M and s_i^m , respectively, for $i = 1, 2$. Therefore, the two best response functions cross each other at least at one point, as shown in Figure 3.2. Each point satisfies (3.14) for $i = 1, 2$.

(ii) By Theorem 3.2, $s_i^*(s_j)$ is increasing in θ, p_i , and b_i and decreasing in h_i . Moreover, it is increasing in s_j , where $s_j^*(s_i)$ is itself increasing in θ, p_j and b_j , and decreasing in h_j , implying the result.

(iii) By Theorem 4 in Cachon and Zhang (2006), if the best response mapping is a contraction on the entire strategy space, there is a unique Nash equilibrium. For a two-player game, this condition reduces to the requirement that the absolute value of the derivative of the best response function of each player must be less than one. In our case, this derivative is positive, so the condition further reduces to (3.23). $\square$

Proof of Proposition 3.2. By Theorem 3.3, if condition (3.18) holds, there exists at least one pure-strategy Nash equilibrium. Consider an arbitrary equilibrium (s_i^e, s_j^e) . Both s_i^e and s_j^e satisfy the first-order condition (3.19), i.e., $\phi_i(s_i^e, s_j^e) = 0$ and

$\phi_j(s_j^e, s_i^e) = 0$. Because of symmetry, $\phi_i(x, y) = \phi_j(x, y) = \phi(x, y)$, so the first-order conditions can be written as $\phi(s_i^e, s_j^e) = (\bar{F}(s_j^e) + \bar{F}(s_i^e))G'(s_i^e) + f(s_i^e)G(s_i^e) = 0$ and $\phi(s_j^e, s_i^e) = (\bar{F}(s_i^e) + \bar{F}(s_j^e))G'(s_j^e) + f(s_j^e)G(s_j^e) = 0$.

Suppose that (s_i^e, s_j^e) is asymmetric, where without loss of generality $s_j^e > s_i^e$. If condition (3.18) holds, then $\phi(x, y)$ is decreasing in x , and since $s_j^e > s_i^e$, we have that $\phi(s_j^e, s_j^e) < \phi(s_i^e, s_j^e)$, where $\phi(s_j^e, s_j^e) = (\bar{F}(s_j^e) + \bar{F}(s_j^e))G'(s_j^e) + f(s_j^e)G(s_j^e)$. Moreover, noting that $\bar{F}(s_j^e) < \bar{F}(s_i^e)$ since $s_j^e > s_i^e$, and $G'(s_j^e) < 0$ since from (3.17) $s_j^e \in (s_j^m, s_j^M)$, we have that $\phi(s_j^e, s_i^e) < \phi(s_j^e, s_j^e)$. To summarize, $s_j^e > s_i^e$ implies that $\phi(s_j^e, s_i^e) < \phi(s_j^e, s_j^e) < \phi(s_i^e, s_j^e)$ which further implies that the first-order conditions, $\phi(s_i^e, s_j^e) = 0$ and $\phi(s_j^e, s_i^e) = 0$, cannot both hold. Therefore, (s_i^e, s_j^e) , $s_j^e > s_i^e$, cannot be a pure-strategy equilibrium, so all equilibria are symmetric, proving (i) and (ii).

For any symmetric equilibrium (s^e, s^e) , the first-order conditions reduce to the single condition $\phi(s^e, s^e) = \hat{\phi}(s^e) = 2\bar{F}(s^e)G'(s^e) + f(s^e)G(s^e) = 0$, which can be rewritten as (3.24). If condition (3.26) holds, where

$$\frac{\partial \hat{\phi}(s)}{\partial s} = 2\bar{F}(s)G''(s) + f'(s)G(s) - f(s)G'(s), \quad (\text{B.1})$$

then $\hat{\phi}(s)$ is decreasing in s , and since $\hat{\phi}(0) > 0$, the first-order condition (3.24) has a unique solution, implying (iii). Note that, when $s_i^e = s_j^e = s^e$, the expected average demand share of each supplier is 50% of the total demand, so his payoff is given by (3.25). $\square$

Proof of Theorem 3.4. Figure B.1 shows a partition of the (s_i, s_j) space into four regions: A , B , C , and D , demarcated by s_i^m , s_i^M , s_j^m , and s_j^M . The diagonal curves traversing the regions represent contour lines along which the expected average demand shares of the suppliers are constant, i.e., they graphically represent function $\pi_i(s_i, s_j)$ defined in (3.11), for different demand share values $\pi_i \in (0, 1)$ of supplier i .

For any pair (s_i, s_j) in region B or C , including the Nash equilibrium under competition (s_i^e, s_j^e) , which is in region B , there exists a pair on the same contour line as (s_i, s_j) that belongs in region A or D (depending on whether the line passes through A or D) and has a higher payoff than (s_i, s_j) , because it is closer to s_i^m and s_j^m , the maximizers of $G_i(s_i)$ and $G_j(s_j)$, respectively. This implies that the optimal




Figure B.1: Constant demand-share contour lines $\pi_i(s_i, s_j)$ for different values of $\pi_i \in (0, 1)$.

active basestock level pair (s_i^c, s_j^c) belongs in region A or D , given by the following expression: $0 \leq s_i^c \leq s_i^m$ and $s_j^m \leq s_j^c < s_j^M$, for $(i, j) = (1, 2)$ or $(i, j) = (2, 1)$. It also implies (3.36). To see which of these two regions (s_i^c, s_j^c) belongs to, consider the following.

From (3.28) and (3.29), the first-order conditions $\partial \Pi(s_i, s_j) / \partial s_i = 0$, $i = 1, 2$, reduce to $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, since the term $\bar{F}(s_j) / [\bar{F}(s_j) + \bar{F}(s_i)]^2$ in (3.28) is positive. From (3.14) and (3.29), the conditions $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, can be written as $[\bar{F}(s_j) + \bar{F}(s_i)]G'_i(s_i) + f(s_i)[G_i(s_i) - G_j(s_j)] = 0$, $i = 1, 2$. At (s_i^m, s_j^m) , these conditions become $\psi_i(s_i^m, s_j^m) = f(s_i^m)[G_i(s_i^m) - G_j(s_j^m)] = 0$, $i = 1, 2$, since $G'_i(s_i^m) = G'_j(s_j^m) = 0$.

(i) If $G_i(s_i^m) = G_j(s_j^m)$, then $\psi_i(s_i^m, s_j^m) = 0$, $i = 1, 2$, so (s_i^m, s_j^m) is a solution of the first-order conditions. If there exists any other solution, say (s'_i, s'_j) , where $s'_i \neq s_i^m$ and/or $s'_j \neq s_j^m$, then from (3.27), $\Pi(s'_i, s'_j) = [\bar{F}(s'_j)G_i(s'_i) + \bar{F}(s'_i)G_j(s'_j)] / [\bar{F}(s'_j) + \bar{F}(s'_i)] < [\bar{F}(s'_j)G_i(s_i^m) + \bar{F}(s'_i)G_j(s_j^m)] / [\bar{F}(s'_j) + \bar{F}(s'_i)] = G_i(s_i^m) = G_j(s_j^m) = \Pi(s_i^m, s_j^m)$. Therefore, (s_i^m, s_j^m) is a global maximizer of $\Pi(s'_i, s'_j)$.

(ii) If $G_i(s_i^m) < G_j(s_j^m)$, then $\psi_i(s_i^m, s_j^m) < 0$ and $\psi_j(s_i^m, s_j^m) > 0$, so neither first-order condition is satisfied, i.e., (s_i^m, s_j^m) is not optimal. In fact, any point (s_i^m, s_j) and (s_i, s_j^m) does not satisfy the first-order conditions. From our previous analysis, the active basestock level of one supplier must be increased above his myopic basestock level and the active basestock level of the other supplier must be reduced below his myopic basestock level. At the optimal active basestock level pair (s_i^c, s_j^c) ,

conditions $\psi_i(s_i^c, s_j^c) = 0$, $i = 1, 2$, become $[\bar{F}(s_j^c) + \bar{F}(s_i^c)]G_i'(s_i^c) + f(s_i^c)[G_i(s_i^c) - G_j(s_j^c)] = 0$ and $[\bar{F}(s_j^c) + \bar{F}(s_i^c)]G_j'(s_j^c) + f(s_j^c)[G_j(s_j^c) - G_i(s_i^c)] = 0$. Suppose that $s_i^m < s_i^c < s_i^M$, $0 \leq s_j^c < s_j^m$. Then, $G_i'(s_i^c) < 0$, $G_j'(s_j^c) > 0$, and the above conditions imply that $G_i(s_i^c) > G_j(s_j^c)$. Then, from (3.27), $G_i(s_i^m) < [\bar{F}(s_j^m)G_i(s_i^m) + \bar{F}(s_i^m)G_j(s_j^m)]/[\bar{F}(s_j^m) + \bar{F}(s_i^m)] = \Pi(s_i^m, s_j^m) < \Pi(s_i^c, s_j^c) = [\bar{F}(s_j^c)G_i(s_i^c) + \bar{F}(s_i^c)G_j(s_j^c)]/[\bar{F}(s_j^c) + \bar{F}(s_i^c)] < G_i(s_i^c)$; however, this is impossible, since $G_i(s_i^m) > G_i(s_i^c)$. Therefore, $s_i^m < s_i^c < s_i^M$, $0 \leq s_j^c < s_j^m$ cannot be true, so it must be the case that $0 \leq s_i^c < s_i^m$, $s_j^m < s_j^c < s_j^M$. This implies that $G_i'(s_i^c) > 0$, $G_j'(s_j^c) < 0$, and $G_i(s_i^c) < G_j(s_j^c)$. Moreover, from (3.27), $G_i(s_i^c) < G_i(s_i^m) < \Pi(s_i^m, s_j^m) < \Pi(s_i^c, s_j^c) < G_j(s_j^c) < G_j(s_j^m)$. $\square$

Proof of Theorem 3.5. (i) From Theorem 3.4, the assumption $G_i(s_i^m) < G_j(s_j^m)$ implies that $s_i^c \in [0, s_i^m)$ and $s_j^c \in (s_j^m, s_j^M)$. As mentioned in the proof of Theorem 3.4, the first-order conditions $\partial\Pi(s_i, s_j)/\partial s_i = 0$, $i = 1, 2$, reduce to $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, which can be written as $-G_i'(s_i^c)/f(s_i^c) = [G_i(s_i^c) - G_j(s_j^c)]/[\bar{F}(s_j^c) + \bar{F}(s_i^c)]$, $i = 1, 2$, implying (3.41). If conditions (3.39) hold, then from (3.30), $\psi_i(s_i, s_j)$ is decreasing in s_i , $i = 1, 2$. There are two cases to consider.

