TESIS DE LA UNIVERSIDAD DE ZARAGOZA

2013

129

Eirini Spiliotopoulou

Information Reliability in supply chains: the case of multiple retailers

Departamento

Zaragoza Logistics Center

Director/es

Gurbuz, Mustafa Cagri

Donohue, Karen

ISSN 2254-7606

Prensas de la Universidad
Universidad Zaragoza

Universidad Zaragoza

1542

Tesis Doctoral

INFORMATION RELIABILITY IN SUPPLY CHAINS: THE CASE OF MULTIPLE RETAILERS

Autor

Eirini Spiliotopoulou

Director/es

Gurbuz, Mustafa Cagri
Donohue, Karen

UNIVERSIDAD DE ZARAGOZA

Zaragoza Logistics Center

2013

UNIVERSIDAD DE ZARAGOZA

TESIS DOCTORAL

FIABILIDAD DE LA INFORMACION EN CADENAS DE SUMINISTRO: EL CASO DE LOS DISTRIBUIDORES MINORISTAS

Eirini Spiliotopoulou

Máster de Ingeniería en Logística y Gestión de la Cadena de Suministro,
Programa Internacional de Logística MIT-Zaragoza
Zaragoza Logistics Center (ZLC), Universidad de Zaragoza (España)

Máster en Administración y Dirección de Empresas (MBA),
Universidad de Económicas y Administración de Empresas de Atenas (Grecia)

Licenciada en Investigación Operativa y Marketing,
Universidad de Económicas y Administración de Empresas de Atenas (Grecia)

20 de Junio de 2013

Autor: Dña. Eirini Spiliotopoulou, Doctorando

Director de tesis: Dr. Mustafa Çagri Gürbüz
Profesor de Gestión de la Cadena de Suministro,
Programa Internacional de Logística MIT-Zaragoza
Zaragoza Logistics Center (ZLC)

Director del Zaragoza Logistics Center (ZLC): Dr. David Gonsalvez

Fiabilidad de la información en cadenas de suministro: el caso de los distribuidores minoristas

por

Eirini Spiliotopoulou

en relación con el cumplimiento parcial de los requisitos para la obtención del título de Doctor en Logística y Gestión de las Cadenas de Suministro

Resumen

En esta tesis doctoral abordamos el estudio relativo al intercambio de información sobre la demanda dentro de una cadena de suministro cuando las partes interactúan de una forma estratégica. Los distribuidores minoristas forman una agrupación y delegan la gestión del inventario (los pedidos y la asignación) a un planificador central benévolo (CP, por sus siglas en inglés). Cada uno de los minoristas debe enfrentarse a una demanda incierta y dispone de información privada sobre ella como consecuencia de su proximidad al mercado; nos centramos en determinar si entre los minoristas y el CP se produce un intercambio fiable de información sobre la demanda. En primer lugar estudiamos el impacto de diversos mecanismos de asignación sobre el comportamiento en materia de pedidos de los minoristas, cuando la cantidad de inventario total en el almacén central es fija. Los minoristas efectúan los pedidos después conocer de manera privada su demanda. Demostramos analíticamente que los minoristas comunicarán sus necesidades reales, es decir, sus demandas realizadas, de acuerdo a una norma de asignación uniforme pero no de acuerdo a otras normas comunes como, por ejemplo, la noma proporcional o lineal; posteriormente, estudiamos una configuración donde la cantidad de inventario agrupado no es fija, sino más bien una variable de decisión, determinada por el CP después de haber solicitado información de demanda prevista de los minoristas. Las asignación del inventario total, en este caso también, se efectúa después de conocerse las realizaciones de demanda final, pero las demandas finales son de conocimiento común. Entonces, los minoristas pueden influir su asignación solo a través de la cantidad de inventario total. Mediante modelos teóricos asociados a tácticas podemos ver que el reconocimiento de la verdad y la confianza no se encuentran en una situación de equilibrio. A continuación, en un entorno de laboratorio controlado que simula la configuración de la cadena de suministro objeto de consideración, estudiamos el impacto de a) la competencia por el inventario común y b) la incertidumbre del mercado sobre la distorsión de la información, la confianza y la eficacia de la cadena de suministro. Nuestros resultados sugieren que existe una confianza continua cuando los incentivos pecuniarios están alineados y cuando no lo están, lo que viene a desmentir los casos teóricos extremos de minoristas completamente dignos de confianza o que no son fiables en absoluto; incluso aunque la información no sea totalmente fiable, el valor de la comunicación es importante. En última instancia, estudiamos el impacto de la propiedad del inventario sobre las motivaciones de las partes implicadas de cara a compartir de manera honrada sus previsiones de demanda; los inventarios específicos tampoco inducen a decir la verdad. Comparamos los inventarios resultantes y los beneficios de acuerdo con la toma de decisiones a nivel local con información más precisa con respecto a la toma de decisiones centralizada, mediante la cual se logra la coordinación de los pedidos.

Director de tesis: doctor Mustafa Cagri Gurbuz
Cargo: Profesor de Gestión de la Cadena de Suministro, Zaragoza Logistics Center

Agradecimientos

En primer lugar, me gustaría dar las gracias más sinceras a mis directores de tesis, los doctores Mustafa Çagri Gürbüz y la doctora Karen Donohue, por sus aportaciones y sus consejos durante todo el proceso; también me gustaría expresar mi agradecimiento al doctor Santiago Kraiselburd por su apoyo en este proyecto desde el primer momento.

Por encima de todo, me gustaría dar las gracias a mi familia y mis amigos, repartidos por todo el mundo, que me han dado todo su apoyo durante estos años de mi vida tan complicados, a la par que gratificantes, como estudiante de doctorado.

Contenidos

1 Introducción **17**

1.1 Motivación 17

1.2 Configuración común 18

1.3 Preguntas de investigación 21

2 Análisis de la literatura especializada **25**

2.1 Decisiones sobre la agrupación de inventarios 25

2.2 Elección de capacidad y asignación 30

2.3 Rendimiento aleatorio 32

2.4 Toma de decisiones central frente a local 34

2.5 Intercambio de información sobre previsiones de la demanda 36

3 Impacto de los mecanismos de asignación del inventario en el intercambio de información sobre la demanda **39**

3.1 Introducción 39

3.2 Configuración y calendario de los eventos 41

3.3 Propiedades de los Pareto-óptimos mecanismos de asignación 44

3.4 Asignación con información simétrica 46

3.5 Asignación con información asimétrica 50

3.6 Comentarios a modo de conclusión 53

4 Intercambio de información sobre previsiones y decisión sobre la cantidad de pedidos: impacto de la competencia sobre los inventarios y la incertidumbre del mercado 55

4.1 Introducción 55

4.2 El modelo analtico 58

4.2.1 Configuración y calendario de los eventos 58

4.2.2 La tctica de intercambio de informacin para $n = 1$ 60

4.2.3 La tctica de intercambio de informacin y asignacin para $n > 1$ 61

4.3 Estudio numérico 69

4.3.1 Conocimiento común 70

4.3.2 Información asimétrica 71

4.4 Hipótesis 74

4.5 Experimentos 76

4.5.1 Diseño y procedimientos experimentales 76

4.5.2 Resultados experimentales 81

4.6 Comentarios a modo de conclusi'on 90

5 Intercambio de informaci'on sobre previsiones y decisi'on sobre la cantidad de pedidos: impacto de la propiedad sobre el inventario 93

5.1 Introducci'on 93

5.2 Configuraci'on y calendario de los eventos para LMI y CPMI 95

5.3 Inventario gestionado a nivel local (LMI) 97

5.4 Inventario gestionado por un planificador central (CPMI) 104

5.4.1 El hecho de decir la verdad, ?' representa un equilibrio? . . . 108

5.5 Análisis numérico 111

5.6 Coordinación de los precios de transferencia 119

5.7 Comentarios a modo de conclusión 122

A Notas 131

B Pruebas **135**

B.1 Pruebas del capítulo 3 135

B.2 Pruebas del capítulo 4 139

B.3 Pruebas del capítulo 5 147

C Estudio numérico adicional **153**

D Instrucciónes del experimento **155**

D.1 Hoja informativa para los directores regionales: caso R_2D_U 155

D.2 Hoja informativa para los planificadores centrales: caso R_2D_U 159

D.3 Hojas resumen 163

E Imágenes del software del experimento **165**

E.1 Imágenes de la pantalla del ordenador de un minorista: caso R_2D_U . 165

E.2 Imágenes de la pantalla del ordenador de un planificado central: caso R_2D_U 168

Lista de figuras

3-1 El calendario de eventos en la táctica de asignación del inventario . . 44

4-1 El calendario de eventos en la táctica de intercambio de información con inventario común 59

4-2 $(Q^1 - \mathbb{E}[Q^{CP}])$ como función de la señal recibida para alta variabilidad de seal, pequeña incertidumbre de mercado y dimensiones asimétricas de dicho mercado 73

4-3 Registros de los minoristas frente a las señales reales de la demanda . 83

4-4 Selección de inventario de los CP frente a las previsiones registradas $(Q(\hat{\theta}))$ 84

5-1 El calendario de los eventos para LMI y CPMI 97

5-2 Beneficio total esperado como función de q_1 y q_2 107

5-3 $q_i^{LMI}(\theta_H) - q_i^{LMI}(\theta_L)$ como función de la incertidumbre del mercado . 116

Lista de tablas

1.1 1.1. Resumen de los captulos principales 24

3.1 Comparación de diversos mecanismos de asignación 52

4.1 Resumen de hallazgos analticos del capítulo 4 69

4.2 Inventario óptimo para el minorista 1 como función de θ_1 ? θ_2 cuando cr=0,5 70

4.3 $Q_f^1 - Q_f^{CP}$, como función de θ_1 y θ_2, cuando cr=0,75 and cr=0,25 . . 71

4.4 Comparación del inventario óptimo para el minorista 1 y el inventario esperado en el sistema si se dice la verdad, como función de θ_1 73

4.5 Diseño experimental 78

4.6 Resumen de estadsticas en el valor $|\hat{\theta} - \theta|$ 82

4.7 Resumen de estadsticas en el valor Q 82

4.8 Definici'on de variables en ecuaciones (4.6)—(4.8) 87

4.9 Impacto de la competencia y de la incertidumbre del mercado sobre el hecho de decir la verdad, la confianza y la eficacia 88

5.1 Elecciones de inventario optimas de acuerdo con LMI: $(q_i^{LMI}(\theta_i) = (q_{iL}, q_{iH}))$ 113

5.2 Elecciones de inventario optimas de acuerdo con LMI y conocimiento común 115

5.3 Comparación del inventario total y de los beneficios esperados de acuerdo con el LMI, el CPMI y el CPMI con toda la información 117

5.4 Comparación del inventario local de acuerdo con el LMI, el CPMI y el CPMI con toda la información 117

A.1 Resumen de notas 133

C.1 Comparación del inventario óptimo para el minorista 1 y el inventario esperado en el sistema si se dice la verdad, como función de θ_1, cuando cr=0,75 y cr=0,25 153

C.2 El inventario que maximiza el beneficio medio del minorista, para diversos valores de $theta_1$, frente al inventario óptimo para el sistema, cuando la demanda es uniforme y discreta 154

Conclusiones y debate

En esta tesis hemos estudiado los aspectos relacionados con el intercambio de información sobre la demanda – previsiones y demanda realizada – dentro de una coalición asociada con la puesta en común de inventarios. Incluso a pesar de que el valor de la información es un tema bien estudiado en la bibliografa especializada relativa a la gestión de operaciones y de cadenas de suministro, normalmente se parte de la base de que el intercambio de información, cuando se produce, se realiza de una manera creíble. Por otro lado, muchos fallos bien documentados en las actividades comerciales se deben a la desinformación en materia de datos privados, como por ejemplo la exageración de pedidos ante la expectativa de escasez del inventario o pedidos imprecisos demasiado optimistas que nunca se materializan. Centramos nuestro trabajo en determinar si el intercambio fiable de información sobre la demanda se produce entre minoristas y un planificador central benévolo (CP) que coordina la asignación de pedidos y de inventario dentro de una coalición asociada con la puesta en común de inventarios. Los minoristas no compiten por la demanda, pero pueden competir por el inventario; cada uno de ellos cuenta con información privada sobre la demanda en su región como consecuencia de la proximidad al mercado, la cual puede ser transmitida de forma fidedigna, o no, al planificador central.

En primer lugar, estudiamos analíticamente el impacto de diversos mecanismos de asignación en lo que respecta al comportamiento de los minoristas al efectuar pedidos, una vez que se ha establecido la cantidad total de inventario en el almacén central. Para hacer esto, en primer lugar mostramos que, cuando la asignación se basa en la demanda realizada (es decir, la demanda realizada pasa a ser conocida para todas las partes interesadas), todas las normas de asignación consideradas, es decir, proporcional, lineal y uniforme, resultan eficaces (excluyen pérdidas) y con un resultado Pareto óptimo. Sin embargo, cuando las demandas realizadas en cada región permanecen en el ámbito privado con respecto a los minoristas, la norma de asignación empleada desempeña un importante papel; solo a través de la norma de asignación uniforme los minoristas podrán informar de sus necesidades reales mediante la presentación de un pedido definitivo equivalente a su demanda realizada. Este resultado es análogo al caso típico de racionamiento de la capacidad, incluso si en la configuración que esté considerándose (a) la asignación se determina después de resolver la incertidumbre de la demanda y (b) cada minorista puede recibir por encima o por debajo de su pedido final. La diferencia fundamental es que la asignación uniforme basada en los pedidos definitivos, en nuestro caso, induce a decir la verdad y resulta óptima desde el punto de vista de Pareto. En el caso del racionamiento de la capacidad, la asignación uniforme no resulta óptima desde el mencionado punto de vista de Pareto, puesto que no es receptiva desde la perspectiva individual, lo cual constituye una condición necesaria para que resulte óptima desde dicho punto de vista. Además, proponemos una norma modificada de asignación uniforme que no solo sea óptima desde el punto de vista de Pareto e induzca a decir la verdad, sino que también garantice a cada minorista un beneficio superior al que habríamos obtenido en un sistema puramente descentralizado de acuerdo con cualquier realización de la demanda.

A continuaciòn, procedemos a efectuar un estudio analítico y experimental del intercambio de la previsión de la demanda entre minoristas y el CP que solicita esta información para establecer el inventario total. Las demandas realizadas pasan a ser de dominio común cuando se produce la asignación, y nos centramos en la táctica de registrar señales de la demanda (previsión) para influir sobre el inventario contenido en el sistema. Valoramos la táctica en los casos en los que se utiliza un mecanismo de asignación proporcional (para la demanda realizada a nivel local): un mecanismo ampliamente utilizado en la práctica y con multitud de propiedades atractivas. Mediante el uso de modelos basados en la teoría de juegos determinamos que, cuando existe una incertidumbre no resuelta en materia de demanda antes de la comunicación, el hecho de decir la verdad y la confianza no forman un equilibrio Bayesiano Perfecto; además, en un sistema de inventario automatizado que toma como datos las previsiones registradas por los minoristas y solicita el nivel de inventario óptimo para toda la coalición, no existe un equilibrio Bayesiano de Nash puro entre los diversos minoristas. Posteriormente, en un entorno de laboratorio controlado que simula la configuración de la cadena de suministro objeto de consideración, estudiamos el impacto de a) la competencia por el inventario común y b) la incertidumbre del mercado sobre la distorsión de la información, la confianza y la eficacia de la cadena de suministro. Nuestros resultados sugieren que existe una confianza continua cuando los incentivos pecuniarios estàn alineados y cuando no lo están, lo que viene a desmentir los casos teóricos extremos de minoristas completamente dignos de confianza o que no son fiables en absoluto. Además, determinamos que tanto la competencia por el inventario común como la incertidumbre en materia de previsiones dañan notablemente tanto el hecho de decir la verdad como la cooperación entre las diversas partes de la cadena de suministro. Ese es el motivo por el que la acumulaciòn de inventario con arreglo a una asimetría de la información puede tener resultados negativos, a pesar de la agregación de riesgos de la demanda. Incluso aunque la información no fuera íntegramente fiable, el valor de la comunicación era significativo en todos nuestros experimentos.