Case (i-1): If $\psi_i(0, s_j^c) > 0$, equations $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, have a unique solution (s_i^c, s_j^c) , satisfying (3.40), with $s_i^c \in (0, s_i^m)$ and $s_j^c \in (s_j^m, s_j^M)$. Conditions (3.39) also imply that at (s_i^c, s_j^c) , $\partial^2\Pi(s_i^c, s_j^c)/\partial s_i^2 < 0$, $i = 1, 2$. From (3.28), (3.29), (3.31): $\partial^2\Pi(s_i^c, s_j^c)/\partial s_i \partial s_j = \{[-f_j(s_j^c)\psi_i(s_i^c, s_j^c) + \bar{F}(s_j^c)\partial\psi_i(s_i^c, s_j^c)/\partial s_j][\bar{F}(s_j^c) + \bar{F}(s_i^c)]^2 + 2[\bar{F}(s_j^c) + \bar{F}(s_i^c)]f(s_j^c)\bar{F}(s_j^c)\psi_i(s_i^c, s_j^c)\}/[\bar{F}(s_j^c) + \bar{F}(s_i^c)]^4 = 0$, since at (s_i^c, s_j^c) , $\psi_i(s_i^c, s_j^c) = 0$, from (3.29) and (3.40), and $\partial\psi_i(s_i^c, s_j^c)/\partial s_j = 0$, from (3.31) and (3.41). Therefore, the determinant of the Hessian at (s_i^c, s_j^c) satisfies:

$$[\partial^2\Pi(s_i^c, s_j^c)/\partial s_i^2][\partial^2\Pi(s_i^c, s_j^c)/\partial s_j^2] - [\partial^2\Pi(s_i^c, s_j^c)/\partial s_i \partial s_j]^2 > 0$$

So, (s_i^c, s_j^c) is a local maximum of $\Pi(s_i, s_j)$; however, since equations $\psi_i(s_i^c, s_j^c) = 0$, $i = 1, 2$, have a unique solution, (s_i^c, s_j^c) is a global maximum.

Case (i-2): If $\psi_i(0, s_j^c) \leq 0$, then $\Pi(s_i, s_j)$ is decreasing in s_i , for $s_i \geq 0$, implying that $s_i^c = 0$ and s_j^c uniquely satisfies the first-order condition $\partial\Pi(0, s_j)/\partial s_j = 0$, which reduces to (3.42).

(ii) We have that $0 = \psi_j(s_i^c, s_j^c) < \phi_j(s_i^c, s_j^c) < \phi_j(s_i^e, s_j^e)$, where the first equality holds from (3.40) and (3.41), the first inequality holds from (3.29), and the last inequality holds from (3.16) and the fact that $s_i^e > s_i^c$ from Theorem 3.4 (ii). From (3.18), $\phi_j(s_i, s_j)$ is decreasing in s_j , and since $\phi_j(s_i^e, s_j^e) > 0$, in order for $\phi_j(s_i^e, s_j^e) = 0$ to hold, it must be that $s_j^e > s_j^c$. $\square$

Proof of Proposition 3.3. (i) The proof follows from Theorem 3.2. First, note that if $f(w)$ is given by (3.46), $f'(w) < 0$; therefore, condition (3.18) immediately holds. To solve equation (3.19), it is sufficient to set the numerator of the right-hand-side of (B.8) at zero, i.e., solve $(\rho_i - \lambda s_i)e^{\lambda s_j} - (e^{\lambda s_i} - \beta_i) = 0$. For convenience, we define $z = \beta_i + \rho_i e^{\lambda s_j}$ and $m = e^{\lambda s_j}$ and rewrite the above equation as $e^{\lambda s_i} + m\lambda s_i = z$, or equivalently $m(z/m - \lambda s_i)e^{-\lambda s_i} = 1$. If, in addition, we set $u = z/m - \lambda s_i$ or $s_i = -(mu - z)/(\lambda m)$, the above equation can be rewritten as $ue^u = e^{z/m}/m$. Finally, by setting $ue^u = e^{z/m}/m = k$, this equation further reduces to $ue^u = k$. From property (ii) of the Lambert W function, the latter equation has a unique solution given by $W(k) = u$. Back substituting k, u, m , and z yields (3.49). (ii) The bounds in (3.50) follow directly from (3.7), (3.8), (3.46), and (3.49). $\square$

Proof of Proposition 3.4. The proof follows from Theorem 3.3. Firstly, recall from the proof of Proposition (3.3) that if $f(w)$ is given by (3.46), condition (3.18) immediately holds, and the solution of equation (3.19) is obtained by solving $(\rho_i - \lambda s_i)e^{\lambda s_j} - (e^{\lambda s_i} - \beta_i) = 0$, which can be rewritten as (3.51). To verify condition (3.23), we use property (iii) of the Lambert W function to compute the first derivative of $s_i^*(s_j)$ as follows:

$$\frac{\partial s_i^*(s_j)}{\partial s_j} = \frac{W\left(e^{\rho_i - \lambda s_j + \beta_i} e^{-\lambda s_j}\right) - \beta_i e^{-\lambda s_j}}{W\left(e^{\rho_i - \lambda s_j + \beta_i} e^{-\lambda s_j}\right) + 1}.$$

Clearly, $\partial s_i^*(s_j)/\partial s_j < 1$, since $-\beta_i e^{-\lambda s_j} < 1$, hence the best response is a contraction mapping on the entire strategy space. $\square$

Proof of Proposition 3.5. The proof follows from Theorem 3.4 and Theorem 3.5. First, note that if $f(w)$ is given by (3.46), $s_i^m = \ln(\beta_i)/\lambda$, from (3.7), and $G_i(s_i^m) = h_i(\rho_i - \ln(\beta_i))/\lambda$, from (B.4). Using the last equality, assumption (3.54) implies that

$G_i(s_i^m) < G_j(s_j^m)$, which, from Theorem 3.4 (ii), further implies that $0 \leq s_i^c < s_i^m$ and $s_j^m < s_j^c < s_j^M$. Therefore, $s_j^c > s_j^m > 0$ and $s_i^c \geq 0$.

Also note that $\psi_i(s_i, s_j)$ in (B.9) is the product of functions $e^{-\lambda(s_i+s_j)}$ and $L_i(s_i, s_j)$, where:

$$L_i(s_i, s_j) = e^{\lambda s_j} [h_i(\rho_i - \lambda s_i) - h_j(\rho_j - \lambda s_j)] - [h_i(e^{\lambda s_i} - \beta_i) + h_j(e^{\lambda s_j} - \beta_j)].$$

The first function is positive and decreasing in s_i . $L_i(s_i, s_j)$ is also decreasing in s_i , since $\partial L_i(s_i, s_j)/\partial s_i = -[\lambda h_i(e^{\lambda s_i} + e^{\lambda s_j})] < 0$. Therefore, for values of s_i for which $L_i(s_i, s_j) > 0$, $\psi_i(s_i, s_j)$ is decreasing in s_i , because it is the product of two positive, decreasing functions. For values of s_i for which $L_i(s_i, s_j) < 0$, $\psi_i(s_i, s_j) < 0$. This implies that the condition $\psi_i(s_i, s_j^c) = 0$ has at most one positive solution. There are two cases to consider.

Case (i): $L_i(0, s_j^c) > 0$. In this case, the condition $\psi_i(s_i, s_j^c) = 0$ has one positive solution $s_i^c > 0$, and therefore $\partial \Pi(s_i^c, s_j^c)/\partial s_i = 0$. To show this, assume that the equations $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, have a unique solution (s_i^c, s_j^c) , where $s_i^c > 0$ (recall that $s_j^c > 0$ by assumption (3.54)). Setting $\psi_i(s_i, s_j) = \psi_j(s_i, s_j) = 0$ in (B.9) yields: $e^{\lambda s_j^c} [h_i(\rho_i - \lambda s_i^c) - h_j(\rho_j - \lambda s_j^c)] = e^{\lambda s_i^c} [h_j(\rho_j - \lambda s_j^c) - h_i(\rho_i - \lambda s_i^c)]$. This equality holds only if the expressions in both square brackets of (B.9) are zero, yielding (3.55) and (3.56). Substituting s_j^c from (3.55) into (3.56) yields:

$$h_i \beta_i + h_j \beta_j - (h_i e^{\lambda s_i^c} + h_j e^{(\Delta p + h_i \lambda s_i^c)/h_j}) = 0.$$

Note that the left-hand side of the above equation is $L_i(s_i^c, s_j^c)$ which, seen as a function of s_i^c , is decreasing in s_i^c , as also mentioned earlier. For this equation to have a positive solution, it must be positive at $s_i^c = 0$, which implies that $\Delta p < h_j \ln(K)$. Reversing the arguments and assuming that $\Delta p < h_j \ln(K)$, leads to the conclusion that the equations $\psi_i(s_i, s_j) = 0$, $i = 1, 2$, have a unique solution (s_i^c, s_j^c) , where $s_i^c, s_j^c > 0$.

Case (ii): $L_i(0, s_j^c) \leq 0$. In this case, the condition $\psi_i(s_i, s_j^c) = 0$ has a zero solution or no solution, and therefore $s_i^c = 0$ and $\partial \Pi(s_i^c, s_j^c)/\partial s_i \leq 0$. To see this, note that if $\Delta p \geq h_j \ln(K)$, then $L_i(0, s_j^c) \leq 0$ and therefore the equation $\psi_i(s_i, s_j^c) = 0$ either has a zero solution or no solution, as mentioned earlier. In this case, $s_i^c = 0$ and

s_j^c uniquely satisfies $L_j(0, s_j^c) = 0$. The solution of this equation is given by (3.57). Again, if we reverse the arguments and substitute s_j^c into the inequality $L_i(0, s_i^c) \leq 0$, leads to the condition $\Delta p \geq h_j \ln(K)$. $\square$

Proof of Theorem 3.6. The bounds in (3.66) can be derived as in the proof of Proposition 3.1.

The proof of (i) and (iii) are similar to the corresponding assertions in Theorem 3.2.

To prove (ii), consider the first derivatives of $\phi_i(\mathbf{s})$ given by:

$$\frac{\partial \phi_i(\mathbf{s})}{\partial s_i} = \prod_{k \neq i} \bar{F}(s_k) G_i''(s_i) + \sum_{l \neq i} \prod_{k \neq l, i} \bar{F}(s_k) (\bar{F}(s_i) G_i''(s_i) + f'(s_i) G_i(s_i)), \quad (\text{B.2})$$

$$\frac{\partial \phi_i(\mathbf{s})}{\partial s_j} = -f(s_j) \left[\prod_{k \neq j, i} \bar{F}(s_k) G_i'(s_i) + \sum_{l \neq j, i} \prod_{k \neq l, j, i} \bar{F}(s_k) (\bar{F}(s_i) G_i'(s_i) + f(s_i) G_i(s_i)) \right], \quad (\text{B.3})$$

for $j \neq i$. The derivative of $s_i^*(\mathbf{s}_{-i})$ with respect to s_j is given by the fraction in (3.68) by using implicit differentiation. The denominator of that fraction is negative by condition (3.66) and the numerator is positive. To see this, set $\phi_i(\mathbf{s}) = 0$ and substitute $[\bar{F}(s_i) G_i'(s_i) + f(s_i) G_i(s_i)]$ from (3.64) into (B.3). The result is:

$$\frac{\partial \phi_i(s_i^*(\mathbf{s}_{-i}), \mathbf{s}_{-i})}{\partial s_j} = - \prod_{k \neq j, i} \bar{F}(s_k) f(s_j) G_i'(s_i^*(\mathbf{s}_{-i})) \left[1 - \frac{\sum_{l \neq j, i} \prod_{k \neq l, j, i} \bar{F}(s_k) \bar{F}(s_j)}{\sum_{l \neq i} \prod_{k \neq l, i} \bar{F}(s_k)} \right],$$

for $j \neq i$. The above expression is positive, because $G_i'(s_i^*(\mathbf{s}_{-i})) < 0$ and the term inside the square brackets is positive. $\square$

Expressions for exponentially distributed demand

If $f(w)$ is given by (3.46), expressions (3.4), (3.5), (3.7), (3.10), (3.14), and (3.29) become:

$$G_i(s_i) = \frac{h_i (\rho_i - \lambda s_i + 1 - \beta_i e^{-\lambda s_i})}{\lambda}, \quad (\text{B.4})$$

$$G_i'(s_i) = -h_i (1 - \beta_i e^{-\lambda s_i}), \quad (\text{B.5})$$

$$s_i^m = \frac{\ln(\beta_i)}{\lambda}, \quad (\text{B.6})$$

$$\Pi_i(s_i, s_j) = \frac{h_i(\rho_i - \lambda s_i + 1 - \beta_i e^{-\lambda s_i}) e^{\lambda s_i}}{\lambda(e^{\lambda s_i} + e^{\lambda s_j})}, \quad (\text{B.7})$$

$$\phi_i(s_i, s_j) = \frac{h_i[(\rho_i - \lambda s_i)e^{\lambda s_j} - (e^{\lambda s_i} - \beta_i)]}{e^{\lambda(s_i + s_j)}}. \quad (\text{B.8})$$

$$\psi_i(s_i, s_j) = \frac{e^{\lambda s_j} [h_i(\rho_i - \lambda s_i) - h_j(\rho_j - \lambda s_j)] - [h_i(e^{\lambda s_i} - \beta_i) + h_j(e^{\lambda s_j} - \beta_j)]}{e^{\lambda(s_i + s_j)}}. \quad (\text{B.9})$$

Moreover, from (3.8), (B.4), and (B.6), we have:

$$s_i^M = \frac{\rho_i + 1 + W(-\beta_i e^{-(\rho_i + 1)})}{\lambda}, \quad (\text{B.10})$$

$$G_i(s_i^m) = \frac{h_i(\rho_i - \ln(\beta_i))}{\lambda}. \quad (\text{B.11})$$

Appendix C

Chapter 4 Supplemental Material

Proof of Theorem 4.1. In the single-period problem, clearly, the optimal selection policy is the revenue-greedy policy, according to which active buyers are served in descending order of their revenue rate. Under this policy, the expected profit function $G(y) = \mathbb{E}[g(y)]$ for any order quantity $y \in \mathcal{B}_0$ is given by (4.13). The first-order difference of $G(y)$, denoted by $G_1(y)$, is:

$$G_1(y) = G(y) - G(y-1) = \sum_{i=y}^n f_{(i-1)}(y-1) q_{(i)} r_{(i)} - c, \quad y \in \mathcal{B}, \quad (\text{C.1})$$

The second-order difference of $G(y)$, denoted by $G_2(y)$, is:

$$G_2(y) = G_1(y+1) - G_1(y) = \sum_{i=y}^n (f_{(i-1)}(y) - f_{(i-1)}(y-1)) q_{(i)} r_{(i)}, \quad (\text{C.2})$$

for $y \in \{1, \dots, n-1\}$. Using conditioning, $f_{(i)}(y)$ can be written as follows:

$$f_{(i)}(y) = \sum_{j=y}^i f_{(j-1)}(y-1) q_{(j)} \prod_{k=j+1}^{n-1} \bar{q}_{(k)}, \quad i \geq y \in \mathcal{B}.$$

Substituting $f_{(i-1)}(y-1)$ from the above expression into (C.2) yields:

$$G_2(y) = \sum_{i=y}^n \left[\left(\sum_{j=y}^{i-1} f_{(j-1)}(y-1)q_{(j)} \prod_{k=j+1}^{n-1} \bar{q}_{(k)} \right) - f_{(i-1)}(y-1) \right] q_{(i)}r_{(i)},$$

where $y \in \{1, \dots, n-1\}$. Collecting terms and rearranging yields:

$$G_2(y) = \sum_{i=y}^n f_{(i-1)}(y-1)q_{(i)} (\eta_{(i)} - r_{(i)}), \quad (\text{C.3})$$

for $y \in \{1, \dots, n-1\}$, where $\eta_{(i)}$ is given by following recursion:

$$\eta_{(i)} = q_{(i+1)}r_{(i+1)} + \bar{q}_{(i+1)}\eta_{(i+1)}, \quad i = 1, \dots, n-1, \quad \text{and } \eta_{(n)} = 0. \quad (\text{C.4})$$

Next, we show by induction that $r_{(i)} > \eta_{(i)}$, $i = 1, \dots, n$. Clearly, $r_{(n)} > \eta_{(n)} = 0$. For $i = 1, \dots, n-1$, assume that $r_{(i+1)} > \eta_{(i+1)}$. Then, $r_{(i)} > r_{(i+1)} = q_{(i+1)}r_{(i+1)} + \bar{q}_{(i+1)}r_{(i+1)} > q_{(i+1)}r_{(i+1)} + \bar{q}_{(i+1)}\eta_{(i+1)} = \eta_{(i)}$, where the first inequality follows from (4.14) and the second follows from the induction hypothesis. Because $r_{(n)} > \eta_{(n)}$, $i = y, \dots, n$, all the terms in (C.3) are negative, and hence $G_2(y) < 0$. Therefore, $G(y)$ is concave and attains a unique maximum at $y^m = \arg \min_{y \in \{0, \dots, n-1\}} \{G_1(y+1) \leq 0\}$ given by (4.15). $\square$