En último lugar, estudiamos el impacto de la propiedad del inventario sobre los alicientes de las partes interesadas para compartir sus previsiones de forma fidedigna. Cuando valoramos dos ubicaciones independientes que deciden a nivel local sobre su nivel de inventarios y toman en consideración la posibilidad de transferir la propiedad del inventario después de realizar las demandas, existe un equilibrio Bayesiano de Nash con características únicas. La elección óptima de un inventario en un emplazamiento se incrementa en su señal de demanda recibida; cuando el CP adopta la decisión de efectuar el pedido, lo que importa es únicamente la cantidad total de inventario. A menos que concurran algunas condiciones especiales de conservación, éste no podrá ser separado entre los dos emplazamientos antes de que se realice la demanda, de forma que los incentivos locales y del sistema queden alineados. Procedemos a comparar en términos numéricos los inventarios y los beneficios resultantes en condiciones de toma de decisiones a nivel local con información más precisa frente a la toma de decisiones centrales en las que se consigue la coordinación de pedidos. Vemos que, cuando el fractil fundamental es alto, la toma de decisiones a nivel central desemboca en inventarios más elevados, mientras que cuando el fractil fundamental es bajo se cumple la circunstancia contraria. Las comparaciones direccionales de beneficios esperados dependen del valor de la información local que se pierde cuando pasamos a la toma de decisiones a nivel central frente al valor adicional de la coordinación del inventario (en el equilibrio *babbling* o equilibrio no informativo).

Este trabajo tiene varias limitaciones, como consecuencia de la complejidad analtica del problema. Estudiamos por separado la estrategia de asignación de inventarios cuando el planificador central no conoce las demandas finales y aquella relacionada con el intercambio de información sobre la previsión de la demanda para influir sobre el nivel de inventario de la coalición. El hecho de saber cuáles serían las interacciones cuando ambos temas se valoraran de manera conjunta sigue siendo una pregunta abierta. Por ejemplo, si la asignación final estuviera vinculada a la previsión registrada, ¿cómo cambiarían las dinàmicas del intercambio de información sobre previsiones? ¿Cuál sería el impacto sobre la actitud de los minoristas a la hora de efectuar pedidos? ¿Sería eficaz la asignación final?

Hay muchas ampliaciones interesantes de esta tesis doctoral; para empezar, podramos analizar con mayor detalle el impacto de los factores del comportamiento en la estrategia de compartir información relativa a las previsiones de la demanda. Cuando el inventario es común, es interesante investigar cómo el tamaño de la coalición y el de la demanda media los minoristas afecta relaciones de confianza. Además, quisiéramos estudiar si el nivel de confiabilidad de los minoristas cambia cuando hay una garantía de cómo se utilizan sus previsiones de la demanda para ajustar el nivel de inventario común. Consideramos esto una pregunta de investigación interesante con implicaciones gerenciales potencialmente muy relevantes. Como investigación futura, estamos planeando ejecutar experimentos adicionales cuando el nùmero de minoristas aumenta a 3 y 4, el CP está automatizado y los minoristas no son idénticos. Sería interesante estudiar, desde un punto de vista experimental, cómo la propiedad sobre el inventario tiene un impacto sobre la estrategia de registro de las previsiones por parte de los minoristas. Incluso aunque los inventarios específicos no alinearan los alicientes individuales y los pecuniarios del sistema, ¿la incertidumbre reducida en relación con la asignaciòn final incrementaría la confianza de los minoristas y mejorara la cooperación?

Un segundo tema de interés es estudiar, mediante experimentos asociados al comportamiento, el impacto de los mecanismos de asignaciòn en las actitudes de formulación de pedidos de los minoristas, para así arrojar luz sobre los componentes que, potencialmente, pudieran "faltar" en esta interacción. Los conceptos de equilibrio, que parten de la base de que las partes interesadas son perfectamente racionales, ¿exageran notablemente la tendencia de los minoristas a pedir más / menos de que lo que necesitan? ¿En virtud de qué mecanismos de asignación es más pronunciada la distorsión en los pedidos?

Otra vía de investigación futura es estudiar el intercambio de informaciones sobre la demanda en una coalición para la puesta en común de inventarios y centrarse en las implicaciones conductuales a la hora de presentar pedidos imprecisos (no vinculantes) frente al intercambio de predicciones sobre la demanda. Una vez más, partimos de la base de que todos los minoristas tienen una mejor información sobre la demanda gracias a su proximidad al mercado, y que la comparten (quizá de manera falsa) con el CP, ya sea en forma de intercambio de previsiones (enviando su señal de demanda) o de una orden no vinculante antes de que la demanda se realizara. ¿Cómo se comparan los niveles de inventario y los beneficios con arreglo a estas dos formas diferentes de intercambio de información?

Una configuración de cadena de suministro distinta donde el intercambio de las previsiones de la demanda desempeña un papel fundamental es la de proveedor – fabricantes / minoristas. Cuando el proveedor es una unidad de negocio independiente con su propio margen de beneficio, ¿cómo cambia la dinàmica de la comunicación? El hecho de saber cuál sería el impacto de confianza y confiabilidad sobre el intercambio de información estratégica, el nivel inventario (o capacidad) y / o la asignación, en una configuración de proveedor único – múltiples minoristas, sigue siendo una pregunta abierta.

Otra ampliación interesante para trabajos futuros es estudiar, tanto a nivel analítico como experimental, las dinámicas de comunicación y de intercambio de información (demanda final o previsiones) en problemas multiperíodo. Cuando se toman en consideración las interacciones repetidas, ¿en virtud de qué condiciones pueden formar un equilibrio sostenible la confianza y el hecho de decir la verdad? Los contratos de relación (p. ej., basados en estrategias de desencadenamiento), ¿inducen a la colaboracin? En estrategias multiperodo, muchos aspectos adicionales pueden desempeñar tambièn un papel importante, como por ejemplo las opiniones y el aprendizaje, las consideraciones relativas a la reputación, la posibilidad de sancionar un comportamiento engañoso o la confianza en las relaciones a largo plazo.

Information Reliability in Supply Chains: the Case of Multiple Retailers

A thesis presented by

Eirini Spiliotopoulou

In partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Logistics and Supply Chain Management

in the
MIT-Zaragoza International Logistics Program

at the
Zaragoza Logistics Center, a research institute associated with the
University of Zaragoza

July 2013

Information Reliability in Supply Chains: the Case of Multiple Retailers

by

Eirini Spiliotopoulou

in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Logistics and Supply Chain Management

Abstract

In this thesis, we study the issue of demand information sharing within a supply chain when parties strategically interact. Multiple retailers form a pooling coalition and delegate inventory management – ordering and allocation – to a benevolent central planner (CP). Each retailer faces uncertain demand and has private information about it due to his proximity to the market. We focus on whether reliable demand information sharing occurs between retailers and the CP. First, we study the impact of various allocation mechanisms on the ordering behavior of retailers, when the quantity of pooled inventory is fixed. Retailers place their orders after demand is privately revealed to them. We determine analytically that retailers will communicate their true needs (i.e., their realized demand) under a uniform allocation rule, but not under other common rules such as proportional or linear. Next we study a setting where the pooled inventory quantity is no longer fixed, but rather a decision variable, set by the central planner after receiving initial demand forecast information from the retailers. The inventory allocation is still determined after final demand realizations are known, but now this final demand information is common knowledge. So, the retailers can only influence their allocation through the quantity of pooled inventory available. Using game theoretical models, we find that this setting has no trust-telling equilibrium. We then study this setting in a controlled laboratory experiment to test the impact of (a) the number of retailers, and (b) level of demand uncertainty on information distortion, trust and supply chain efficiency. Our results suggest that a continuum of trust exists both when pecuniary incentives are aligned or misaligned, refuting the extreme theoretical cases of fully trustworthy or fully non-trustworthy retailers. Even when information is not fully reliable, the value of communication is significant. Last, we study the impact of inventory ownership on the incentives of the players to truthfully share their demand forecasts and find that dedicated inventories do not induce truth-telling either. We compare resulting inventories and profits under local decision making with more accurate information versus central decision making where coordination of orders is achieved.

Thesis Supervisor: Dr. Mustafa Cagri Gurbuz
Title: Professor of Supply Chain Management, Zaragoza Logistics Center

Thesis Supervisor: Dr. Karen Donohue
Title: Associate Professor of Supply Chain and Operations,
Carlson School of Management, University of Minnesota

Acknowledgments

First, I would like to sincerely thank my advisors Dr. Mustafa Cagri Gurbuz and Dr. Karen Donohue for their input and guidance throughout the process. I would also like to thank Dr. Santiago Kraiselburd for his support for this project in its early stage.

Above all, I deeply thank my family and my friends, dispersed around the globe, who kept supporting me during these difficult but rewarding years of my life as a PhD student.

Author . . . .

Eirini Spiliotopoulou, PhD Candidate
MIT-Zaragoza International Logistics Program

Thesis Supervisor . . . .

Dr. Mustafa Cagri Gurbuz
Professor of Supply Chain Management, Zaragoza Logistics Center

Thesis Supervisor . . . .

Dr. Karen Donohue
Associate Professor of Supply Chain and Operations,
Carlson School of Management, University of Minnesota

Director, Ph.D. Program . . . .

Dr. Maria Jesus Saenz
Professor of Supply Chain Management, Zaragoza Logistics Center

Contents

1 Introduction **17**

1.1 Motivation 17

1.2 Common setting 18

1.3 Research questions 21

2 Literature Review **25**

2.1 Inventory pooling decisions 25

2.2 Capacity choice and allocation 30

2.3 Random yield 32

2.4 Central versus local decision making 34

2.5 Demand forecast information sharing 35

3 Impact of Inventory Allocation Mechanisms on Demand Information Sharing **39**

3.1 Introduction 39

3.2 Setting and timing of events 41

3.3 Properties of Pareto optimal allocation mechanisms 44

3.4 Allocation with symmetric information 46

3.5 Allocation under asymmetric information 50

3.6 Concluding remarks 53

4 Forecast Information Sharing and the Order Quantity Decision: Impact of Inventory Competition and Market Uncertainty **55**

4.1 Introduction 55
4.2 The analytical model 58
4.2.1 Setting and timing of events 58
4.2.2 The information sharing game for $n = 1$ 60
4.2.3 The information sharing and allocation game for $n > 1$ 61
4.3 Numerical study 69
4.3.1 Common knowledge 70
4.3.2 Asymmetric information 71
4.4 Hypotheses 74
4.5 Experiments 76
4.5.1 Experimental design and procedures 76
4.5.2 Experimental results 81
4.6 Concluding remarks 90

5 Forecast Information Sharing and the Order Quantity Decision: Impact of Inventory Ownership 93

5.1 Introduction 93
5.2 The setting 95
5.3 Locally Managed Inventory (LMI) 97
5.4 Central Planner Managed Inventory (CPMI) 104
5.4.1 Is truth-telling an equilibrium? 108
5.5 Numerical analysis 111
5.6 Coordinating transfer prices 119
5.7 Concluding remarks 122

6 Conclusions and Discussion 125

A Notation 131

B Proofs 135

B.1 Proofs of chapter 3 135
B.2 Proofs of chapter 4 139

B.3 Proofs of chapter 5 147

C Additional Numerical Analysis **153**

D Experiment Instructions **155**

D.1 Handout to regional managers: Case R_2D_U 155

D.2 Handout to central planners: Case R_2D_U 159

D.3 Summary sheets 163

E Snapshots from the Experiment Software **165**

E.1 Snapshots of a retailer's computer screen: Case R_2D_U 165

E.2 Snapshots of a central planner's computer screen: Case R_2D_U 168

List of Figures

3-1 The timing of events in the inventory allocation game 43

4-1 The timing of events in the information sharing game with common inventory 59

4-2 $(Q^1 - \mathbb{E}[Q^{CP}])$ as a function of the signal received, for high signal variability, small market uncertainty and asymmetric market sizes . . 73

4-3 Retailers' reported versus actual demand signal 83

4-4 CP's inventory choice versus reported forecasts $(Q(\hat{\theta}))$ 84

5-1 The timing of events in the information sharing game with common inventory 97

5-2 Total Expected Profit as a function of q_1 and q_2 107

5-3 $q_i^{LMI}(\theta_H) - q_i^{LMI}(\theta_L)$ as a function of market uncertainty 116

List of Tables

1.1 Summary of main chapters 24

3.1 Comparison of various allocation mechanisms 52

4.1 Summary of analytical findings of Chapter 4 69

4.2 Optimal inventory for retailer 1 as a function of θ_1 and θ_2 when cr=0.5 70

4.3 $Q_f^1 - Q_f^{CP}$, as a function of θ_1 and θ_2, for cr=0.75 and cr=0.25 71

4.4 Comparison of the optimal inventory for retailer 1 and the expected inventory in the system under truth-telling, as a function of θ_1 73

4.5 Experimental Design 78

4.6 Summary Statistics on Reported $|\hat{\theta} - \theta|$ 82

4.7 Summary Statistics on Total Order Quantity Q 82

4.8 Variable Definition in Equations (4.6)—(4.8) 87

4.9 Impact of competition and market uncertainty on truth-telling, trusting and efficiency 88

5.1 Optimal inventory choices under LMI: $(q_i^{LMI}(\theta_i) = (q_{iL}, q_{iH}))$ 113

5.2 Optimal inventory choices under LMI and common knowledge 115

5.3 Comparison of total inventory and expected profits under LMI, CPMI and CPMI with full information 116

5.4 Comparison of local inventory under LMI, CPMI and CPMI with full information 117

A.1 Summary of notation 133

C.1 Comparison of the optimal inventory for retailer 1 and the expected inventory in the system under truth-telling, as a function of θ_1, for cr=0.75 and cr=0.25 153

C.2 The inventory that maximizes average profit for retailer 1, for various values of θ_1, and the system optimal inventory level for discrete uniform demand 154

Chapter 1

Introduction

1.1 Motivation

One of the critical strategic decisions in supply chain management is whether to decentralize inventory decisions or centralize them at a corporate or supply chain network level. The supply chain management literature often advocates for centralized planning because of its ability to pool inventory risk (e.g. [15, 18]) or coordinate inventory orders to maximize aggregate performance [2, 46]. However, in practical settings, the realized benefits of centralization may be less than analytical inventory models would suggest for a variety of reasons. Possible hidden costs of centralization include managing increased complexity, addressing local incentive conflicts and coordinating decision authority and information structure [1].