Proof of Proposition 4.1. To show that $y^m(\alpha') \geq y^m(\alpha)$ for any two satisfaction state vectors α' and α such that $\alpha' \geq \alpha$, it suffices to show that $y^m(\alpha') \geq y^m(\alpha)$ for α' and α such that $\alpha'_j = 1$ and $\alpha_j = 0$, for some $j \in \mathcal{B}$, and $\alpha'_i = \alpha_i$, $i \neq j$. To simplify notation, let $f'_{(k)}(\cdot)$ and $f_{(k)}(\cdot)$ denote the p.m.f. of $D_{(k)}(\alpha')$ and $D_{(k)}(\alpha)$, respectively. To show that $y^m(\alpha') \geq y^m(\alpha)$, it suffices to show that $G'_1(y) \geq G_1(y)$, where $G'_1(y)$ and $G_1(y)$ denote the first-order differences $G_1(\alpha', y)$ and $G_1(\alpha, y)$, respectively. From (C.1), and assuming that $\alpha'_j = 1$ and $\alpha_j = 0$, for some $j \in \mathcal{B}$, and $\alpha'_i = \alpha_i$, $i \neq j$, we define the difference $\Delta G_1 = G'_1(y) - G_1(y)$ and we have:

$$\Delta G_1 = \begin{cases} \Delta q_{(j)}f_{(j-1)}(y-1)r_{(j)} + \sum_{i=j+1}^n (f'_{(i-1)}(y-1) - f_{(i-1)}(y-1))q_{(i)}r_{(i)}, & j \geq y, \\ 0, & j < y, \end{cases} \quad (\text{C.5})$$

$$\Delta G_1 = 1_{\{j \geq y\}} \left[\Delta q_{(j)} f_{(j-1)}(y-1) r_{(j)} + \sum_{i=j+1}^n (f'_{(i-1)}(y-1) - f_{(i-1)}(y-1)) q_{(i)} r_{(i)}, \right]$$

We will only consider the case $j \geq y$ because, for $j < y$, $\Delta G_1 = 0$. Let $D_{(i) \setminus \{j\}} = \sum_{k=1, k \neq j}^i d_{(k)}$ and let $f_{(i) \setminus \{j\}}(k)$ denote the p.m.f. of $D_{(i) \setminus \{j\}}$. Clearly, $f_{(i) \setminus \{j\}}(k) = f_{(i)}(k)$, $i < j$, and

$$f_{(j) \setminus \{j\}}(k) = f_{(j-1)}(k), \quad (\text{C.6})$$

$$f_{(i) \setminus \{j\}}(k) = f_{(i-1) \setminus \{j\}}(k-1) q_{(i)} + f_{(i-1) \setminus \{j\}}(k) \bar{q}_{(i)}, \quad i > j. \quad (\text{C.7})$$

We can express the p.m.f.'s in (C.5) as follows: $f'_{(i)}(y-1) = f_{(i) \setminus \{j\}}(y-2) q'_{(j)} + f_{(i) \setminus \{j\}}(y-1) \bar{q}'_{(j)}$ and $f_{(i)}(y-1) = f_{(i) \setminus \{j\}}(y-2) q_{(j)} + f_{(i) \setminus \{j\}}(y-1) \bar{q}_{(j)}$, $i \geq j$. Substituting the difference of the above expressions into (C.5), also substituting $f_{(j-1)}(k)$ from (C.6), yields:

$$\Delta G_1 = \Delta q_{(j)} (f_{(j) \setminus \{j\}}(y-1) r_{(j)} + \sum_{i=j+1}^n (f_{(i-1) \setminus \{j\}}(y-2) - f_{(i-1) \setminus \{j\}}(y-1)) q_{(i)} r_{(i)}).$$

Substituting $f_{(i-1) \setminus \{j\}}(y-1)$ from (C.7) into the above expression, unfolding the recursion, and collecting terms yields:

$$\Delta G_1 = \Delta q_{(j)} (f_{(j) \setminus \{j\}}(y-1) (r_{(j)} - \eta_{(j)}) + \sum_{i=j+1}^n f_{(i-1) \setminus \{j\}}(y-2) q_{(i)} (r_{(i)} - \eta_{(i)})),$$

where $\eta_{(i)}$ is given by (C.4). As was shown in the proof of Proposition 1, $r_{(i)} > \eta_{(i)}$, $i \in \mathcal{B}$; therefore, all the terms in the above equation are positive, and hence $\Delta G_1 > 0$. $\square$

Proof of Proposition 4.2. Consider the states $\alpha' = (\alpha'_1, \alpha'_2, \dots, \alpha'_n)$ and $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, such that $\alpha'_j = \alpha_j$, $\forall j \neq i$, and, $\alpha'_i = 0, \alpha_i = 1$, for some $i \in \mathcal{B}$, then $\alpha' \geq \alpha$. Consider two sample paths, a nominal path starting from α and following the optimal ordering policy and buyer selection policy, and an alternative path starting from α' and following the actions of the nominal path. If buyer i is

active in the nominal path, she is also active in the alternative path. If the firm serves (does not serve) this buyer in the nominal path, it also serves (does not serve) her in the alternative path. Therefore, her satisfaction state becomes 1 (0) in both sample paths, and the two paths become identical thereafter. If buyer i is inactive in the nominal path, she is either also inactive or active in the alternative path. In the first case, her satisfaction state in both sample paths remains unchanged and the argument is repeated in the next period. In the second case, her satisfaction state in the nominal path remains unchanged, i.e., 0. In the alternative path, if there is no excess capacity, the firm does not serve the buyer. In this case, her satisfaction state becomes 0, and the total satisfaction states in the two sample paths become identical. If there is excess capacity, the firm serves the buyer, receiving a reward r_i , and the satisfaction state of the buyer remains unchanged. In this case, the satisfaction state of the buyer in both sample paths remains unchanged and the argument is repeated. To summarize, the satisfaction state of buyer i in the alternative path either becomes the same as that in the nominal path or is greater than that in the nominal path. $\square$

Proof of Proposition 4.3. From (4.1), $\gamma_i \geq \gamma_j$ implies $q_i(1)q_j(0) \geq q_i(0)q_j(1)$, which implies $P(d_i(1) + d_j(0) = 2) \geq P(d_i(0) + d_j(1) = 2)$, which implies $P(d_i(1) + d_j(0) \geq 2) \geq P(d_i(0) + d_j(1) \geq 2)$. From (4.2), $\bar{\gamma}_i \geq \bar{\gamma}_j$ implies $\bar{q}_i(1)\bar{q}_j(0) \leq \bar{q}_i(0)\bar{q}_j(1)$, which implies $P(d_i(1) + d_j(0) = 0) \leq P(d_i(0) + d_j(1) = 0)$, which implies $P(d_i(1) + d_j(0) \geq 1) \geq P(d_i(0) + d_j(1) \geq 1)$. Therefore, $\gamma_i \geq \gamma_j$ and $\bar{\gamma}_i \geq \bar{\gamma}_j$ imply $d_i(1) + d_j(0) \geq_{st} d_i(0) + d_j(1)$. $\square$

Proof of Proposition 4.4. If $y(\alpha) = n - 1$, $\forall \alpha$, the buyer selection policy matters only when all buyers are active. In this case, the question is not who to select but who to leave out. Suppose that when all buyers are active, buyer j has the lowest priority and is left out. In this case, she becomes dissatisfied, and all other buyers become satisfied, i.e., the satisfaction state vector becomes $\alpha_j = (\alpha_1, \dots, \alpha_n : \alpha_j = 0, \alpha_k = 1, k \in \mathcal{B} \setminus \{j\})$. Once in state α_j , the satisfaction state vector will remain in α_j until buyer j is served. Given that she has the lowest priority, this will happen only if she is active and at least one of the other buyers is inactive. The probability of this event is $q_j(0)F_{-i}^1(n-2) = q_j(0)[1 - \prod_{k \in \mathcal{B} \setminus \{i\}} q_k(1)]$. When this event happens,

buyer j becomes satisfied and all other buyers remain satisfied, i.e., the satisfaction state vector becomes $\alpha_1 = (\alpha_1, \dots, \alpha_n : \alpha_k = 1, k \in \mathcal{B})$. Once in state α_1 , the satisfaction state will remain in α_1 until all buyers are active again and, as before, buyer j is not selected. The probability of this event is $\prod_{k \in \mathcal{B}} q_k(1)$. When this event happens, buyer j becomes dissatisfied, all other buyers remain satisfied, and the cycle repeats. Therefore, when buyer j has the lowest priority, the buyer satisfaction state vector is a Markov chain with two states, α_j and α_1 , and transition probabilities $p_{\alpha_1 \alpha_j} = \prod_{k \in \mathcal{B}} q_k(1)$ and $p_{\alpha_j \alpha_1} = q_j(0)[1 - \prod_{k \in \mathcal{B} \setminus \{j\}} q_k(1)]$. The stationary probabilities of these two states are $\pi_{\alpha_j} = \prod_{k \in \mathcal{B}} q_k(1) / (\prod_{k \in \mathcal{B}} q_k(1) + q_j(0)[1 - \prod_{k \in \mathcal{B} \setminus \{j\}} q_k(1)])$ and $\pi_{\alpha_1} = 1 - \pi_{\alpha_j}$. Every time the Markov chain makes a transition to state α_j , buyer j is not served, and every time it makes a transition to state α_1 , buyer j is served. Therefore, the expected contribution of buyer j to the firm's revenue is $\pi_{\alpha_1} q_j(1) r_j = (1 - \pi_{\alpha_j}) q_j(1) r_j$. The expected contribution of the other buyers, who always remain satisfied, is $\sum_{k \in \mathcal{B} \setminus \{j\}} q_k(1) r_k$. Thus, the total average expected profit of the firm when buyer j has the lowest priority is $\sum_{k \in \mathcal{B} \setminus \{j\}} q_k(1) r_k + (1 - \pi_{\alpha_j}) q_j(1) r_j - (n - 1) c$, which, after some algebraic manipulations, can be rewritten as (4.18), where z_j is given by (4.17) for $i = j$. So, if the firm wants to maximize its average expected profit, it must assign the lowest priority to the buyer with the smallest value of z_i . Thus, z_i is the index, and the optimal average expected profit is given by (4.18). $\square$