One key assumption in most of these models is that the proposed central decision maker can view local demand information at the same level of detail and accuracy as local decision makers when she sets inventory level(s) or when she allocates total inventory to satisfy realized demands. In reality the information that is available to different decision makers within a supply chain, or even within the same company varies [1]. Retailers or local managers may have more information or "feel" for local market conditions because of their proximity to the market. A central decision maker on the other hand, especially later in the demand cycle of a product, may have more information about its popularity because she observes demand information across many locations.

When there is information asymmetry among supply chain parties, reliable information sharing between the better-informed party and the decision maker is critical for the supply chain to operate effectively [11]. Many supply chain failures have been well-documented in the literature due to the ordering behavior of downstream parties that does not reflect their true market needs (e.g., order inflation when capacity shortage is anticipated) [8–10, 38] or the inability of the involved parties to credibly share forecasts (e.g. overoptimistic soft orders that never materialize) [42,44,54]. The latter has a two-fold negative effect: inflated forecasts may lead to an over-investment in inventory or capacity, with possibly serious operational and financial consequences for supplier's profitability, or to persistent shortages due to the lack of credibility of the shared information. For example, overoptimistic forecasts from its customers led Cisco, a major networking equipment supplier, to write off $2.1 billion in excess inventory in 2001 [42]. According to New York Times (March 3, 2007), when UPS cancelled its orders for the Airbus A380 freighters, Airbus had to launch a cost-cutting plan with an expected loss of 10,000 jobs. On the other hand, "Boeing admits that it has been a significant job to persuade suppliers to invest in enough capacity to meet future demand" (The Economist, January 2012).

The goal of this research is to determine whether reliable demand information sharing (realized demand or forecasts) occurs in a setting where a number of retailers form a pooling coalition and strategically interact. We study the impact of inventory ownership, inventory ordering and allocation rights, and behavioral factors, such as trust and trustworthiness on information transmission among supply chain parties.

1.2 Common setting

This thesis focuses on the setting where a group of n retailers pool together their inventories at a central location, managed by a single central planner, to more efficiently satisfy their uncertain local demand. For convenience, we refer to the central planner as *she* and a retailer as *he*. Retailers face the problem of procuring, before the start of the selling season, the inventory that they will stock to serve their stochastic demand for the period. In other words, each retailer faces a classic newsvendor problem [47]. The newsvendor problem applies in a broad array of settings characterized by substantial demand forecast error and a single order or production decision. For example, it could represent the case of fashion apparel retailers, manufacturers of high-technology products (e.g., computers or smart phones) or specific promotional periods. The central planner has the right to set the inventory for the pooling coalition and / or allocate total inventory to its members after local demands are realized.

Given the type of product retailers are trading, the central planner has one opportunity to order inventory before the selling season with no possibility of further replenishments. She procures inventory from an uncapacitated external supplier. The central planner is benevolent, interested in maximizing expected aggregate profits (sum of retailer profits). We assume that retailers within the network do not compete for demand. This is typical in many industries where retailers have exclusive territory rights (e.g., car dealers, franchising). The same dynamics arise in the case of a single company that owns multiple stores that operate in different markets, where each store is a separate business unit (local profit maximizer). Hence, the term "retailer", "regional manager" or "location" are used interchangeably for the purposes of this thesis.

Each retailer has private information about his local demand due to his proximity to the market (e.g., better knowledge of local market conditions and trends). Retailers are homogeneous in their cost parameters. However, they may have heterogeneous expectations about local demand (i.e., different market sizes, competition from substitute products, promotion plans, geographic variation)[1]. Each retailer gains more

[1]We keep our analytical work general allowing for different market sizes at each location but, for comparison purposes and to gain some useful insights, we assume retailers have homogeneous expectations about demand for much of the numerical analysis and experimental work.

knowledge about his local demand as time unfolds, which is captured through a private demand signal known only to him. Cost parameters and demand distributions (before the signal) are common knowledge to all players.

Demand at each location is the sum of three elements: average demand that represents the market size at each location, local private information (demand signal), which represents additional knowledge that becomes available only to the retailer just before the beginning of the selling season, and market uncertainty (demand fluctuations). The central planner knows the distribution of the local demand signal but the retailer learns its exact realization at the beginning of the period. The retailer may obtain this information because of his proximity to the market [42]. Both retailer and the central planner know the average market size at each location and also the distribution of market uncertainty. We assume that demand signals and market uncertainty are independent across markets and that demand signal and market uncertainty in the same region are independent as well.

When retailers decide to form an inventory pooling coalition managed by a central planner there are a number of dimensions to be considered: 1) Who decides on the level of inventory held centrally? Is the decision taken locally or centrally? 2) How will inventory be allocated after demands are realized? Are inventory re-allocations ("transshipments") allowed? 3) When inventory level and allocation decisions are taken by the CP, what is the available information they are based on? If both inventory and allocation decisions are delegated to the central planner, and both demand signal and demand realizations remain retailers' private knowledge, the setting can be described as follows. Retailers receive their demand signal and they communicate their forecast to the central planner who determines total inventory. Demands are realized and retailers simultaneously submit their orders to the central warehouse who allocates total inventory among retailers. This is a double information asymmetry case. The central planner does not know demand signals but she sets inventory based on the forecast information solicited by retailers and she does not observe demand realizations and thus she bases inventory allocation on the final orders placed by the retailers.

In this thesis, we study the issue of inventory allocation when final demands remain retailers' private knowledge as well as the issue of demand forecast information sharing between retailers and the central planner who sets inventory for the coalition. Due to the analytical complexity and mathematical intractability of the general problem, we study these two separately. In chapter 3 we assume that final demands remain private knowledge, and we study the ordering behavior of retailers under various allocation mechanisms. For chapters 4 and 5, we shift focus to the quality of forecast sharing and its influence on the quality of the inventory decision. For this purpose, we assume that realized demands become common knowledge to all players. This captures the dynamics of a single firm operating in many geographical markets or a third party logistics provider that monitors sales data at a retail location and decides on its inventory (Vendor Managed Inventory principles).

1.3 Research questions

As aforementioned, the thesis consists of three main chapters, each one studying a different aspect of demand information sharing within a supply chain. In chapter 3 we focus on the influence of allocation mechanisms on quality of demand information passed through final orders. In chapter 4, we study the forecast information sharing game, when inventory is common and inventory allocation is assigned proportionally to realized demands. We focus on how competition for common inventory, market uncertainty and behavioral factors affect information sharing within the pooling coalition. In chapter 5, we explore the effect of inventory ownership, employing the notion of claims, on retailers' incentives to share truthfully their demand forecasts. We also compare central to local (each retailer decides his own inventory while lateral reallocations are allowed) decision making regimes, studying the resulting system inventories and profits. Next, we present in more detail the specific business problem each chapter relates to and the specific research questions it addresses.

When a number of retailers form an inventory pooling coalition, the issue of allocation after demands are realized becomes very important. Allocation, or rationing of limited common inventory, is an ongoing issue in many industries ranging from automobiles, to pharmaceuticals and toys [8]. In the presence of demand and supply imbalances, retailers may have an incentive to misreport their needs in order to gain a more favorable allocation. In chapter 3, we assume the pooled inventory quantity is given and focus on how this inventory is allocated between retailers when the central planner must rely on retailers to pass final demand information. We consider various allocation mechanisms and their effect on retailers' ordering behavior after local demands are realized. The main questions include: (a) Which allocation mechanisms are Pareto optimal when realized demands become common knowledge? (b) Do Pareto optimal allocation mechanisms exist that are also truth-inducing when realized demands remain retailers' private knowledge? (c) Under which allocation schemes will retailers truthfully report their local realized demands (through order placement) to the central planner who allocates inventory?

As already mentioned, many supply chain failures (e.g., excessive inventory, insufficient capacity to meet demand) are due to the inability of supply chain parties to credibly share demand information. Chapter 4 focuses on the demand forecast sharing game between retailers and the central planner who decides, at the beginning of the selling season, the common inventory quantity to be held for the coalition. Credible demand forecast sharing is critical in this context given that the central planner has less accurate information about demand in each of the regions. Recent research has shown that reliable information sharing in a supply chain depends not only on parties' incentives but also on behavioral factors between supply chain parties [42]. Hence, in chapter 4, we first use game theoretical models to study players' pecuniary incentives in such a setting and then we conduct laboratory experiments to study the impact of behavioral factors, such as trust and trustworthiness on the communication. In this work we define trust as the CP's willingness to rely on retailers' forecasts to determine total inventory. Trustworthiness is measured by the difference between a retailer's real forecast and the one he reports to the CP. We also study, both analytically and experimentally, the role of supply chain environment (i.e., number of retailers, market uncertainty) on the result of the communication. The main research questions are: (a) Do retailers truthfully transmit their demand forecasts (private information) to the central planner when they compete for inventory? (b) Does the CP incorporate retailers' transmitted forecast information in her inventory decision (i.e., trust the shared demand information)? (c) How are communication dynamics affected by the supply chain environment, i.e., number of retailers, level of forecast uncertainty, retailers' trust in the CPs inventory decision process?

In chapter 5, we continue studying the demand forecast sharing game between locations and the central planner, keeping the same information structure. In contrast to chapter 4 where the inventory held centrally is common, in this chapter we examine the impact of inventory ownership on individual players' incentives. We study the quality of communication when each region has a dedicated inventory quantity and unilateral inventory reallocations are allowed after demands are realized. Unilateral transshipments (or reallocations of centrally held inventories) is a common practice in many industries such as OEMs [46] and pharmaceuticals. We first study the impact of inventory decision rights placement (locally versus centrally) on the resulting inventory levels and profits. We then examine whether inventory ownership may align retailers' and system's incentives so that, in the case of central decision making, truthful forecast information sharing can be expected. The main research questions are: (a) What is the impact of inventory ownership on the reliability of information transmitted? To be more specific, do the dynamics of information sharing change when each retailer has a dedicated inventory quantity held centrally? (b) What are the resulting inventory levels when inventory decision rights are placed locally (i.e., to retailers) versus centrally (i.e., the right to decide on inventories is transferred to the central planner)?

Table 1.1 summarizes the settings we consider in the following chapters, comparing them along the dimensions of decisions studied and information transmitted. The pooling coalition (retailers and the central planner) is the system we are considering. For example, when we refer to “endogenous” inventory quantity, we mean that it is a decision variable for either retailers or the central planner. Before presenting the main chapters of the thesis, in the next chapter we review the related literature, mainly within the field of operations and supply chain management but not exclusively, and we position our research within each research stream.

	Inventory Quantity	Allocation	Information Asymmetry
Chapter 3	Exogenous	Various rules	Final demands
Chapter 4	Endogenous	Proportional	Demand forecasts (signals)
Chapter 5	Endogenous	Dedicated Inventories w/ reallocation possibility	Demand forecasts (signals)

Table 1.1: Summary of main chapters

Chapter 2

Literature Review

In this chapter, we review the streams of literature this thesis is mainly related to, considering both analytical and experimental papers. We first review the literature related to inventory pooling decisions and associated behavioral issues. We next examine related literature on capacity choice and allocation, random yield, central versus local decision making and demand forecast information sharing, focusing on the issue of information credibility. We comment on how our work fits into each of these fields and why it differs from previous research.

2.1 Inventory pooling decisions

Benefits of inventory pooling

Pooling, the strategy of aggregating demand streams, has long been studied under two main contexts: reducing product variety (i.e., through SKU rationalization) or geographical variety (i.e. inventory centralization). A third dimension of aggregation is that of demand aggregation across time (i.e., order consolidation). Reducing variety will generally allow a company to provide the same customer service level at a lower cost, or to improve it without incurring extra costs. Eppen (1979) [18] in his seminal work showed that total holding and stockout costs are lower when demand is aggregated, under the assumption of independent and normal demand distributions and identical cost parameters (one-period setting). Eppen and Schrage (1981) [19] extend the result to the multiple period problem while Corbett and Rajaram (2006) [15] generalize Eppen's model to (almost) arbitrary multivariate dependent demand distributions. In this thesis we study the benefits of pooling when there is information asymmetry. We assume that a central planner has less information about the demand distributions than that of the decision makers in the decentralized case (due to her distance from the markets). To the best of our knowledge this is the first attempt to incorporate information asymmetry in the problem of inventory pooling.

Sharing the benefits of inventory pooling

Another stream of papers studies how the benefits of inventory pooling should be shared among retailers, so that the pooling coalition is stable, in the sense that all participating retailers have higher expected profits. Hartman et al. (2000) [26] employ cooperative game theory to prove the existence of a cost allocation scheme so that all retailers are better off (incur lower costs) by pooling their inventories together, under the assumption of symmetric and independent demand distributions, identical overage and underage costs, or a multivariate normal demand distribution. In other words, they show that, under these assumptions, the core of the inventory centralization game is non-empty, implying stability of the system. They do not study which cost allocation schemes or mechanisms of inventory rationing are on the core of the game. Muller et al. (2002) [40] extend this result for all possible joint distributions of random demands and give sufficient conditions under which the core is a singleton and conditions under which at a core allocation every newsvendor shares a nonnegative cost. Gerchak and Gupta (1991) [22] study whether "popular" cost allocation methods, namely cost allocation by demand volume, by individual safety stock requirements, by incremental contribution to joint costs, and proportional to stand-alone costs, are potentially unfair (in the sense that a retailer may be charged more under consolidation than his cost would be under a dedicated inventory system).

Only apportioning costs according to each retailer's stand-alone costs is found to be "fair" in centralized continuous review inventory systems. Robinson (1993) [45] uses cooperative game theory and, through a counterexample, proves that this basis of allocation is not always in the core of the game; some customers may be worse off when new customers join up. He uses the Shapley value allocation rule to determine "fair" cost allocations in which no one is worse off after consolidation [48]. This allocation ensures that each customer is charged somewhere between their incremental and stand-alone costs.

Kemahlioglu-Ziya and Bartholdi (2011) [31] study a two-stage supply chain consisting of a supplier and multiple retailers and find that if savings of centralized inventory are allocated among supply chain members by Shapley value, the pooling coalition is farsightedly stable. This allocation may not always belong to the core of the centralization game if retailers are not identical but it always coordinates the supply chain and distributes profits in a "fair" way. They move one step further by identifying a quantity allocation rule (a modified linear allocation) of shared inventory where players' expected after-pooling profits are equal to their Shapley value allocations plus their before-pooling profits and the supplier carries the supply-chain-optimal level of inventory.

In this stream of research, cost or profit allocation is based on full information. In our work we abstract from the issue of excess profit allocation by assuming that all retailers join the inventory pooling mechanism. This is representative of situations where retail stores belong to the same parent company or are otherwise constraint to accept a pooling policy. In other words, we study the information sharing dynamics within a coalition, after its synthesis has been set. In chapter 3, where different allocation mechanisms are compared, we identify one that guarantees each individual retailer profit's is at least as high as in the decentralized system.

Behavioral issues in inventory pooling decisions

It has long been acknowledged in the decision making (i.e. Tversky and Kahneman, 1974 [56]; Thaler, 1980 [55]) and economics literature (Mullainathan and Thaler, 2000) [39] that decision makers are not perfectly rational expected profit maximizers. Their decisions may systematically deviate from optimality for a variety of reasons. For example, people exhibit systematic biases in their judgments and rely on a limited number of heuristic principles to assess uncertain situations [56], they have limited cognitive abilities and bounded self-interest [39], they may have other social considerations, e.g., fairness, reputation, social norms [14]. In the field of operations and supply chain management, there has been a growing interest in behavioral research. Behavioral research studies the effects of human behavior in processes and performance by incorporating social and cognitive psychology considerations [17]. For a comprehensive review of this stream, please see Bendoly et al. 2010 [3] and Donohue and Siemsen 2010 [17]. We continue with reviewing behavioral literature that is specifically related to inventory and pooling decisions.

There is a vast literature on how people make newsvendor problem decisions [4, 5, 7, 47]. We focus on multi-player newsvendor settings. Su (2008) [53] builds a decision model based on the quantal choice model (best decision need not always be made but better decisions are made more often) and applies it to the newsvendor setting. He predicts that when inventory for multiple locations is held centrally, apart from the benefits associated with variance reductions, behavioral benefits also exist as inventory centralization helps by pooling decision errors across locations ("supply uncertainty" pooling). Lavaro and Corbett (2003) [33] study analytically and through simulations the pooling effect in the context of SKU rationalization when inventory policies are suboptimal and demand is non-normal (single planning period). They find that the value of pooling may be negative when the inventory policy in use is suboptimal while it varies little across the distributions they studied. More recently, Ho et al. (2010) [28] study experimentally the ordering behavior in multilocation inventory systems and they show that systematic biases eliminate the risk-pooling benefit when the demand across stores is strongly correlated. They propose a behavioral theory based on reference dependence (psychological aversion to leftovers is greater than the disutility of stockouts) to explain / predict ordering behavior in a multi-location newsvendor framework, under both a centralized and a decentralized inventory structure.

Kremer (2007) [32] studies experimentally the impact of secondary markets (opportunity of inventory rebalancing after demands are realized) on newsvendor's ordering decisions. Inventory reallocations are possible, at exogenous or endogenous market prices determined through a market-clearing auction mechanism, allowing for demand risk pooling. He finds that the option to trade units in the secondary market increases supply chain profitability in all cases. More interestingly, it also has a beneficial impact on inventory decisions ahead of the selling season as it induces average order quantities to regress towards system-optimal levels. Behavioral issues related to transshipment as a pooling strategy are studied as well by Bostian et al. (2012) [6]. Unlike Kremer's work, transshipment decisions are considered automatic (both quantities and prices) and the main focus is on whether behavioral bias in ordering nullify the risk-pooling benefit of transshipments. They find that transshipment is a behaviorally robust risk-pooling technique in the sense that it gives an even greater benefit in the presence of behavioral biases than in the absence of these biases.

Our work considers behavioral factors in information sharing within an inventory pooling coalition. On the one hand, we study the role of trust and trustworthiness in a pooling coalition where information asymmetry among players is present. We compare the predictions of analytical models based solely on pecuniary payoffs to the results of controlled laboratory experiments with human subjects. On the other hand, we show that under information asymmetry and when behavioral factors are taken into consideration, the benefit of inventory pooling may, in practice, be inexistent or even negative, even when demands across locations are independent.

2.2 Capacity choice and allocation

The framework employed to study the dynamics of a centralized inventory system when retailers strategically interact closely parallels that of capacity choice and allocation in a single supplier, multiple retailer (or manufacturer) context. One main characteristic of the capacity choice and allocation literature is that it employs a two-stage supply chain in contrast to the classical inventory pooling literature where retailers decide to collude among themselves (i.e. forming an aggregate retailer) to increase total profits. Cachon and Lariviere (1999, 1999 b) [8,10] study how the capacity allocation scheme affects retailers' orders and supply chain performance when the supplier has limited capacity and the retailers have private information about their optimal stocking levels. They find that under any Pareto allocation mechanism (e.g., proportional or linear), all retailers truthfully reporting their optimal allocations is not a dominant Bayesian equilibrium and that a manipulable mechanism (not truth revealing) may lead to higher capacity and higher profits for everyone. Our work differs in that "the central planner", who decides about the inventory (rather than capacity) to be held centrally, is a total system profit maximizer. Cachon and Lariviere (1999 a) [9], in a two period setting and with two retailers, find that linking a retailer's current allocation to his previous sales rate ("turn and earn" allocation) does not generally coordinate the system. Retailers sell more but in equilibrium no one gains an advantage. Lu et al. (2010) [37] extend this work to an infinite horizon game with multiple retailers and find a richer set of equilibria. "Turn and earn" allocation may reduce demand variability placed on the supplier as retailers absorb local demand fluctuations.

Compared to this stream of literature, the timing of the allocation in our work, and therefore the information it is based on, differs. In the capacity rationing literature, manufacturers usually order from a capacitated supplier a quantity at the start of the selling season, before their demands are revealed. In the pooling coalition case considered in this thesis, total quantity is procured from an uncapacitated external supplier at the beginning of the selling season and it is held centrally until regional demands are realized. Hence, the final allocation happens after demand uncertainty is resolved because of the different supply chain tier considered. In our work, allocation by the central planner is based on realized demands if they are known (chapters 5 and 6) or on final orders that are placed by the retailers after they see their final demands (chapter 3). In chapter 3, we consider inventory allocation mechanisms that are similar to the capacity allocation mechanisms studied in these papers (modified accordingly to represent our setting), while in chapters 4 and 5 we take the allocation rule as given and we focus instead on the issue of optimal central inventory choice for the coalition.

Li et al. (2011) [35] consider situations where firms buy options to use the capacity of a supplier. They explore whether the supplier benefits from providing transfer rights with these options such that a firm that cannot use all its purchased capacity can sell it to another firm that may need it. They examine what happens when a buyer requires more or less than her reserved capacity under three policies: a) individual reservations are final and excess demand at any buyer is lost, b) any unused capacity by each buyer returns to the supplier and he has the right to sell it to another buyer if there is demand for it and c) each buyer owns her reserved capacity and can resell any unused capacity to another buyer after demands are realized. They find that providing transfer rights to the buyers can be better for the supplier than reserving the transfer rights to itself in a wide variety of situations because the buyers value their reserved capacity more which in turn allows the supplier to charge higher reservation prices. Their setting parallels ours, mainly the one considered in Chapter 5, in the sense that they are comparing final “individual capacity reservations” to “pooled” capacity with minimum guarantees (allowing “capacity transshipments”) and they are focusing on the issue of who has the right to transfer unused capacity. In the case where the supplier owns the property rights to reserved capacity, and both buyers need capacity above their reserved quantities, the uniform capacity allocation rule is employed in order to exclude retailers’ strategic ordering and to keep the model tractable. We are considering a central total system profit maximizer instead and we are focusing on the gaming between the players in the presence of information asymmetry.

The only work, to the best of our knowledge, that explicitly models and experimentally estimates behavioral factors in an allocation game, is that of Chen et al. (2012) [13]. They consider a setting of complete information, capacity shortage, known demands and proportional allocation. They find that retailers do not order that much more than what they need, as game theory would predict as a result of their strategic interaction. They propose a model of bounded rationality, based on the quantal response equilibrium, to explain the observed ordering behavior in the lab. For our experimental work, we base the allocation of total inventory on realized demands (instead of orders placed) and we focus on the information sharing game between the retailers and the central planner who sets the inventory quantity.

2.3 Random yield

In our setting, the final allocation that a retailer or a region gets is random because it is based either on realized demands at two or more locations (chapters 4 and 5) or on retailers' final orders that, in turn, depend on random demand realizations (chapter 3). The allocation that each retailer receives can be thought of as a random yield of the total inventory quantity and hence our work is also related to the random yield literature. The yield factor will depend on the allocation rule employed. Analytical models of determining lot sizes when production or procurement yields are random can be classified in two main categories: single-stage continuous time review models where demand is constant or random and periodic review (discrete time) models for single or multiple production stages and single or multiple periods (known or random demand). Yano and Lee (1995) [59] provide a comprehensive review on quantitatively-oriented approaches for determining lot sizing with random yields. Our setting is closer to the discrete time, single period, single run model studied first by Shih (1980) [49]. He focuses on the case where yield uncertainty is caused by defective units and he shows that when the distribution of the % of defectives is known and invariant with the production level (stochastically proportional yield) the expected cost is convex and a closed form solution for the optimal procured quantity can be found. Gerchak et al. (1988) [23] study periodic review production models with variable yield and uncertain demand but they start their work with the single period case. It is the first paper that provides a complete analysis of a general profit maximizing single-period model with variable yield and uncertain demand and they show that the profit function is concave in the lot size. Henig and Gerchak (1990) [27] study the structure of periodic review policies in the presence of random yield and they start as well by studying the single period case. They show that for linear holding and shortage costs (as in our case) the probability of shortage under random yield is no smaller than the probability of shortage under certain yield. Even if in our work allocation at each location is a random part of the total inventory, the modeling of "yield" uncertainty, through the allocation mechanism, is quite distinct than the classic random yield literature (e.g., the number of good units follows the binomial distribution, the fraction of good units (yield rate) follows a known distribution etc).

It is worth mentioning that in one of the earliest papers in the random yield literature [50], it is acknowledged that the randomness of the amount received from an order may come from many sources of uncertainty, including human administrative errors. Silver (1976) refers to the ratio of the expected amount received to the lot size as the "bias factor". In our setting, there are two sources of randomness in the yield; one resulting from the fact that retailers in a pooled setting will receive a random fraction of the total inventory and a second resulting from errors in decision making (order quantity from the central planner may not always be the one that maximizes expected profit). Our modeling work is based on the assumption that the optimal order quantity is always ordered but experimental results do not exclude that ordering error may also be present.

2.4 Central versus local decision making

The analysis of classical distribution systems usually assumes a single (central) decision maker or that decisions are somehow coordinated by a central agency. However, distribution systems may consist of multiple decision makers with their own objectives and without an explicit mechanism that coordinates decisions. Along these lines, Rudi et al. (2001) [46] consider a two-location inventory model where transshipments are allowed, and examine how this possibility affects optimal inventory levels at each location. They consider local decision making and contrast it to central decision making by comparing the resulting inventory levels. Anupindi et al. (2001) [2] introduce a general framework for the analysis of decentralized distribution systems. They consider a setting of both local and central inventories and they study together inventory ordering and allocation decisions. Our work (chapter 5) closely parallels these papers with respect to the dynamics that arise. Even if inventories are held centrally, each unit of inventory "belongs" to a specific location (retailer) and change of ownership for units of inventory is allowed after demands are realized (lateral inventory reallocation). The critical difference is that local decision makers have more accurate information about their local demands due to their proximity to the market (compared to the central decision making case). Therefore, our work could be considered as an extension in the sense that it incorporates information asymmetry.

Anand and Mendelson (1997) [1] model different coordination structures within a monopolistic firm that operates in multiple horizontal markets. They jointly consider decision rights (who decides what) and information structure (who knows what), and quantify the value of collocating decision rights with specific knowledge. They discriminate between information that can be transmitted and local knowledge that cannot be communicated, and study the information structure jointly with the decision making structure. They consider a firm with increasing marginal cost and linear in price demand curves. They take the system perspective, assuming that all players want to maximize expected overall profits, but also consider transfer-pricing schemes that address incentive conflicts between local and overall profitability. In our work, we take the information structure as given, we assume that all available local knowledge can be efficiently communicated if the right incentives are in place and we focus on whether truthful sharing of private information is expected, when retailers operate in competitive markets (exogenous price, constant cost and high demand uncertainty).

In similar lines, Netessine and Rudi (2003) [41] study centralized and competitive inventory models with demand substitution. They compare under these two settings optimal inventory stocking policies for a given product line when demand substitution is consumer driven. Our work analyzes centralized versus competitive inventory models for a single product but multiple retailers that they do not directly compete with for demand. However, they do strategically interact for inventory allocation in case of a centralized system. Jiang et al. (2010) [30] study the optimal stocking levels of multiple newsvendors that compete for demand but they have asymmetric information about future demand realizations, and this information is limited to knowledge of the support of the demand distribution. They characterize the Nash-equilibrium stocking quantities by applying a robust optimization methodology. They find that a competing newsvendor does not necessarily benefit from having better information about its own demand distribution than its competitor. This work, related to a decentralized competitive inventory setting relaxes the assumption of complete and symmetric information that game-theoretic models in operations management primarily rely on. Similarly, we relax the assumption of complete and symmetric information in a distinct setting: we consider a two stage supply chain when the central planner plays an active role in setting total inventory while, in both local and central decision making cases, retailers do not compete for demand but strategically interact for inventory allocation.

2.5 Demand forecast information sharing

While there has been substantial work in operations and supply chain management regarding the value of sharing (demand) information, it is usually assumed that the exchange of information is truthful and credible. For a comprehensive review on the gains of information sharing see Chen (2003) [12]. For example, Li (2002) [34] and Zhang (2002) [60] both study the issue of information sharing in a two-level supply chain where there is one upstream and multiple downstream competing firms, which have better information either about demand or their costs. In these papers, firms operate in a duopoly market and compete for demand (in contrast to the newsvendor setting with seperate markets we are considering) but most importantly, if players agree to share information it is assumed that they will do it truthfully. The main focus is on how information sharing affects the profitability of the firms and whether firms will engage in such activities. Cachon and Lariviere (2001) [11] study demand forecast sharing in a supply chain with one manufacturer and one supplier. The player who is closer to the market (the manufacturer), as in our case, has better forecast information. However, he has an incentive to inflate his forecast to induce the supplier to build more capacity. They study contracts that allow the supply chain to share demand forecasts credibly under two contract compliance regimes: forced and voluntary compliance of the supplier. They find that it is always in the interest of the manufacturer with high demand forecast to share credibly his forecast (separating equilibrium), but doing so creates system inefficiencies compared to the full information case, as credible signaling comes at a cost.

The issue of whether information sharing within a supply chain is truthful and to what extent, both in theory and in practice, has recently received some attention in the literature. Li and Zhang (2008) [36] consider the issue of truthful information sharing in a setting where one manufacturer supplies to multiple retailers competing in price. They study the effect of confidentiality in information sharing and they find that if retailers share their private information about demand confidentially with the retailer, truth-telling is an equilibrium and supply chain coordination is achieved.

Our work is very close in spirit and could be considered an extension of Ozer et al. (2011) [42]. In this paper, the authors explicitly consider the issue of forecast information sharing in a single supplier — single manufacturer context where the manufacturer has better information about the stochastic demand. They experimentally test the analytical predictions of game theory in such a setting (cheap talk, uninformative communication) and find that there is a continuum of trust when people share information. Trust among the decision makers significantly affects the outcome of cheap-talk forecast communication that is not totally uninformative in practice and improves channel efficiency. Our paper extends their work to multiple retailers while the incentives of the central planner are distinct. We do not consider manufacturing and capacity decisions but inventory decisions instead within a pooling coalition. A closely related, follow-up working paper compares forecast information sharing in China and the U.S., investigating the effect of the cultural and institutional heterogeneity in trust and trustworthiness in this context [43]. They find that Chinese consistently exhibit lower trust and trustworthiness than their U.S. counterparts.