Proof of Theorem 4.2. When $n = 2$, buyer selection matters only when $y = 1$ and both buyers are active. In this case, if the firm selects buyer i over j , the satisfaction state vector becomes $(\alpha_i = 1, \alpha_j = 0)$, regardless of its initial value, implying that the optimal selection policy is an index policy. Assume for the moment that the optimal ordering policy is an FOQ policy, i.e., $y^*(\alpha) = y^z$, where $y^z = 0, 1$, or 2 . If $y^z = 1$, then from Proposition 4.4, the optimal selection policy is an index policy u^z with index for buyer i given by (4.17), which for $n = 2$ reduces to (4.22). If $y^z = 0$ or 2 , buyer selection is irrelevant as was mentioned earlier and any policy including u^z is optimal.

Next, we will show that the optimal ordering policy satisfies the monotonicity property in Conjecture 1, namely, $y^*(\alpha') \geq y^*(\alpha), \alpha' \geq \alpha$, and then, we will show that in fact $y^*(\alpha) = y^z$, where $y^z = 0, 1$, or 2 . To this end, suppose that the firm uses

index policy $\mathbf{u}^z$ and without loss of generality assume that $z_i > z_j$, $i \neq j \in \mathcal{B} = \{1, 2\}$, i.e., buyer i has priority over j . Then, the optimality equation (4.10) can be written as $\Pi = \max_{y \in \mathcal{B}_0} \{\mathbb{E}_{\mathbf{d}}[g(y, \mathbf{u}^*)] + G(\boldsymbol{\alpha}, y, \mathbf{u}^*)\}$, where:

$$\mathbb{E}_{\mathbf{d}}[g(y, \mathbf{u}^*)] = 1_{\{y \neq 0\}}(r_i q_i(\alpha_i) - c) + 1_{\{y=2\}}(r_j q_j(\alpha_j) - c) + 1_{\{y=1\}} r_j \bar{q}_i(\alpha_i) q_j(\alpha_j), \quad (\text{C.8})$$

$$G(\boldsymbol{\alpha}, y, \mathbf{u}^*) = \mathbb{E}_{\mathbf{d}}[V(\Phi(\boldsymbol{\alpha}, \mathbf{u}^* \mathbf{d}))] - V(\boldsymbol{\alpha}). \quad (\text{C.9})$$

Consider the optimal order quantity when both buyers are dissatisfied, $y^*(0, 0)$. There are two cases.

Case 1: $y^*(0, 0) = 0$. In this case, state $\boldsymbol{\alpha} = (0, 0)$ is absorbing. This case arises if the firm cannot be profitable in any state, so it is better off setting $y^*(\boldsymbol{\alpha}) = 0$, $\forall \boldsymbol{\alpha}$, and hence $\Pi = 0$. Clearly, the monotonicity property holds, and in fact, the optimal ordering policy is an FOQ policy with FOQ $y^z = 0$.

Case 2: $y^*(0, 0) > 0$. In this case, the firm can be profitable in some states, so $\Pi > 0$. We will show that $y^*(\boldsymbol{\alpha}) > 0$, $\boldsymbol{\alpha} \neq (0, 0)$. To this end, assume that the reverse is true, i.e., $y^*(\boldsymbol{\alpha}) = 0$, for some $\boldsymbol{\alpha} \neq (0, 0)$. There are three subcases to consider. **Subcase 2-i:** $\boldsymbol{\alpha} = (1, 1)$, so $G(\boldsymbol{\alpha}, 0, \mathbf{u}^*) = q_1(1)\bar{q}_2(1)(V(0, 1) - V(1, 1)) + \bar{q}_1(1)q_2(1)(V(1, 0) - V(1, 1)) + q_1(1)q_2(1)(V(0, 0) - V(1, 1))$. **Subcase 2-ii:** $\boldsymbol{\alpha} = (0, 1)$, so $G(\boldsymbol{\alpha}, 0, \mathbf{u}^*) = q_2(1)(V(0, 0) - V(0, 1))$. **Subcase 2-iii:** $\boldsymbol{\alpha} = (1, 0)$, so $G(\boldsymbol{\alpha}, 0, \mathbf{u}^*) = q_1(1)(V(0, 0) - V(1, 0))$. In all subcases, $\mathbb{E}_{\mathbf{d}}[g(0, \mathbf{u}^*)] = 0$ by (C.8) and $G(\boldsymbol{\alpha}, 0, \mathbf{u}^*) \leq 0$ by Proposition 4.2, which means that the firm makes no profit or even incurs losses. However, this is impossible, given that the assumption $y^*(0, 0) > 0$ implies $\Pi > 0$; therefore, $y^*(\boldsymbol{\alpha}) > 0$, $\boldsymbol{\alpha} \neq (0, 0)$. To show that $y^*(\boldsymbol{\alpha}') \geq y^*(\boldsymbol{\alpha})$, $\boldsymbol{\alpha}' \geq \boldsymbol{\alpha}$, it suffices to show that $\arg \max_{y \in \{1, 2\}} \{\mathbb{E}_{\mathbf{d}}[g(y, \mathbf{u}) + G(\boldsymbol{\alpha}', y, \mathbf{u}^*)]\} \geq \arg \max_{y \in \{1, 2\}} \{\mathbb{E}_{\mathbf{d}}[g(y, \mathbf{u}) + G(\boldsymbol{\alpha}, y, \mathbf{u}^*)]\}$, $\boldsymbol{\alpha}' \geq \boldsymbol{\alpha}$. To this end, we must show that $\Delta_{\boldsymbol{\alpha}'} g(\mathbf{u}) \geq \Delta_{\boldsymbol{\alpha}} g(\mathbf{u})$ and $\Delta G(\boldsymbol{\alpha}', \mathbf{u}^*) > \Delta G(\boldsymbol{\alpha}, \mathbf{u}^*)$, $\boldsymbol{\alpha}' \geq \boldsymbol{\alpha}$, where $\Delta_{\boldsymbol{\alpha}} g(\mathbf{u}) = \mathbb{E}_{\mathbf{d}}[g(2, \mathbf{u}) - g(1, \mathbf{u})]$ and $\Delta G(\boldsymbol{\alpha}, \mathbf{u}^*) = G(\boldsymbol{\alpha}, 2, \mathbf{u}^*) - G(\boldsymbol{\alpha}, 1, \mathbf{u}^*)$. Clearly, $\Delta_{\boldsymbol{\alpha}} g(\mathbf{u}) = r_j q_i(a_i) q_j(a_j) - c$ from (C.8) and $\Delta G(\boldsymbol{\alpha}, \mathbf{u}^*) = q_i(\alpha_i) q_j(a_j)(V(1, 1) - V(1, 0))$ from (C.9), which implies the result.

To complete the proof, note that any ordering policy satisfying $y(\boldsymbol{\alpha}') \geq y(\boldsymbol{\alpha})$, $\boldsymbol{\alpha}' \geq \boldsymbol{\alpha}$, drives $\boldsymbol{\alpha}_t$ to an absorbing state (or set of two states) that can also be reached under

an FOQ policy. Specifically, if $y(\alpha) = 0, 1$ or $2, \forall \alpha$, then α_t is absorbed in state $\alpha = (0, 0)$, the set of states $\alpha = (1, 0)$ or state $\alpha = (1, 1)$, respectively. Therefore, the initial assumption that $y^*(\alpha) = y^z$, where $y^z = 0, 1$, or 2 is verified.