Ren et al. (2010) [44] study forecast sharing in a long term supplier-customer relationship. They prove that a truth-sharing outcome can emerge as an equilibrium through the repeated supply chain relationship, even when a linear price contract is employed between the supplier and the customer. They identify a multi-period review strategy profile that not only can induce truthful information sharing, but also system coordination. We employ instead a single-period setting and we study how behavioral factors, i.e., trust, may lead to truthful communication even in the absence of repeated interactions. Terwiesch et al. (2005) [54] empirically study, in a specific industry (semiconductor equipment industry), the demand forecast sharing process between one buyer and a number of suppliers. They find that suppliers penalize the buyer for unreliable forecasts by providing lower service level, while, in turn, the buyer overly inflates his forecasts to the suppliers with poor service history.

There has been some additional research recently that examines the role of trust and trustworthiness in information sharing in different operational settings. For example, Inderfurth et al. (2013) [29] shed more light on the role of trust and trustworthiness in communication under asymmetric information within a supply chain, by studying the effect of pre-play communication in a principal-agent setting. They conduct a laboratory experiment and they find that information sharing, reduces inefficiencies in the supply chain despite the cheap-talk dynamics, if there is a certain amount of trust. In a similar context, Voigt and Inderfurth (2012) [57] study analytically the effect of communication and trustworthiness on supply chain performance. They allow pre-play communication in a supplier-buyer relationship that falls into the principal-agent category and assume that a proportion of the agents are always truth-tellers.

In this thesis we consider the issue of forecast sharing in a cheap talk setting without mechanism design contracts or repeated interactions. We focus on whether truthful information is an equilibrium when multiple retailers being part of a pooling coalition strategically interact and compete for inventory. We study the effect of competition, market uncertainty and inventory ownership on players' incentives for truthful forecast sharing. We then study the role of trust and trustworthiness on the result of communication and cooperation within the supply chain by conducting a laboratory experiment.

Chapter 3

Impact of Inventory Allocation Mechanisms on Demand Information Sharing

3.1 Introduction

When a number of retailers form an inventory pooling coalition, the issue of how inventory is allocated after demands are realized becomes very important. In the absence of a pricing mechanism that would balance supply and demand, inventory rationing through quantity limits (upper or lower) in case of shortage or surplus are necessary. In such cases though, retailers may have an incentive to misreport their needs in order to gain a more favorable allocation. Allocation, or rationing of limited common inventory, is an ongoing issue in many industries where there is only one opportunity for production or procurement, before the start of the selling season. For example, allocation mechanisms have been employed in the fashion apparel, consumer electronics and, automotive industries [8, 9, 38].

The issue of capacity rationing in a supply chain with one supplier selling to multiple retailers is a phenomenon well-studied in the literature [8–11]. When several retailers compete for limited capacity, a broad class of allocation mechanisms are prone to manipulation. When inventory is limited, prior analytical research has shown that retailers may order more than what they need to gain a higher allocation [10]. Our setting differs from this prior research in three ways: (a) the manager of the pooling coalition (CP) is interested in maximizing total supply chain profits (she is not a seperate business entity trying to maximize her own profits), (b) when allocation takes place, retailers already know their realized demands (not their optimal stocking levels under demand uncertainty), and (c) retailers are responsible for both under-stocking and over-stocking costs of the coalition. In contrast to the case of capacity rationing, total inventory of the coalition is held in the central warehouse until local demands are realized. After demands are realized and final orders are placed to the central warehouse, retailers may receive less or more than their final order. Retailers collectively assume all demand uncertainty risk, as the central planner is not a separate unit with its own financial objectives.

In this chapter, we abstract from the issue of how the quantity in the central warehouse is set (i.e. the level of inventory held at the CP is for now assumed to be given) and study how retailers pass orders under different allocation mechanisms. We study three allocation mechanisms that are commonly used in practice and analyzed in the literature. These include *proportional, linear* and *uniform*. Perhaps the most intuitive mechanism is the proportional allocation where each retailer receives a proportional amount of his announced demand (or order when demands are not common knowledge). Linear allocation gives each retailer his final demand (or order) plus / minus a common quantity when there is inventory surplus / shortage. Under uniform allocation, each retailer gets the same quantity, under some conditions. No retailer gets more than what he asks for when total orders exceed total inventory and no retailer gets less than what he asks when the reverse is true. Under common knowledge, these conditions guarantee that there are no unsold units when there is inventory shortage and that there is no unmet demand when there is inventory surplus. Furthermore, we introduce a new allocation rule, the *modified uniform allocation* rule, which has some attractive properties in this setting. This allocation mechanism connects retailers' initial (soft) orders to the final allocation they receive after demands have been realized.

The research questions we address in this chapter include: (a) Which allocation mechanisms are Pareto optimal under common knowledge of realized demands? (b) When final demands remain retailers' private knowledge, under which allocation rules, if any, will retailers truthfully report these to the central planner (who allocates inventory)? (i.e. place a final order equal to their realized demand) (c) Do Pareto optimal allocations that induce truthful demand telling exist?

The rest of the chapter is organized as follows: we start by further defining the supply chain setting we are studying and the timing of events. Then, we identify properties that Pareto optimal allocation mechanisms have in this setting. Next, we introduce the various allocation mechanisms we are considering in this chapter. We first study the case of symmetric information (the CP knows the realized demands when she allocates inventory) which serves as a benchmark. Then, we analyze the asymmetric information case. Each retailer knows his realized demand (and only his) while the CP does not see actual demands but only the orders retailers place.

3.2 Setting and timing of events

A number of retailers n have decided to form a pooling coalition to better satisfy their uncertain demand. At the beginning of the season, each retailer receives a signal θ_i about the demand at his region ($i = 1...n$). Demand at each region is given by $d_i = \mu_i + \theta_i + \epsilon_i$, where μ_i is a positive constant representing the average demand at location i and θ_i is retailer's i private information about demand, a zero-mean random variable with cumulative distribution function $F_i(\cdot)$, probability density function $f_i(\cdot)$ and support $[\underline{\theta_i}, \overline{\theta_i}]$ (adopted from Ozer, 2011 [42]). Market uncertainty is represented by ϵ_i, a zero-mean random variable with cdf $G_i(\cdot)$, probability density function $g_i(\cdot)$ and support $[\underline{\epsilon_i}, \overline{\epsilon_i}]$. Both retailer i and the central planner know μ_i and $G_i(\cdot)$. Retailer i also knows at the beginning of each selling season the realization of θ_i, while the central planner knows only its distribution $F_i(\cdot)$. We assume that signal and market uncertainty are independent across regions and that within a region market uncertainty is independent of the signal received (i.e., $\text{Cov}(\epsilon_i, \epsilon_j) = 0$, $\text{Cov}(\theta_i, \theta_j) = 0$ $\forall i, j = 1, 2...n$, $i \neq j$ and $\text{Cov}(\epsilon_i, \theta_j) = 0$ $\forall i, j = 1, 2...n$). Cost parameters are also common knowledge to all players. Each retailer receives p for each unit he sells and pays c for each unit he receives from the central warehouse. The salvage value of the product at the end of the selling season is normalized to zero.

In this chapter, we focus on the information transmission and retailers' incentives at the inventory allocation stage. Hence, the level of inventory held by the CP is for now assumed to be a given quantity (i.e., it is a parameter). Central inventory can be determined in many ways, e.g., as the sum of local inventory level calculations, or set optimally by the CP given his full knowledge of the demand distributions at the time that the inventory level decision is made. In all cases though, we assume that the level of common inventory is independent of the allocation mechanism employed. For the rest of the chapter we consider the case where inventory level is determined as the sum of retailers' initial (soft) orders q_i^a. It is important to point out that these initial orders could be decided in a number of ways (e.g., optimal newsvendor quantity for each retailer) but retailers do not have freedom in choosing the way q_i^a is calculated. Hence allocation does not change the way retailers behave. In this sense, we ignore at this stage the interaction between the allocation mechanism and the total inventory set. The same assumption would hold if the CP announces the allocation mechanism along with the inventory level set for the coalition at the same time.

The central planner (CP) is responsible for managing total inventory at the central warehouse and for distributing it to retailers after their local demands are realized, according to a publicly announced mechanism. After he observes his realized demand, each retailer places a second, final, order ($q_i^b \geq 0$) to the central warehouse. We denote by $\boldsymbol{\alpha}$ the vector of quantities sent to retailers. The timing of events, depicted graphically in Figure 3-1, is as follows:

0. Before the period begins, the central planner announces the allocation mechanism she will use once demands are realized (or after retailers have placed their final orders in the asymmetric information case).

1. At the beginning of the period, each retailer i observes a private signal θ_i (only his) about his demand at location i and places an order q_i^a to the central planner.

2. The central planner calculates total inventory to be held centrally, $Q = \sum_i^n q_i^a$. This quantity becomes common knowledge to all players.

3. Demand is realized. In the symmetric information case after local demands are realized they become common knowledge to all players. In the asymmetric information case, where only retailers know their own d_i's, retailers simultaneously place a second order q_i^b to the central planner.

4. The allocation of inventory to retailers ($\boldsymbol{\alpha}$) is calculated according to the posted allocation mechanism based on Q and the available information. The central warehouse fulfills the orders submitted by sending to each retailer α_i. Each retailer is charged with $c \cdot \alpha_i$.

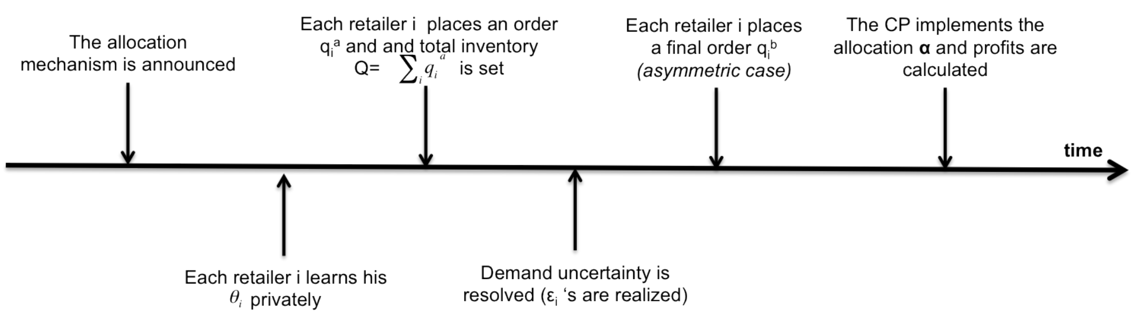

Figure 3-1: The timing of events in the inventory allocation game

Before we analyze the symmetric and asymmetric information case, we begin by formally defining a Pareto optimal allocation mechanism and identifying some properties of such mechanisms in this context.

3.3 Properties of Pareto optimal allocation mechanisms

An allocation mechanism in this case, where demand uncertainty is resolved before allocations are determined, is *efficient* when it excludes wastage. System efficiency implies that there are no unsold units when total demand is equal to or higher than Q and there is no locally unsatisfied demand when total demand is lower than Q. In other words, no retailer faces a shortage while at the same time another retailer faces a surplus (i.e., it cannot be that $\alpha_i > d_i$ and $\alpha_j < d_j$ for some $i, j = 1...n$).

We denote by $\mathbf{d}$ and $\mathbf{q}^b$ the vector of local demand realizations and the vector of final orders placed by retailers, respectively. An allocation mechanism is *Pareto optimal* if it maximizes the sum of retailers profits and thus system profits, assuming in the asymmetric information case that all retailers truthfully report their realized demand by ordering $\mathbf{q}^b = \mathbf{d}$. Please note that initial orders q_i^a are soft orders used to determine total inventory and are not sent to retailers (q_i^b is *not* on the top of q_i^a).

Proposition 3.1:
(a) A Pareto optimal allocation mechanism satisfies: $\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = \frac{\partial \pi_j(\alpha_j, d_j)}{\partial \alpha_j} \;\; \forall i, j$.
(b) An allocation mechanism is Pareto optimal *if and only if* it is efficient.
All proofs are provided in Appendix B.1.

Next we show that in our setting, a Pareto optimal allocation mechanism is not necessarily individually responsive. An individually responsive mechanism ensures that if a retailer is receiving not zero but positive allocation ($\alpha_i > 0$), his allocation quantity increases when his demand (or his final order quantity in the asymmetric information case) increases, unless he has already been allocated the total quantity. Similarly, under such a mechanism, allocation quantity decreases when demand (or order quantity) decreases.

Following the definition of Cachon and Lariviere (1999) [10], an allocation mechanism is *individually responsive* if, for all i, $0 < \alpha_i(d_i) < Q$ implies

$$\alpha_i(\hat{d}_i, d_{-i}) > \alpha_i(d_i, d_{-i}), \quad \hat{d}_i > d_i \tag{3.1}$$

Proposition 3.2 states that being individually responsive is not a necessary condition for Pareto optimality in our setting.

Proposition 3.2: A Pareto optimal allocation mechanism is not necessarily individually responsive.

This is an important result because the (modified) uniform allocation mechanism, which is not individually responsive, is not excluded from the set of potentially Pareto optimal allocations. This result is different from the result of Cachon and Lariviere (1999) [10] who show that when allocation of scarce capacity to retailers is done before demand uncertainty is resolved if an allocation mechanism is not individually responsive it cannot be Pareto optimal.

The intuition behind their result is that a Pareto mechanism "must recognize the smallest change in every retailer's marginal valuation of stock and thus must be individually responsive". If a retailer has a higher optimal stocking level, it is optimal to receive a higher allocation. In the setting under consideration, because allocation is decided after demand uncertainty is resolved at each location (and retailers have identical cost parameters), the incremental value of an additional unit of stock is constant; it is $p - c$ when there is inventory shortage and $-c$ in case of total inventory surplus. Hence, responsiveness is not a necessary condition for allocation optimality (unlike the capacity allocation setting considered in papers [8–10]).

3.4 Allocation with symmetric information

We start by introducing the structure of the three allocation functions of interest and checking whether they are Pareto optimal. We build on the allocation rules for capacity rationing proposed by Cachon and Lariviere (1999, 199b) [8, 10] and we modify them so that $\sum_{i=1}^{n} \alpha_i = Q$. Please note that in the context under consideration a retailer may be allocated a quantity higher than his final demand, if total demand is lower than the total quantity held centrally.

Proportional allocation

$$\alpha_i(\mathbf{d}) = \frac{d_i}{\sum_{i=1}^{n} d_i} Q \tag{3.2}$$

Under symmetric information, it is trivial to show that proportional allocation is efficient and therefore Pareto optimal.

Linear allocation Index retailers in decreasing order of their demand, i.e. $\{d_1 \geq d_2 \geq ... \geq d_n\}$.

Case 1: $\sum_{i=1}^{n} d_i \geq Q$, retailer i is allocated $\alpha_i(\mathbf{d}, \tilde{n})$, where

$$\alpha_i(\mathbf{d}, \tilde{n}) = \begin{cases} d_i - \frac{1}{\tilde{n}}(\sum_{j}^{\tilde{n}} d_j - Q) & \text{for} \quad i \leq \tilde{n} \\ 0 & \text{for} \quad i > \tilde{n} \end{cases} \tag{3.3}$$

and $\tilde{n}$ is the largest integer such that $\alpha_{\tilde{n}}(\mathbf{d}, \tilde{n}) > 0$.

Case 2: $\sum_{i=1}^{n} d_i < Q$, retailer i is allocated $\alpha_i(\mathbf{d})$, where

$$\alpha_i(\mathbf{d}) = d_i + \frac{1}{n}(Q - \sum_{i=1}^{n} d_i) \tag{3.4}$$

Linear is also an efficient allocation because if $\sum_{i=1}^{n} d_i \geq Q$, then $\alpha_i \leq d_i \; \forall \; i$. Similarly, when $\sum_{i=1}^{n} d_i < Q$, it is guaranteed that $\alpha_i \geq d_i \; \forall \; i$. Therefore wastage is excluded in both cases. Although this is an allocation mechanism that maximizes the sum of retailer profits, when there is inventory shortage it may assign zero inventory to retailers with low demand [8]. Consequently, implementing this allocation rule by itself may not satisfy the individual rationality constraints of the retailers and hence may not encourage inventory centralization [31].

Uniform allocation Index retailers in decreasing order of their demand, i.e. $\{d_1 \geq d_2 \geq ... \geq d_n\}$.

Case 1: $\sum_{i=1}^{n} d_i \geq Q$, retailer i is allocated $\alpha_i(\mathbf{d}, \hat{n})$, where

$$\alpha_i(\mathbf{d}, \hat{n}) = \begin{cases} \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} d_j) & \text{for} \quad i \leq \hat{n} \\ d_i & \text{for} \quad i > \hat{n} \end{cases} \tag{3.5}$$

and $\hat{n}$ is the largest integer such that $\alpha_{\hat{n}}(\mathbf{d}, \hat{n}) < d_{\hat{n}}$.

Case 2: $\sum_{i=1}^{n} d_i < Q$, retailer i is allocated $\alpha_i(\mathbf{d}, \hat{n})$, where

$$\alpha_i(\mathbf{d}, \hat{n}) = \begin{cases} d_i & \text{for} \quad i < \hat{n} \\ \frac{1}{n-\hat{n}+1}(Q - \sum_{j=1}^{\hat{n}-1} d_j) & \text{for} \quad i \geq \hat{n} \end{cases} \tag{3.6}$$

and $\hat{n}$ is the smallest integer such that $\alpha_{\hat{n}}(\mathbf{d}, \hat{n}) > d_{\hat{n}}$.

Again, as in linear allocation, the uniform allocation rule guarantees that when $\sum_{i=1}^{n} d_i \geq Q$, $\alpha_i \leq d_i \; \forall \; i$ and when $\sum_{i=1}^{n} d_i < Q$, $\alpha_i \geq d_i \; \forall \; i$. Therefore it is efficient and Pareto optimal. Uniform allocation favors retailers with low demand in the case of inventory shortage, and retailers with high demand in the case of inventory surplus.

These allocation mechanisms are based solely on realized demands and not on retailers' initial orders q_i^a. Hence, they do not allow for comparisons between the profit of a retailer as part of the pooling coalition and what he could have earned in a decentralized system, assuming that he would have ordered q_i^a. The mechanisms also do not provide any minimum inventory or profit guarantee to individual retailer.

We are interested in identifying an allocation mechanism that provides a *guaranteed* and a *maximum* allocation. A guaranteed allocation is the amount of inventory that the retailer is assured to receive if he wants it. Maximum allocation is the retailer's largest possible allocation given his demand. A mechanism that satisfies condition (3.7) provides such guarantees and gives a retailer a greater control over his own destiny.

$$\min(d_i, q_i^a) \leq \alpha_i(d_i, q_i^a, Q) \leq \max(d_i, q_i^a) \;\; \forall i \;\; \& \quad \sum_{i=1}^{n} \alpha_i = Q \tag{3.7}$$

An allocation mechanism that satisfies condition (3.7) guarantees retailer i a quantity at least equal to his initial order and up to his demand, when his realized demand is higher than his initial order. When his local demand is lower than his order quantity, retailer's demand is guaranteed to be satisfied but he may be allocated and charged a quantity less than his initial order. An allocation mechanism satisfying these properties is Pareto efficient because it excludes wastage and eliminates all profitable trade among retailers (total profit maximizing).

An allocation mechanism that satisfies (3.7) also guarantees a retailer a profit larger or equal to what he would have earned under a decentralized system.

Proposition 3.3: Each retailer's profit when his is part of a pooling coalition that sets $Q = \sum q_i^a$ and the CP allocates inventory according to a mechanism that satisfies condition (3.7) is larger or equal to his profit under a pure decentralized system.

In Proposition 3.3 we assume that in a decentralized system (separate newsvendors) each retailer will set inventory q_i^a. Please note that we put no restrictions on how q_i^a is calculated. We assume though that it is the same and it does not depend on the allocation function. It follows immediately from Proposition 3.3. that total profits as well will be larger or equal to sum of individual profits in a decentralized setting. This is due to the benefit of excess demand and stock rebalancing opportunity after demands are realized (unilateral change of inventory ownership).

None of the allocation mechanisms considered so far satisfies condition (3.7) for arbitrary demand realizations. One possible way to implement an allocation satisfying the above inequality is to initially charge each retailer $c \cdot q_i^a$ to cover the quantity he ordered and then allow "allocation changes" after individual demands are realized. Each additional inventory unit above q_i^a is received at a cost c and sent for a revenue of c as well. In this case the same dynamics arise as in a decentralized inventory system where transshipments are allowed and their cost is zero.

We continue by introducing a specific allocation rule, the *modified uniform* allocation, which satisfies condition (3.7). This rule uses each retailer's initial order q_i^a as a starting point to determine his allocation. In case of shortage, shortage is divided equally among retailers that have $d_i > q_i^a$ as long as their resulting allocation is above q_i^a. In case of surplus, additional units are equally divided among those that have $d_i < q_i^a$ as long as the resulting allocation is lower than q_i^a. By construction, this allocation rule is efficient and Pareto optimal. We formally define the modified uniform allocation rule as follows:

Modified uniform allocation Define $g_i = d_i - q_i^a$ and index retailers in decreasing order of the difference between their final and their initial order, i.e. $\{g_1 \geq g_2 \geq ... \geq g_n\}$.

Case 1: $\sum_{i=1}^n d_i \geq Q$, retailer i is allocated $\alpha_i(\mathbf{q}^a, \mathbf{d}, \hat{n})$, where

$$\alpha_i(\mathbf{q}^a, \mathbf{d}, \hat{n}) = \begin{cases} q_i^a + \frac{1}{\hat{n}}(Q - \sum_{j=1}^{\hat{n}} q_j^a - \sum_{j=\hat{n}+1}^{n} d_j) & \text{for} \quad i \leq \hat{n} \\ d_i & \text{for} \quad i > \hat{n} \end{cases} \tag{3.8}$$

and $\hat{n}$ is the largest integer such that $\alpha_{\hat{n}}(\mathbf{q}^a, \mathbf{d}, \hat{n}) < d_{\hat{n}}$.

Case 2: $\sum_{i=1}^n d_i < Q$, retailer i is allocated $\alpha_i(\mathbf{q}^a, \mathbf{d}, \hat{n})$, where

$$\alpha_i(\mathbf{q}^a, \mathbf{d}, \hat{n}) = \begin{cases} d_i & \text{for} \quad i < \hat{n} \\ q_i^a - \frac{1}{n-\hat{n}+1}(\sum_{j=1}^{\hat{n}-1} d_j + \sum_{j=\hat{n}}^{n} q_j^a - Q) & \text{for} \quad i \geq \hat{n} \end{cases} \tag{3.9}$$

and $\hat{n}$ is the smallest integer such that $\alpha_{\hat{n}}(\mathbf{q}^a, \mathbf{d}, \hat{n}) > d_{\hat{n}}$.

3.5 Allocation under asymmetric information

In the asymmetric information case, the central planner determines the allocation of total quantity without observing local demand at each retailer's location. Instead, each retailer, after observing his local demand d_i, places a final order q_i^b to the central warehouse. The main question of this section is: will retailers truthfully report their realized demands to the central planner ($q_i^b = d_i \ \forall i$)? Also, how do the various allocation mechanisms influence the report?

We consider the allocation rules introduced in the previous section, modified to reflect the case where final demands are not known to the central planner. Under asymmetric information, the structure of these allocation mechanisms remains unchanged but they become a function of $\mathbf{q}^b$ instead of $\mathbf{d}$ (i.e. d_i is replaced by $q_i^b \ \forall i$).

All allocation mechanisms studied in 3.4 are Pareto optimal under common knowledge, but the question remains if these allocation mechanisms are inducing retailers to truthfully report their realized demands. To answer this question, we use the concept of dominant strategy equilibrium. We define $q_i^b(d_i)$ to be a function mapping from $[\underline{d_i}, \overline{d_i}]$ to $[0, Q]$. This function defines a strategy for player i, dictating an order for each possible demand realization. Similarly, q_{-i}^b denotes the vector of orders submitted by all retailers but retailer i. The function $q^{b*}(d) = \{q_1^{b*}(d_1), q_2^{b*}(d_2), ..., q_n^{b*}(d_n)\}$ forms a dominant equilibrium for all i and $\mathbf{d}$, *if and only if*:

$$\pi_i(\alpha_i(q_i^{b*}(d_i), q_{-i}^b(d_{-i}))) \geq \pi_i(\alpha_i(q_i^b(d_i), q_{-i}^b(d_{-i}))) \ \forall \ q_i^b, q_{-i}^b \in [0, Q] \tag{3.10}$$

In a dominant strategy equilibrium, each retailer has a strategy that maximizes his profit regardless of the strategies of the other retailers. We are interested in strategies where retailers order their optimal quantities, their true needs, in a dominant equilibrium ($q_i^{b*}(d_i) = d_i$). Under a proportional or linear allocation rule, a retailer has an incentive to inflate his final order above his demand when he thinks there will be inventory shortage and order less than his realized demand when he believes there will be inventory surplus. This results from the observation that under these two allocation mechanisms (a) no retailer receives exactly what he orders, unless $\sum_{i=1}^{n} q_i^b = Q$, and (b) each retailer can influence his allocation to his benefit by modifying his order above or below his demand. On the other hand, uniform allocation is truth inducing: all retailers are always better off placing a final order equal to their realized demand. More formally, we have the following results.

Proposition 3.4: Truthfully reporting their realized demands is not a dominant strategy equilibrium for retailers under proportional or linear allocation mechanisms.

Proposition 3.5: Truthfully reporting their realized demands is a dominant strategy equilibrium for retailers under uniform allocation mechanism.

The intuition behind this result is that when the uniform allocation rule is employed and there is inventory shortage, a retailer can increase his allocation by untruthfully placing an order above his demand only when he belongs to those that receive an allocation equal to their final order. Thus, by placing a final order above his realized demand, the retailer will receive additional units that he cannot sell. The retailers that receive an allocation lower than their final order cannot increase their allocation by ordering more. Similarly, when there is inventory surplus in the system, a retailer can reduce his allocation by ordering less only when he is among the ones who get a quantity equal to their final order. In this case, placing an order lower than the realized demand results in lost sales and hence lost profits for the retailer.

Even if the uniform allocation rule incentivizes retailers to truthfully place a final order equal to their realized demands, it does not guarantee that each retailer gains a profit higher than what he would have gained under a pure decentralized system (i.e. inequality (3.7) may not hold). We show that the modified uniform allocation rule, introduced in section 3.4, not only assures condition (3.7) holds for each retailer and any demand realization, but it is also truth-inducing.

Proposition 3.6: The modified uniform allocation rule is a truth revealing Pareto optimal allocation mechanism that guarantees each retailer a profit larger or equal to his profit under a decentralized system.

Table 3.1 provides a summary of the comparison of the various allocation mechanisms along the dimensions of Pareto optimality (under common knowledge), truth-inducing property and inventory / profit guarantees.

	Proportional	Linear	Uniform	Modified Unifrom
Pareto optimal (under common knowledge)	Yes	Yes	Yes	Yes
Truth Inducing	No	No	Yes	Yes
Quantity & Profit Guarantees	No	No	No	Yes

Table 3.1: Comparison of various allocation mechanisms

In this thesis we assume that the coalition is already formed. If this is not the case, only the modified uniform allocation guarantees that each retailer is not worse off as part of the coalition, for any demand realization. Table 3.1 results hold under the assumption that the inventory level is exogenous, an assumption relaxed in chapters 4 and 5. Therefore, the proportional allocation rule studied in chapter 4 may lead to different results as total inventory becomes an endogenous decision and allocation is based on common knowledge of demand realizations.

3.6 Concluding remarks

Acknowledging that the issue of allocation is present and may play an important role in the dynamics of an inventory pooling coalition formed by multiple retailers, in this chapter we studied the impact of three well-known allocation mechanisms on retailers' strategic behavior and resulting profits. We started by identifying the properties of Pareto optimal allocation mechanisms in the setting under consideration and showing that any allocation mechanism that excludes wastage is Pareto optimal. We continued by showing that when inventory allocation is based on realized demands (when assuming that these are common knowledge) all three allocation schemes we are considering (i.e. proportional, linear, uniform) are Pareto efficient. On the other hand, when the CP bases inventory allocation on final orders placed by the retailers (asymmetric information about local demands), only the uniform and the introduced modified uniform mechanism induce truth-telling and result in Pareto optimal outcomes. Under both these cases, retailers truthfully reporting their realized demands (i.e. place a final order equal to local demand) is a dominant strategy equilibrium. The latter allocation mechanism has the additional attractive property of guaranteeing each retailer and for any demand realization a profit larger or equal to what he would have earned in a decentralized system.

In this chapter, information sharing relates to reporting local realized demands from the retailers to the CP when the latter allocates total inventory. In the following chapters, we focus on demand forecast information sharing between retailers and the CP when the CP sets also the total inventory level. In that case, the CP solicits private demand information from the various markets to inform the inventory decision. When they report their forecasts, retailers consider the impact on the inventory level decision.

Chapter 4

Forecast Information Sharing and the Order Quantity Decision: Impact of Inventory Competition and Market Uncertainty

4.1 Introduction

When we consider the case of multiple retailers that form a pooling coalition and delegate inventory management to a central planner (CP), reliability of demand forecast sharing becomes crucial. The present chapter analyzes the forecast-sharing game between multiple retailers and a central planner who sets common inventory under information asymmetry. In particular, early demand signals are known by retailers but the CP needs this information to make appropriate inventory decisions for the coalition.

Reliable demand information sharing in a supply chain depends on parties' incentives as well as on behavioral factors such as trust and trustworthiness between supply chain parties [42]. For example, in a typical supplier-retailer setting, where a supplier solicits private forecast information from a retailer to set his capacity, standard game theory predicts that parties do not cooperate and the only equilibrium is uninformative. However, recent research reveals that in controlled laboratory experiments parties cooperate even in the absence of reputation-building mechanisms and complex contracts (e.g., [42]). The underlying reason for cooperation seems to be supplier's trust in retailer and the retailer's associated trustworthiness.