To find the exact value of y^z , it suffices to compare the average expected profit Π^{y, u^z} under index policy u^z and FOQ policy $y(\alpha) = y, \forall \alpha$, for $y = 0, 1, 2$. Clearly, if $y = 0$, both buyers are always dissatisfied and $\Pi^{0, u^z} = 0$. Similarly, if $y = 2$, both buyers are always satisfied and $\Pi^{2, u^z} = R - 2c$, where R is given by (4.19). Finally, if $y = 1$, $\Pi^{1, u^z} = R - R_j - c$ from (4.18), where R_j is given by (4.20). Moreover, it can be shown that if $\Pi^{1, u^z} < \Pi^{0, u^z} = 0$, then $\Pi^{2, u^z} < \Pi^{1, u^z} < \Pi^{0, u^z}$. Therefore, $y^z = 0$ if $\Pi^{0, u^z} > \Pi^{1, u^z}$, $y^z = 2$ if $\Pi^{2, u^z} > \Pi^{1, u^z}$, and $y^z = 1$ if $\Pi^{1, u^z} > \max(\Pi^{0, u^z}, \Pi^{2, u^z})$, leading to the conditions in Table 4.1. $\square$

Proof of Proposition 4.5. Substituting $\varphi(\alpha_i, d_i, u_i)$ and $\Phi(\alpha, \mathbf{d}, \mathbf{u})$ from (4.7)–(4.8) and $\hat{V}^\lambda(\alpha)$ and $\hat{\Pi}^\lambda$ from (4.27)–(4.28) into (4.26), and simplifying terms, yields:

$$\sum_{i \in \mathcal{B}} \hat{\Pi}_i^\lambda + \hat{V}_i^\lambda(\alpha_i) = \mathbb{E}_{\mathbf{d}} \left[\max_{\mathbf{u} \in \mathcal{U}(\mathbf{d})} \left\{ \sum_{i \in \mathcal{B}} (r_i - \lambda) u_i + \hat{V}_i^\lambda(u_i + (1 - d_i) \alpha_i) \right\} \right].$$

Interchanging summation with maximization and expectation in the r.h.s. yields (4.29). $\square$

Proof of Proposition 4.6. Carrying out the expectation in (4.29) yields:

$$\hat{\Pi}_i^\lambda + \hat{V}_i^\lambda(\alpha_i) = \max_{u_i \in \{0, 1\}} \left\{ q_i(\alpha_i) [(r_i - \lambda) u_i + \hat{V}_i^\lambda(u_i)] \right\} + \bar{q}_i(\alpha_i) \hat{V}_i^\lambda(\alpha_i). \quad (\text{C.10})$$

Let $h_i(u_i)$ denote the term inside the maximization as a function u_i . Then $h_i(1) - h_i(0) = q_i(\alpha_i) (r_i - \lambda) + \hat{V}_i^\lambda(1) - \hat{V}_i^\lambda(0)$. Substituting $\hat{V}_i^\lambda(1)$ and $\hat{V}_i^\lambda(0)$ from (4.31) yields the difference: $h_i(1) - h_i(0) = q_i(\alpha_i) [r_i - \lambda] + (r_i - \lambda)^+ \gamma_i / (1 - \gamma_i)$. If $r_i > \lambda$, this difference is positive which implies that $u_i = 1$ is optimal. If $r_i \leq \lambda$, the difference is negative which implies that $u_i = 0$ is optimal. In both cases, $\hat{\Pi}_i^\lambda$ given by (4.32) and $\hat{V}_i^\lambda(1)$ and $\hat{V}_i^\lambda(0)$ given by (4.31) verify (C.10) and therefore (4.29). $\square$

Proof of Proposition 4.7. From (4.24), (4.28) and (4.32) we have: $\hat{\Pi}^{y, \lambda} = (\lambda - c)y(\alpha) + \sum_{i \in \mathcal{B}} (r_i - \lambda)^+ q_i(1)$. Clearly, $\hat{\Pi}^{y, \lambda}$ is continuous, piecewise linear, and convex

in λ , with $(\partial \hat{\Pi}^{y,\lambda} / \partial \lambda)^+ = y(\alpha) - \sum_{k=1}^{i-1} q(k)(1)$ for $r_{(i)} \leq \lambda < r_{(i-1)}$, $i = 1, \dots, n+1$, where $r_0 \rightarrow \infty$, by convention. The value of λ that minimizes (4.33), denoted by λ^* , is the smallest λ for which $(\partial \hat{\Pi}^{y,\lambda} / \partial \lambda)^+ \geq 0$, which can be expressed as (4.34). $\square$

Proof of Corollary 4.3. If $y(\alpha) = n-1$, $\forall \alpha$, then there are two cases to consider.

Case 1: $\sum_{k=1}^n q(k)(1) > n-1$. From (4.34) and (4.38), $i^* = n, \lambda^* = r_{(n)}, l_{(i)}^* = r_{(i)} + (r_{(i)} - r_{(n)}) \gamma_{(i)} / (1 - \gamma_{(i)}) > r_{(i)}$, $i = 1, \dots, n-1$, and $l_{(n)}^* = r_{(n)}$. **Case 2:** $\sum_{k=1}^n q(k)(1) \leq n-1$. From (4.34) and (4.38), $i^* = n+1, \lambda^* = r_{(n+1)} = 0$, and $l_{(i)}^* = r_{(i)} + r_{(i)} \gamma_{(i)} / (1 - \gamma_{(i)}) = r_{(i)} / (1 - \gamma_{(i)}) = s_{(i)}$, $i \in \mathcal{B}$. For $\mathcal{B} = \{1, 2\}$, case 1 yields: $l_{(1)}^* > r_{(1)} > r_{(2)} = l_{(2)}^*$; therefore, the buyer selection policy is revenue-greedy, which is what we would get if we set $l_{(1)}^* = r_{(1)}$. $\square$

Proof of Proposition 4.8. If $d_i = 0$, the solution of DP (4.42) is $u_i = 0$, since $u_i \leq d_i$. If $d_i = 1$, there are two cases to consider. **Case 1:** $D_{-i} < y(\alpha)$. In this case, the term inside the maximization in (4.42) is $\tilde{V}_i^{\theta_i}(0)$, for $u_i = 0$, and $r_i + \tilde{V}_i^{\theta_i}(1)$, for $u_i = 1$. Clearly, $u_i = 1$ is optimal, because $r_i + \tilde{V}_i^{\theta_i}(1) - \tilde{V}_i^{\theta_i}(0) = \theta_i$ from (4.41) and $\theta_i > 0$, since θ_i is defined as the subsidy that must be given to the firm to make it indifferent between selecting vs. not selecting buyer i . **Case 2:** $D_{-i} \geq y(\alpha)$. In this case, the term inside the maximization of (4.42) is $\tilde{V}_i^{\theta_i}(0)$ for, $u_i = 0$, and $r_i - \theta_i + \tilde{V}_i^{\theta_i}(1)$, for $u_i = 1$, which equals $\tilde{V}_i^{\theta_i}(0)$ from (4.41); therefore, both $u_i = 0$ and $u_i = 1$ are optimal. This is expected because θ_i is defined as the subsidy that must be given to the firm to make it indifferent between selecting vs. not selecting buyer i . Therefore, in both cases, $u_i = 1$ is optimal. From the above analysis, DP (4.42) for $u_i = 1$ can be written as follows, after carrying out the expectation and rearranging terms: $\tilde{\Pi}_i^{\theta_i} + q_i(\alpha_i) \tilde{V}_i^{\theta_i}(\alpha_i) = q_i(\alpha_i) [r_i + \tilde{V}_i^{\theta_i}(1)] F_{-i}(y(\alpha) - 1) + \tilde{V}_i^{\theta_i}(0) F_{-i}(y(\alpha) - 1)$, $\alpha_i = 0, 1$, $i \in \mathcal{B}$. This set of equations has multiple solutions. For this reason, we set $\tilde{V}_i^{\theta_i}(0) = 0$ and we write the above expression as follows: $\tilde{\Pi}_i^{\theta_i} + q_i(\alpha_i) \tilde{V}_i^{\theta_i}(\alpha_i) = q_i(\alpha_i) F_{-i}(y(\alpha) - 1) [r_i + \tilde{V}_i^{\theta_i}(1)]$. For $\alpha_i = 0, 1$, this expression can be written as $\tilde{\Pi}_i^{\theta_i} = q_i(0) F_{-i}(y(\alpha) - 1) [r_i + \tilde{V}_i^{\theta_i}(1)]$ and $\tilde{\Pi}_i^{\theta_i} = q_i(1) F_{-i}(y(\alpha) - 1) r_i - q_i(1) \tilde{V}_i^{\theta_i}(1) [1 - F_{-i}(y(\alpha) - 1)]$, respectively. The solution of these equations is given by (4.43) and (4.44). Substituting the solution into (4.41) yields (4.45). $\square$

Proof of Corollary 4.4. If $y(\alpha) = n-1$, $\forall \alpha$, state $\alpha_1 = (\alpha_1, \dots, \alpha_n : \alpha_k = 1, k \in \mathcal{B})$

can be reached under any buyer selection policy. When α_1 is reached, $F_{-i}^1(n-2) = 1 - \prod_{k \in \mathcal{B} \setminus \{i\}} q_{(k)}(1)$. From (4.17) and (4.45), this implies that $\theta_i(\alpha_1) = z_i$, $i \in \mathcal{B}$. Under the active-constraint index policy, if all buyers are active in state α_1 , the buyer with the smallest index, say j , is left out and becomes dissatisfied, and all other buyers remain satisfied, i.e., the state becomes $\alpha_j = (\alpha_1, \dots, \alpha_n : \alpha_j = 0, \alpha_k = 1, k \in \mathcal{B} \setminus \{j\})$. In state α_j , $F_{-j}^1(n-2)$ remains unchanged. From (4.45), this implies that the index of buyer j remains unchanged, i.e., $\theta_j(\alpha_j) = \theta_j(\alpha_1) = z_j$. For $i \neq j$, $F_{-i}(n-2) = P(d_j(0) + \sum_{k \in \mathcal{B} \setminus \{i,j\}} d_k(1) \leq n-2)$ which, from (4.45) and the fact that $d_j(1) \geq_{st} d_j(0)$, implies that the index of buyer i becomes larger, i.e., $\theta_i(\alpha_j) \geq \theta_i(\alpha_1) = z_i$. Therefore, in state α_j , the buyer with the smallest index is still j . This means that if all buyers are active in state α_j , buyer j is left out and remains dissatisfied, while all other buyers remain satisfied, i.e., the satisfaction state vector remains α_j . On the other hand, if buyer j is active and at least one of the other buyers is inactive, then all active buyers, including j , are satisfied, and the satisfaction state vector becomes α_1 . This behavior is identical to that under the optimal buyer selection policy given by Proposition 4.4. $\square$

Bibliography

- Abdel-Aal, M.A., Selim, S.Z., 2019. Robust optimization for selective newsvendor problem with uncertain demand. *Comput Ind Eng* 135, 838–854. doi:10.1016/j.cie.2019.06.047.
- Adelman, D., Mersereau, A.J., 2013. Dynamic capacity allocation to customers who remember past service. *Manage Sci* 59, 592–612.
- Alicke, K., Barriball, E., Foster, T., Mauhourat, J., Trautwein, V., 2022. Taking the pulse of shifting supply chains. [online] McKinsey & Company URL: <https://www.mckinsey.com/capabilities/operations/our-insights/taking-the-pulse-of-shifting-supply-chains>.
- Anderson, E.T., Fitzimons, G.J., Simester, D.I., 2006. Measuring and mitigating the cost of stockouts. *Manage Sci* 52, 1751–1763.
- Andrews, I., Barron, D., 2016. The allocation of future business: Dynamic relational contracts with multiple agents. *Amer Econ Rev* 106, 2742–2759.
- Armony, M., Plambeck, E.L., 2005. The impact of duplicate orders on demand estimation and capacity investment. *Manage Sci* 51, 1505–1512.
- Babich, V., Burnetas, A.N., Ritchken, P.H., 2007. Competition and diversification effects in supply chains with supplier default risk. *M&SOM-Manuf Serv Op* 9, 123–146. doi:10.1287/msom.1060.0122.
- Bagnoli, M., Bergstrom, T., 2005. Log-concave probability and its applications. *Econ Theor* 26, 445–469.

- Bassok, Y., Anupindi, R., 1997. Analysis of supply contracts with total minimum commitment. *IIE Trans* 29, 373–381.
- Bavafa, H., Leys, C.M., Örmeci, L., Savin, S., 2019. Managing portfolio of elective surgical procedures: A multidimensional inverse newsvendor problem. *Oper Res* 67, 1543–1563. doi:10.1287/opre.2019.1848.
- Benjaafar, S., Elahi, E., Donohue, K.L., 2007. Outsourcing via service competition. *Manage Sci* 53, 241–259.
- Bertsimas, D., Mišić, V.V., 2016. Decomposable Markov decision processes: A fluid optimization approach. *Oper Res* 64, 1537–1555. doi:10.1287/opre.2016.1531.
- Beyer, D., Chang, F., Sethi, S.P., Taksar, M., 2010. *Markovian Demand Inventory Models*. New York, NY: Springer.
- Blazenko, G.W., Vandezande, K., 2003. Corporate holding of finished goods inventories. *J Econ Bus* 55, 255–266. doi:10.1016/S0148-6195(03)00023-7.
- Boon, M.A., van der Mei, R.D., Winands, E.M., 2011. Applications of polling systems. *Surv Oper Res and Manage Sci* 16, 67–82. doi:10.1016/j.sorms.2011.01.001.
- Bosch, 2020. Purchasing and logistics: Partner and training: Cooperation. URL: http://purchasing.bosch.com/en/de/partner_training/cooperation/cooperation.html.
- Brown, D.B., Smith, J.E., 2020. Index policies and performance bounds for dynamic selection problems. *Manage Sci* 66, 3029–3050. doi:10.1287/mnsc.2019.3342.
- Cachon, G., Lariviere, M.A., 1999. Capacity allocation using past sales: When to turn-and-earn. *Manage Sci* 45, 621–769. doi:10.1287/mnsc.45.5.685.
- Cachon, G., Zhang, F., 2006. Procuring fast delivery: Sole sourcing with information asymmetry. *Manage Sci* 52, 881–896.
- Cachon, G., Zhang, F., 2007. Obtaining fast service in a queuing system via performance-based allocation of demand. *Manage Sci* 53, 408–420.

- Campo, K., Gijsbrechts, E., Nisol, P., 2003. The impact of retailer stockouts on whether, how much and what to buy. *Int J Res Mark* 20, 273–286.
- Carr, S., Lovejoy, W., 2000. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Manage Sci* 46, 912–927. doi:10.1287/mnsc.46.7.912.12036.
- Çetinkaya, S., Parlar, M., 1998. Optimal myopic policy for a stochastic inventory problem with fixed and proportional backorder costs. *Eur J Oper Res* 110, 20–41.
- Chahar, K., Taaffe, K., 2009. Risk averse demand selection with all-or-nothing orders. *Omega* 37, 996–1006. doi:10.1016/j.omega.2008.11.004.
- Chen, J., Rubin, H., 1986. Bounds for the difference between median and mean of gamma and poisson distributions. *Stat Prob Lett* 4, 281–283. doi:10.1016/0167-7152(86)90044-1.
- Chinchuluun, A., Karakitsiou, A., Mavrommati, A., 2008. Game theory models and their applications in inventory management and supply chain, in: Pareto optimality, game theory and equilibria. Springer, pp. 833–865. doi:10.1007/978-0-387-77247-9_29.
- Choi, S., Ketzenberg, M., 2018. An inverse newsvendor model to set the optimal number of customers in a capacitated environment. *Int J Prod Econ* 196, 189–197. doi:10.1016/j.ijpe.2017.11.017.
- Christopher, M., Lee, H., 2004. Mitigating supply chain risk through improved confidence. *Int J Phys Distr Log* 34, 388–396.
- Deng, T., Shen, Z.J.M., Shanthikumar, J.G., 2014. Statistical learning of service-dependent demand in a multiperiod newsvendor setting. *Oper Res* 62, 1064–1076.
- Dion, P.A., Banting, P.M., 1995. Buyer reactions to product stockouts in business to business markets. *Indus Mark Manage* 24, 341–350. doi:10.1016/0019-8501(95)00018-6.

- Durango-Cohen, E.J., Li, C.H., 2017. Modeling supplier capacity allocation decisions. *Int J Prod Econ* 184, 256–272. doi:10.1016/j.ijpe.2016.10.013.
- Elahi, E., 2013. Outsourcing through competition: What is the best competition parameter? *Int J Prod Econ* 144, 370–382. doi:10.1016/j.ijpe.2013.02.023.
- Fergani, Y., 1976. A Market-Oriented Stochastic Inventory Model. Phd dissertation. Department of Industrial Engineering, Stanford University.
- Fisher, M., Keil, F.C., 2018. The binary bias: A systematic distortion in the integration of information. *Psychol Sci* 29, 1846–1858. doi:10.1177/0956797618792256.
- Fisher, M., Newman, G.E., Dhar, R., 2018. Seeing stars: How the binary bias distorts the interpretation of customer ratings. *J Consum Res* 45, 471–489. doi:10.1093/jcr/ucy017.
- Fitzsimons, G.J., 2000. Consumer response to stockouts. *J Consum Res* 27, 249–266.
- Fredrickson, B.L., Kahneman, D., 1993. Duration neglect in retrospective evaluations of affective episodes. *J Pers Soc Psychol* 65, 45–55. doi:10.1037/0022-3514.65.1.45.
- Gallego, G., Katircioglu, K., Ramachandran, B., 2007. Inventory management under highly uncertain demand. *Oper Res Let* 35, 281–289. doi:10.1016/j.orl.2006.03.012.
- Gans, N., 2002. Customer loyalty and supplier quality competition. *Manage Sci* 48, 207–221.
- Gaur, V., Park, Y.H., 2007. Asymmetric consumer learning and inventory competition. *Manage Sci* 53, 227–240.
- Gilbert, S.M., Weng, Z.K., 1998. Incentive effects favor nonconsolidating queues in a service system: The principal–agent perspective. *Manage Sci* 44, 1662–1669. doi:10.1287/mnsc.44.12.1662.

- Goodman, P.S., Chokshi, N., 2021. How the world run out of everything. The New York Times URL: <https://www.nytimes.com/2021/06/01/business/coronavirus-global-shortages.html>.
- Ha, A.Y., Li, L., Ng, S.M., 2003. Price and delivery logistics competition in a supply chain. *Manage Sci* 49, 1139–1153. doi:10.1287/mnsc.49.9.1139.16567.
- Hall, J.M., Porteus, E.L., 2000. Customer service competition in capacitated systems. *M&SOM-Manuf Serv Op* 2, 144–165.
- Homburg, C., Droll, M., Totzek, D., 2008. Customer prioritization: does it pay off, and how should it be implemented? *J Market* 72, 110–130. doi:10.1509/jmkg.72.5.110.
- Hong, Y., 2013. On computing the distribution function for the poisson binomial distribution. *Comput Stat Data An* 59, 41–51. doi:10.1016/j.csda.2012.10.006.
- Intel, 2020. Supplier guide. URL: https://supplier.intel.com/static/supplier/supplier_guide.pdf.
- Jackson, B.B., 1985. Build customer relationships that last. *Harvard Bus Rev* 11, 120–128.
- Jin, Y., Ryan, J., 2012. Price and service competition in an outsourced supply chain. *Prod Oper Manage* 21, 331–344.
- Jing, X., Lewis, M., 2011. Stockouts in online retailing. *J Marketing Res* 48, 342–354.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–292.
- Kalai, E., Kamien, M.I., Rubinovitch, M., 1992. Optimal service speeds in a competitive environment. *Manage Sci* 38, 1154–1163. doi:10.1287/mnsc.38.8.1154.
- Klein, R., Kolb, J., 2015. Maximizing customer equity subject to capacity constraints. *Omega* 55, 111–125. doi:10.1016/j.omega.2015.02.008.