In the case of multiple retailers trading the same product, an additional level of complexity is added due to their strategic interactions. Retailers compete for common inventory and this may affect their incentives when sharing demand information. In such a case, the free-rider problem may arise. The alignment of system and individual incentives becomes important as well as the ability of the CP to induce or detect truthfulness in forecast sharing. For example, observing multiple forecasts may increase the CP's ability to detect truthful information sharing[1] but on the other hand the increased competition among retailers may distort incentives or harm trust.

The same dynamics arise in the case of a single company that owns multiple stores and decides to hold inventory centrally, under the assumption that each store is a separate business unit that strives to maximize local profit. Each branch has private information regarding local market factors that may or may not transmit truthfully to the central decision maker who sets and controls the common inventory.

The research focus of this chapter is on whether reliable forecast information sharing occurs when multiple retailers form an inventory pooling coalition managed by a benevolent CP who (a) sets common inventory level before demand is realized and (b) allocates inventory to the retailers after demand uncertainty is resolved in proportion to their realized demands. As in the previous chapter, we consider the case of a benevolent CP in the sense that she is interested in maximizing total system profits (sum of retailer profits).

[1]In our work we do not model this situation, the CP cannot make inferences whether retailers are lying or not.

Furthermore, we want to study what the impact of the size of the pooling coalition and the market uncertainty is on the level of trust and trustworthiness of supply chain parties. We define trust as the CP's willingness to rely on retailers' forecasts to determine total inventory. Trustworthiness, similarly to Ozer et al. (2011) [42], is measured by the difference between a retailer's real forecast and the one he reports to the CP.

To be more specific, we are interested in answering the following questions: (a) Do retailers transmit truthfully their demand forecasts to the CP in such a setting where they strategically interact and compete for common inventory? (b) Does the CP trust the retailers' shared demand information? (c) Does retailers' truthfulness depend on their forecast accuracy? (d) Does the number of retailers forming the coalition have an impact on their trustworthiness and the CP's level of trust? To answer these research questions, we first study players' incentives and the result of their strategic interaction using game theoretical models. After developing a series of analytical results, we develop and execute a series of laboratory experiments to test the theories implied by these results.

The rest of this chapter is organized as follows: first we analyze the forecast communication game with one-time interaction to obtain the standard model prediction. We proceed by doing an extensive numerical analysis to gain some insights on the analytical results. Next, we present four hypotheses in forecast information sharing and cooperation, established based on the analytical results of the game theoretic model and on the existing literature about trust, trustworthiness and human behavior biases. We test these hypotheses in a controlled laboratory environment. The last section describes the experimental procedure, analysis and findings.

4.2 The analytical model

4.2.1 Setting and timing of events

We consider, again, a setting where a group of n retailers each face uncertain demand (newsvendor type problem) and have private information (receive a signal) about it due to their proximity to the market. We denote by $\tilde{d}_i$ the demand that retailer i faces, after he receives his signal θ_i, a random variable with mean $\mu_i + \theta_i$, cumulative distribution function $L_i(\cdot)$, probability density function $l_i(\cdot)$ and support $[\mu_i + \theta_i + \underline{\epsilon_i}, \mu_i + \theta_i + \overline{\epsilon_i}]$, and by D_i the random variable representing total demand from the point of view of retailer i after he receives θ_i.

Each retailer reports his signal (maybe untruthfully) to the CP ($\hat{\theta}_i$). The CP makes a single ordering decision (Q) for the whole coalition after she solicits demand information from each retailer and before demand is realized. She allocates the total inventory quantity to retailers after demand uncertainty is resolved, in proportion to their realized demands. In contrast to the previous chapter, demand realizations at the end of the period are common knowledge to all players. Let d_i be the realized demand at location i and α_i denote the allocation retailer i gets. Then, we define *proportional allocation* as follows: $\alpha_i(\mathbf{d}) = \frac{d_i}{\sum_{i=1}^{n} d_i} Q$.

The proportional allocation rule is widely used in practice and easy to enforce, as it is perhaps the most intuitive scheme [8]. Furthermore, it has several attractive properties. First, as we showed in chapter 3, when realized demands become common knowledge before allocation, proportional rule is a Pareto optimal mechanism from the system's perspective. It is an efficient mechanism, in the sense that it excludes wastage (there are no unsold units when total demand is equal to or higher than Q and there is no locally unsatisfied demand when total demand is lower than Q). Hence, it maximizes total system profit. Second, because total inventory is assigned to retailers (even when inventory exceeds total demand), the profit function for the retailer is not monotonically increasing in Q. Third, the proportional rule provides the same service level to all retailers, a property expected in a system with identical retailers. Last, the proportional allocation rule is chosen for analytical tractability.

The timing of the events, described schematically in Figure 4-1, is as follows:

0. Before the period begins, the size of the pooling coalition and the allocation rule are announced.

1. At the beginning of the period each retailer observes his (and only his) private demand signal (θ_i).

2. Each retailer sends his forecast (reports a demand signal $\hat{\theta}_i$) to the CP.

3. The CP sets the inventory level of the system (Q) and this quantity becomes common knowledge to all players.

4. Local demands (d_i's) are realized and revealed to all players, the allocation of inventory to retailers is done according to the proportional rule and profits are calculated.

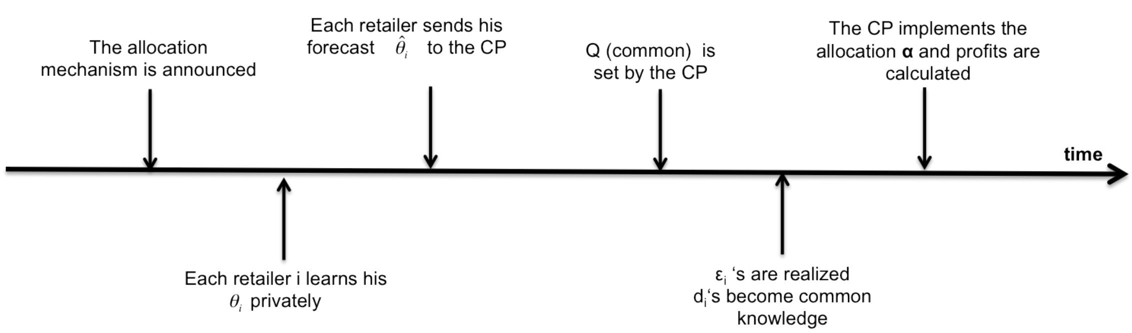

Figure 4-1: The timing of events in the information sharing game with common inventory

Compared to Figure 3-1, the timing of events is very similar. The main differences are: (a) retailers send their forecasts to the CP and not their soft orders and (b) retailers do not place a final order to the central warehouse before allocation is calculated. In terms of decision rights, in the case described in Figure 3-1, the CP just aggregates soft orders placed by the retailers to set the quantity (she has no decision rights on the inventory level).

We first study the information sharing game when there is a single retailer and a CP ($n = 1$) and, second, when multiple retailers form a coalition ($n > 1$). In the case of multiple retailers, we present both the case where the coalition has an automated ordering system in place and the case where there exists a CP with decision rights on setting inventory. We also study the special cases of infinitely many retailers and the case where demand uncertainty is resolved for each retailer before information transmission.

4.2.2 The information sharing game for $n = 1$

We start our analysis with the simplest setting of one retailer to gain intuition on the dynamics between a retailer and the CP when there is no competition for inventory. In this case, there is only one retailer who sends his signal to the CP who in turn decides the quantity to be held for the retailer. For given Q and known θ_i, the retailer's and the CP's expected profits are given by:

$$\pi_i(Q, \theta_i) = \mathbb{E}_{\epsilon_i}[p \min[(\mu_i + \theta_i + \epsilon_i), \alpha_i(Q)] - c\alpha_i(Q)] \tag{4.1}$$

$$\Pi(Q, \theta_i) = \mathbb{E}_{\epsilon_i}[p \min[(\mu_i + \theta_i + \epsilon_i), Q]] - cQ \tag{4.2}$$

The allocation of retailer i after demand is realized reduces to $\alpha_i(Q) = Q$ and hence $\pi_i(Q, \theta_i) = \Pi(Q, \theta_i)$. If the CP knew θ_i, she would maximize her expected profit by setting inventory $Q(\theta_i) = \mu_i + \theta_i + G_i^{-1}(\frac{p-c}{p})$. This is the quantity that maximizes retailer's expected profit, when he has received θ_i. Thus, in this interaction, the retailer has no incentive to distort the report of his forecast, and the CP has no reason not to consider the reported forecast as credible.

Observation 4.1: *In the case of a single retailer, agents' interests coincide and a truthful information sharing equilibrium exists.*

This setting is different from the supplier-manufacturer (or retailer) case considered in Ozer et al. (2011) [42] because the retailer's profit function is not monotonically increasing in the CP's inventory choice Q. Even if he incurs no direct cost by reporting θ_i, in the case of just one retailer his forecast indirectly works as an enforceable order. The CP allocates the retailer the total quantity she orders no matter if it is higher than retailer's demand. The retailer assumes all demand uncertainty risk while the CP has no individual profit margin; she is benevolent instead, trying to maximize retailer's profit.

Please note that this equilibrium is not unique; other equilibria exist in such a setting which vary based on the CP's belief concerning how $\hat{\theta}_i$ and θ_i are related. For example, for any δ_i, the retailer reporting $\hat{\theta}_i = \theta_i + \delta_i$ and the CP setting $Q(\hat{\theta}_i) = \mu_i + \hat{\theta}_i - \delta_i + G_i^{-1}(\frac{p-c}{p})$ constitute a BNE. In such a case the retailer is partially trustworthy and the CP is partially trusting. The resulting equilibrium though is totally informative in the sense that the CP can infer from the forecast reported the true value of θ_i.

4.2.3 The information sharing and allocation game for $n > 1$

In the case of multiple retailers, the issue of allocation becomes important. The ultimate allocation that each retailer gets is a function of the realized demands and the total inventory set by the CP: $\alpha_i(\mathbf{d}, Q) = \frac{d_i}{\sum_{i=1}^{n} d_i} Q$. We note that there is no dedicated inventory to each retailer and no quantity commitment based on his forecast. A retailer's allocation is a (random)[2] fraction of the inventory held centrally. Hence, a retailer may influence the allocation he expects to get after the demands are realized only through influencing the total quantity that is held centrally. In other words, when retailer i reports a signal $\hat{\theta}_i$, his allocation depends on $\hat{\theta}_i$ only through Q that may depend on the signal sent by retailer i. For given Q and θ_i's, retailer's i and

[2]At the time of forecast information transmission allocation is a random fraction of the total quantity because final demands have not been yet realized.

CP's expected profits are:

$$\pi_i(Q,\theta_i) = \mathbb{E}_{\epsilon,\theta_{j,j\neq i}}[p\min[(\mu_i+\theta_i+\epsilon_i),\bar{\alpha}_i(\mathbf{d},Q)] - c\bar{\alpha}_i(\mathbf{d},Q)] \tag{4.3}$$

$$\Pi(Q,\boldsymbol{\theta}) = p\mathbb{E}_{\boldsymbol{\epsilon}}[\min[\sum_{i=1}^{n}(\mu_i+\theta_i+\epsilon_i),Q]] - cQ \tag{4.4}$$

where $\bar{\alpha}_i(\mathbf{d},Q) = \frac{\tilde{d}_i}{D_i}Q$ and $\boldsymbol{\theta} = [\theta_1,\theta_2,...,\theta_n]$.

Common knowledge We begin by studying the common knowledge case where all players have the same information about local and total demand (i.e. $\boldsymbol{\theta}$ is common knowledge to all players). This case serves as a benchmark and a way to better understand the incentives of the players who have private information in the general setting.

When the CP knows the demand signals at each retail location, she faces a newsvendor problem with total demand equal to the convolution of all individual retailers' demand distributions updated according to the realized demand signals. She maximizes Equation (4.4) by setting inventory as:

$$Q_f^{CP}(\boldsymbol{\theta}) = \sum_{i=1}^{n}(\mu_i+\theta_i) + (G_1\circ G_2...\circ G_n)^{-1}\left(\frac{p-c}{p}\right) \tag{4.5}$$

We then study how the quantity given by equation (4.5) compares to the quantity that maximizes each retailer's profit, when he has the same information about total demand. We denote by Q_f^i the optimal quantity to be held centrally from the point of view of retailer i (the quantity that maximizes retailer's i expected profit) when he has the same information about total demand like the CP who sets Q (retailer knows the realized demand signals in all retail locations).

Lemma 4.1: *When demand signals are common knowledge, the CP will set $Q_f^{CP} = Q_f^i$ iff $\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\sum_{i=1}^n \tilde{d}_i > Q_f^i] = \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\sum_{i=1}^n \tilde{d}_i < Q_f^i]$.*
All proofs are in Appendix B.2.

Lemma 4.1 shows that with the same information about total demand, individual retailers' and CP's incentives may not coincide; the CP sets an inventory level that is different than the optimal inventory from the point of view of a retailer i unless the expected ratio of own demand to total demand at Q_f^i is the same in case of shortage and in case of surplus.

Corollary 4.1: *When* $\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i} | \sum_{i=1}^n \tilde{d}_i > Q_f^i] > \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i} | \sum_{i=1}^n \tilde{d}_i < Q_f^i]$ *then* $Q_f^i > Q_f^{CP}$. *When the reverse is true,* $Q_f^i < Q_f^{CP}$.

Corollary 4.1 states that when the expected ration of own demand to total demand in case of shortage is higher than the expected ratio in case of surplus, the optimal total inventory from retailer's i point of view is higher than the system optimal inventory. If the expected ration of own demand to total demand is higher when there is surplus compared to when there is shortage, retailer i prefers a lower total inventory than the system optimal quantity.

Automated inventory ordering system

Under common knowledge, even if the desired quantity by retailer i may be different than the quantity set by the CP, the retailer does not have the chance to influence CP's decision by misreporting demand information. But, under asymmetric information, due to this discrepancy in incentives, retailers, anticipating how the CP will set common inventory, may have an incentive to misreport to her their private information. In the context of strategic communication and information sharing, in which a better informed sender sends a possibly noisy signal to a receiver, who then takes an action that determines the welfare of both, the amount of information that the sender shares is related to the similarity of agents' interests [16].

We begin with the case of an automated central inventory ordering system that takes as input the forecasts reported by the retailers and automatically calculates and orders the optimal inventory level for the whole coalition, based on the demand information provided. The total quantity ordered is given by the equation $Q^{CP}(\hat{\boldsymbol{\theta}}) =$

$\sum_{i=1}^{n}(\mu_i+\hat{\theta}_i)+(G_1\circ G_2...\circ G_n)^{-1}(\frac{p-c}{p})$ and satisfies the sufficient optimality condition $(p-c)\Pr[\hat{D}>Q^{CP}]=c\Pr[\hat{D}<Q^{CP}]$, where $\hat{D}$ denotes the total demand, updated according to the reported signals. In this case, retailer i has an incentive to distort his report of θ_i, *if and only if* Q^{CP} is different than Q^i that maximizes Equation (4.3). As the next theorem formally states, when such an automated inventory ordering system is in place, we expect $Q^{CP}\neq Q^i\ \forall i$.