- Koos, S.E., Shaikh, N.I., 2019. Dynamics of consumers' dissatisfaction due to stock-outs. *Int J Prod Econ* 208, 461–471.
- Lau, H.S., 1997. Simple formulas for the expected costs in the newsboy problem: An educational note. *Eur J Oper Res* 100, 557–561. doi:10.1016/j.ejor.2014.06.007.
- Li, H., Zhang, H., Fine, C.H., 2013. Dynamic business share allocation in a supply chain with competing suppliers. *Oper Res* 61, 280–297.
- Li, J., Wang, S., Cheng, T.C.E., 2010. Competition and cooperation in a single-retailer two-supplier supply chain with supply disruption. *Int J Prod Econ* 124, 137–150. doi:10.1016/j.ijpe.2009.10.017.
- Li, L., 1992. The role of inventory in delivery-time competition. *Oper Res* 38, 182–197.
- Liberopoulos, G., Deligiannis, M., 2022. Optimal supplier inventory control policies when buyer purchase incidence is driven by past service. *Eur J Oper Res* 300, 917–936. doi:10.1016/j.ejor.2021.09.002.
- Liberopoulos, G., Tsikis, I., 2007. Competing for customer goodwill on product availability. *Proc. 6th Internat. Conf. on Analysis of Manufacturing Systems*, May 11-16, Lunteren, The Netherlands , 44–50.
- Liu, L., Shang, W., Wu, S., 2007. Dynamic competitive newsvendors with service-dependent demands. *M&SOM-Manuf Serv Op* 9, 84–93.
- Liyanage, L.H., Shanthikumar, J.G., 2005. A practical inventory control policy using operational statistics. *Oper Res Lett* 33, 341–348. doi:10.1016/j.orl.2004.08.003.
- Lu, L.X., Lariviere, M.A., 2012. Capacity allocation over a long horizon: The return on turn-and-earn. *M&SOM-Manuf Serv Op* 14, 1–163. doi:10.1287/msom.1110.0346.

- Mahajan, S., van Ryzin, G., 2001. Stocking retail assortments under dynamic consumer substitution. *Oper Res* 49, 334–351. doi:10.1287/opre.49.3.334.11210.
- Marshall, G., 2020. Search and wholesale price discrimination. *RAND J Econ* 51, 346–374. doi:10.1111/1756-2171.12317.
- Minner, S., 2003. Multiple-supplier inventory models in supply chain management: A review. *Int J Prod Econ* 81, 265–279. doi:10.1016/S0925-5273(02)00288-8.
- Mittal, V., Kamakura, W.A., 2001. Satisfaction repurchase intent, and repurchase behavior: Investigating the moderating effect of customer characteristics. *J Market Res* 28, 131–142.
- Nagarajan, M., Sošić, G., 2008. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *Eur J Oper Res* 187, 719–745. doi:10.1016/j.ejor.2006.05.045.
- Nahmias, S., Olsen, T.L., 2015. *Production and Operations Analysis*, 7E. Long Grove, IL: Waveland.
- Olsen, T.L., Parker, R.P., 2008. Inventory management under market size dynamics. *Manage Sci* 54, 1805–1821.
- Oracle, 2020. Peoplesoft enterprise performance management 9.1: Peoplesoft scorecard: Understanding supplier rating system scorecard. URL: https://docs.oracle.com/cd/E41507_01/epm91pbr3/eng/epm/pbsc/concept_UnderstandingSupplierRatingSystemScorecard-9916c5.html.
- Oracle, 2021. Oracle Fusion Cloud Supplier Qualification Management. URL: <https://www.oracle.com/a/ocom/docs/applications/erp/oracle-supplier-qualification-management-cloud-ds.pdf>.
- Rasmussen, R.V., Trick, M.A., 2008. Round robin scheduling—a survey. *Eur J Oper Res* 188, 617–636. doi:10.1016/j.ejor.2007.05.046.

- Riedl, D.F., Kaufmann, L., Zimmermann, C., Perols, J.L., 2013. Reducing uncertainty in supplier selection decisions: Antecedents and outcomes of procedural rationality. *J Oper Manage* 31, 24–36.
- Robinson, L.W., 2016. Appropriate inventory policies when service affects future demands, in: Gavirneni, S. (Ed.), *Cross-Functional Inventory Research*. World Scientific. World Scientific-Now Series in Business 10, pp. 135–167.
- Rossi, R., Prestwich, S., Tarim, A., Hnich, B., 2014. Confidence-based optimisation for the newsvendor problem under binomial, poisson and exponential demand. *Eur J Oper Res* 239, 674–684. doi:10.1016/j.ejor.2014.06.007.
- Samsung, 2020. About samsung: Sustainability: Supply chain. URL: <https://www.samsung.com/us/aboutsamsung/sustainability/supply-chain/>.
- SAP, 2022a. SAP Ariba: Running supplier performance management reviews. URL: <https://help.sap.com/doc/91529a66c9be486783a4498ad9b9e8f8/cloud/en-US/SpmUse.pdf>.
- SAP, 2022b. Sap ariba strategic sourcing solutions: Running supplier performance management reviews. URL: https://help.sap.com/docs/ARIBA_SOURCING/605fe6c015834314b3acae8bbd7f699/a9b703611b63489f97f3c5fb54460fdd.html?locale=en-US.
- SAP, 2022c. Sap erp support package versions: Material requirements planning (pp-mrp): Master data (pp-mrp-bd): Quota arrangements: Allocation quota arrangement. URL: https://help.sap.com/docs/SAP_ERP_SPV/85d3fce10e264972a0155c8b46ecf93b/8097b6535fe6b74ce10000000a174cb4.html?version=6.17.21&locale=en-US.
- Saracoglu, I., Topaloglu, S., Keskindurk, T., 2014. A genetic algorithm approach for multi-product multi-period continuous review inventory models. *Expert Syst Appl* 41, 8189–8202.

- Sen, A., Zhang, A.X., 1999. The newsboy problem with multiple demand classes. *IIE Trans* 31, 431–444. doi:10.1080/07408179908969846.
- Sheffi, Y., 2020. The new (ab) normal: reshaping business and supply chain strategy beyond covid-19. MIT CTL Media.
- Siegel, A.F., Wagner, M.R., 2021. Profit estimation error in the newsvendor model under a parametric demand distribution. *Manage Sci* 67, 4863–4879. doi:10.1287/mnsc.2020.3766.
- Silbermayr, L., 2020. A review of non-cooperative newsvendor games with horizontal inventory interactions. *Omega* 92, 102148. doi:10.1016/j.omega.2019.102148.
- Singer, P., Helic, D., Taraghi, B., Strohmaier, M., 2014. Detecting memory and structure in human navigation patterns using markov chain models of varying order. *PLOS ONE* 9, e114952.
- Svoboda, J., Minner, S., Yao, M., 2021. Typology and literature review on multiple supplier inventory control models. *Europ J Oper Res* 293, 1–23. doi:10.1016/j.ejor.2020.11.023.
- SYSPRO, 2022. SYSPRO Preferred Supplier Solutions Factsheet. URL: <https://www.syspro.com/dl/FS/SYSPRO-Preferred-Supplier-FS.pdf>.
- Taafe, K., Romeijn, E., Tirumalasetty, D., 2008. A selective newsvendor approach to order management. *Nav Res Logist* 55, 769–784. doi:10.1002/nav.20320.
- Tang, C.H., 2006. Perspectives in supply chain risk management. *Int J Prod Econ* 103, 451–488. doi:10.1016/j.ijpe.2005.12.006.
- Tang, C.H., Tomlin, B., 2008. The power of flexibility for mitigating supply chain risks. *Int J Prod Econ* 116, 12–27.
- Topaloglu, H., 2009. Using Lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. *Oper Res* 57, 637–649. doi:10.1287/opre.1080.0597.

- Tversky, A., Kahneman, D., 1991. Loss aversion in riskless choice: A reference-dependent model. *Q J Econ* 106, 1039–1061.
- Ülkü, M.A., Gürler, Ü., 2018. The impact of abusing return policies: A newsvendor model with opportunistic consumers. *Int J Prod Econ* 203, 124–133. doi:10.1016/j.ijpe.2018.05.016.
- Usmani, R.A., 1994. Inversion of a tridiagonal jacobi matrix. *Linear Algebra Appl* 212–213, 413–414.
- Valinsky, J., 2020. One in five Wendy's is out of beef, analyst says. *CNN Business* URL: <https://edition.cnn.com/2020/05/05/business/wendys-beef-shortage/index.html>.
- Wang, Y., Eallace, S.W., Shen, B., Choi, T.M., 2015. Service supply chain management: A review of operational models. *Europ J Oper Res* 247, 685–698.
- Wang, Y.H., 1993. On the number of successes in independent trials. *Stat Sinica* 3, 295–312. URL: <https://www.jstor.org/stable/24304959>.
- Wetzel, H.A., Hammerschmidt, M., Zablah, A.R., 2014. Gratitude versus entitlement: A dual process model of the profitability implications of customer prioritization. *J Market* 78, 1–19. doi:doi.org/10.1509/jm.12.0167.
- Whittle, P., 1988. Restless bandits: Activity allocation in a changing world. *J App Prob* 25, 287–298. doi:10.2307/3214163.