Theorem 4.1: *Under an automated inventory ordering system, the optimal quantity from the point of view of retailer i solves the equation*

$$(p-c)\Pr[D_i>Q^i]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i>\bar{\alpha}_i(\mathbf{d},Q^i)]=c\Pr[D_i<Q^i]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i<\bar{\alpha}_i(\mathbf{d},Q^i)]$$

and consequently $Q^{CP}\neq Q^i\ \forall i$, unless $\int\int_{\tilde{d}_i>\bar{\alpha}_i}\bar{r}_i\Gamma(\bar{r}_i,\tilde{d}_i)d\bar{r}_i d\tilde{d}_i=\Pr[\hat{D}>Q^i]\mathbb{E}[\bar{r}_i]$, where $\bar{r}_i=\frac{\tilde{d}_i}{D_i}$ and Γ is the joint distribution of $\bar{r}_i$ and $\tilde{d}_i$. Hence, retailer i may have an incentive to distort his reported signal.

Theorem 4.1 shows that the optimal total quantity for retailer i will differ from the one that the automated system sets, for two reasons. First, each retailer has partial information about total demand (knows only his demand signal) while the automated system uses reported signals about all local markets (which in turn may be distorted). To be more specific, total demand for retailer i is characterized by the random variable $D_i=[\sum_{i=1}^{n}\mu_i+\theta_i]+[\sum_{j=1,j\neq i}^{n}\theta_j+\sum_{i=1}^{n}\epsilon_i]$ where the random part is $\sum_{j=1,j\neq i}^{n}\theta_j+\sum_{i=1}^{n}\epsilon_i$, while the automated system sets inventory based on $\hat{D}=\sum_{i=1}^{n}(\mu_i+\hat{\theta}_i)+\sum_{i=1}^{n}\epsilon_i$, where the random part is $\sum_{i=1}^{n}\epsilon_i$. Thus, for any given Q, retailer i has different belief about the probability of inventory shortage (or surplus) for the coalition than the one the automated system calculates ($\Pr[D_i>Q]\neq\Pr[\hat{D}>Q]$). This is also the case when all the reported signals to the system are true (when $\hat{\theta}_i=\theta_i\ \ \forall i$, $\hat{D}=\sum_{i=1}^{n}(\mu_i+\hat{\theta}_i)+\sum_{i=1}^{n}\epsilon_i$, a different random variable than D_i).

Second, the expected ratio of own demand to total demand may be different in case of shortage and in case of surplus for a given Q (similar to the symmetric information case). When $\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d}, Q^{CP})] \neq \mathbb{E}[\frac{\tilde{d}_i}{D_i}|d_i < \bar{\alpha}_i(\mathbf{d}, Q^{CP})]$, this is another reason why $Q^i \neq Q^{CP}$, incentivizing retailer i to misreport his demand signal to the automated inventory system. The expected ratio of own demand to total demand determines how much the expected allocation to retailer i will increase when the central inventory is increased by one unit. For example, a retailer with $\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d}, Q^{CP})] > \mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d}, Q^{CP})]$ expects to receive more out of one unit of additional central inventory when he needs it (case of shortage) than when he can't sell it (case of surplus).

Next, we study whether a (pure strategy) Bayesian Nash equilibrium exists among retailers when reporting their demand forecasts in this setting. The Bayesian Nash equilibrium is both the natural generalization of the Nash equilibrium to games with incomplete information and a natural extension of the concept of rational-expectations equilibrium to situations where strategic interactions are important [16]. A Bayesian equilibrium requires that each retailer maximizes his profit in expectation, assuming other retailers follow the same Bayesian equilibrium. In this context, a player's type is defined by his demand signal θ_i (private information). Let $\hat{\theta}_i(\theta_i)$ be the reporting strategy of retailer i, a function mapping from $[\underline{\theta_i}, \overline{\theta_i}]$ to $[\underline{\theta_i}, \overline{\theta_i}]$, dictating a reported signal for each possible type. Similarly, $\hat{\theta}_{-i}(\theta_{-i})$ denotes the vector of signals reported by all retailers but retailer i. The functions $\hat{\theta}^*(\theta) = \{\hat{\theta}_1^*(\theta_1), ..., \hat{\theta}_n^*(\theta_n)\}$ form a Bayesian equilibrium if for all i,

$$\hat{\theta}_i^*(\theta_i) \in \arg\max_{\hat{\theta}_i} \int_{\theta_{-i}} \mathbb{E}_\epsilon[P_i(Q(\hat{\theta}_i, \hat{\theta}_{-i}^*(\theta_{-i})), (\theta_i, \theta_{-i})] f(\theta_{-i}|\theta_i) d d\theta_{-i}$$

where f denotes the joint probability function of θ_{-i} and $P_i = p \min[d_i, \alpha_i(\mathbf{d}, Q)] - c\alpha_i(\mathbf{d}, Q)$.

Theorem 4.2: *Under an automated inventory ordering system, a pure strategy Bayesian Nash equilibrium does not exist among retailers when reporting their demand forecasts.*

Central Planner with inventory decision rights

In order to study the information sharing (and allocation) game with multiple retailers and a human CP that makes the ordering decision, we employ the concept of Perfect Bayesian equilibrium. In such an equilibrium, the central planner and the retailers maximize their respective expected profits by responding optimally to each other's strategy, taking into account their actions' implications on the beliefs about the reported signals. Although each retailer's forecast sharing necessarily precedes CP's action in time (setting of total inventory), because the CP observes only the signal (and not the signal reporting rule), retailers' choice of reporting their signal and CP's choice of total inventory are strategically "simultaneous". The CP though, after observing $\hat{\boldsymbol{\theta}}$, updates her beliefs about $\boldsymbol{\theta}$ and bases her choice of inventory on the posterior distribution of $b(\cdot|\hat{\theta})$ over $[\underline{\theta_i}, \overline{\theta_i}]^n$.

We show that a situation where (a) retailers truthfully communicate to the central planner their private demand signal θ_i and (b) the CP considers the reported forecasts as credible, is not a perfect Bayesian equilibrium (PBE).

Theorem 4.3: *Let $\phi(\hat{\theta}_i|\theta_i)$ denote retailer's i reporting strategy given θ_i, $Q(\hat{\theta})$ denote the total quantity to be held determined by the central planner given $\hat{\theta}$ and $b(\theta|\hat{\theta})$ the central planner's posterior belief about θ after observing $\hat{\theta}$.Then,*

- $\phi(\hat{\theta}_i|\theta_i) = \theta_i$
- $Q(\hat{\theta}) = \sum_{i=1}^{n}(\mu_i + \hat{\theta}_i) + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$ *and*
- $b(\theta|\hat{\theta}) = \hat{\theta}$

do not constitute a perfect Bayesian equilibrium. In other words, a perfectly informative bayesian equilibrium of forecast sharing where retailers are fully trustworthy and the CP is fully trusting does not exist.

Theorem 4.3 shows that a completely informative, truthful information sharing PBE among retailers and the central planner does not exist. Moreover, using a similar argument, a perfectly informative equilibrium of forecast information sharing where the CP can infer with certainty from the signal reported the true signal of a retailer cannot exist either. On the other extreme, it is easy to show that a babbling equilibrium exists. In a babbling equilibrium the sender's strategy is independent of his type and the receiver's strategy is independent of signal [51]. In our setting, in this uninformative equilibrium each retailer's report $\hat{\theta}_i$ is independent of θ_i. The CP has no update about θ and determines the optimal quantity based on her prior belief about θ: i.e., the CP sets inventory level $Q = \sum_{i=1}^{n} \mu_i + (F_1 \circ ... \circ F_n \circ G_1 \circ ... \circ G_n)^{-1}(\frac{p-c}{p})$.

Infinitely many retailers

The dynamics that arise in an extremely multitudinous pooling coalition can be approximated by considering the case of infinitely many retailers.

Lemma 4.2: *When $n \to \infty$, if $\sup\{f'^{\otimes(n-1)}(t)|t \geq 0\} \to 0$ as $n \to \infty$, where $f^{\otimes(n-1)}$ is the pdf of $D_{-i} = \sum_{j \neq i} d_j$, then $\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d}, Q)] = \mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d}, Q)] = 0$.*

As the number of retailers forming the coalition approaches infinity and under the condition, loosely speaking, that the pdf of D_{-i} flattens as n increases, the expected ratio of own demand to total demand goes to zero for every retailer, both in case of inventory surplus and in case of shortage. This result implies that a perfectly informative equilibrium of truth telling is sustainable both when the ordering system is automated and when a CP is in charge of setting the common inventory, as the next theorem formally states.

Theorem 4.4: *When $n \to \infty$, a) in the automated inventory system case, $\phi(\hat{\theta}_i|\theta_i) = \theta_i \ \forall i$ form a Bayesian Nash Equilibrium and b) in the case of a CP with decision rights, $\phi(\hat{\theta}_i|\theta_i) = \theta_i$, $Q(\hat{\theta}) = \sum_{i=1}^{n}(\mu_i + \hat{\theta}_i) + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$ and $b(\theta|\hat{\theta}) = \hat{\theta}$, constitute a Perfect Bayesian Nash equilibrium.*

Demand uncertainty is resolved before information transmission

In this section, we consider the special case where demand uncertainty is resolved when retailers receive their private demand signal. In other words, retailers' forecasts are perfectly accurate and demand at each retailer location is $D_i = \mu_i + \theta_i$. This case allows us to study trustworthiness and trusting behavior in demand forecast sharing in a very simple setting with multiple retailers and privately known deterministic demands. Recent research shows that when demand is deterministic and capacity is binding, Nash equilibrium predictions substantially exaggerate retailers' tendency to strategically order more than they need in order to gain a more favorable allocation [13].

In case of deterministic demand, the reasoning behind forming a coalition may not be hedging against demand uncertainty but lower procurement prices due to volume discounts. As before, each retailer, after he receives his demand signal θ_i (private information), inputs to the automated inventory system or sends to the CP a forecast $\hat{\theta}_i$ who in turn sets the common inventory Q. After Q is set, local demands are revealed and the automated system or the CP allocate the common quantity according to the proportional allocation rule. We note that when the common inventory is set, demand uncertainty is resolved but the central system or the CP do not have this information. Local demand information is private knowledge to each retailer who may share it truthfully or not.

Theorem 4.5: *When demand uncertainty is resolved before information transmission then a) in the automated inventory system case, $\phi(\hat{\theta}_i|\theta_i) = \theta_i \;\; \forall i$ form a Bayesian Nash Equilibrium and b) in the case of a CP with decision rights, $\phi(\hat{\theta}_i|\theta_i) = \theta_i$, $Q(\hat{\theta}) = \sum_{i=1}^{n}(\mu_i + \hat{\theta}_i)$ and $b(\theta|\hat{\theta}) = \hat{\theta}$, constitute a Perfect Bayesian Nash equilibrium.*

Theorem 4.5 suggests that retailers truthfully inputing to the automated inventory system their accurate demand forecasts is a sustainable equilibrium among them. Moreover, in the case of a human CP with decision rights, a truthful information sharing PBE exists (retailers are fully trustworthy and the CP is fully trusting) when retailers' demand signals are perfectly accurate. We note that also in this case multiple equilibria exist, under different belief structures; retailers may be partially trustworthy and the CP partially trusting. For example, for any δ_i, $\phi(\hat{\theta}_i|\theta_i) = \theta_i + \delta_i$, $Q(\hat{\theta}) = \sum_{i=1}^{n}(\mu_i + \hat{\theta}_i - \delta_i)$ and $b(\theta_i|\hat{\theta}_i) = \hat{\theta}_i - \delta_i$ constitute a BNE.

	Automated system	**CP with decision rights**
$n = 1$	-	Truth-telling/ trusting is PBE
General Case ($n > 1$, market uncertainty)	Pure strategy BNE does not exist	Truth-telling/ trusting is **not** PBE Babbling equilibrium exists
$n \to \infty$	Truth-telling/ trusting is BNE	Truth-telling/ trusting is PBE
No demand uncertainty when info sharing	Truth-telling/ trusting is BNE	Truth-telling/ trusting is PBE

Table 4.1: Summary of analytical findings of Chapter 4

4.3 Numerical study

To complement our analytical findings, we present a numerical study of how the optimal quantity from the point of view of retailer i differs from the optimal quantity for the whole coalition, under common knowledge and asymmetric information. The misalignment between each retailer's and the system's optimal quantities will determine retailers' incentives to misreport their forecasts (as long as an automated inventory system is in place or the CP is – at least partially – trusting).

We consider the case of two retailers ($i = 1, 2$) and an automated inventory ordering system. Retailers are symmetric in their cost parameters. Each unit received by the retailer costs $c = 1$, and can be sold at unit price, $p = 2$. Demand at the two locations is distributed independently ($\mathrm{Cov}(\theta_i, \theta_j) = 0$, $\mathrm{Cov}(\epsilon_i, \epsilon_j) = 0$ for $i, j = 1, 2$ and $i \neq j$ and $\mathrm{Cov}(\theta_i, \epsilon_j) = 0$, for $i, j = 1, 2$). We compute explicit solutions for the total order quantity that retailer i would prefer. We first consider the case where retailer i knows the signal received by the other retailer and then the case where he knows only its demand distribution.

4.3.1 Common knowledge

Suppose average demand at each location, before the market signal, is 10 units ($\mu_1 = \mu_2 = 10$) and market uncertainty at the two locations is independent and distributed normally, with mean 0 and variance 1 ($\epsilon_i \sim N(0,1)$). The optimal inventory quantity for retailer 1 (Q_f^1) and the difference with the system optimal inventory (Q_f^{CP}) are reported in Table 4.2 for different values of θ_1 and θ_2.

	Q_f^1							$Q_f^1 - Q_f^{CP}$						
	θ_2							θ_2						
θ_1	-3	-2	-1	0	1	2	3	-3	-2	-1	0	1	2	3
-3	14.000	15.010	16.018	17.026	18.032	19.038	20.043	0.000	0.010	0.018	0.026	0.032	0.038	0.043
-2	14.992	16.000	17.008	18.014	19.020	20.025	21.030	-0.008	0.000	0.008	0.014	0.020	0.025	0.030
-1	15.986	16.993	18.000	19.006	20.011	21.016	22.020	-0.014	-0.007	0.000	0.006	0.011	0.016	0.020
0	16.982	17.989	18.995	20.000	21.005	22.009	23.013	-0.018	-0.011	-0.005	0.000	0.005	0.009	0.013
1	17.980	18.986	19.991	20.996	22.000	23.004	24.008	-0.020	-0.015	-0.009	-0.004	0.000	0.004	0.008
2	18.978	19.986	20.988	21.992	22.996	24.000	25.003	-0.022	-0.017	-0.012	-0.008	- 0.004	0.000	0.003
3	19.977	20.982	21.986	22.990	23.994	24.997	26.000	-0.023	-0.018	-0.014	-0.010	- 0.006	- 0.003	0.000

Table 4.2: Optimal inventory for retailer 1 as a function of θ_1 and θ_2 when cr=0.5

As these numbers suggest, when the signal received by retailer 1 is higher than the signal received by retailer 2, retailer's 1 optimal inventory quantity is always smaller than the optimal inventory quantity of the coalition. From corollary 4.1, this implies that for retailer 1 the expected ratio of own demand to total demand given surplus is larger than the expected ratio given a shortage. The reverse is true when the signal received by retailer 1 is smaller than that of retailer 2. As the difference between the two signals increases, so does the difference between the optimal inventory from the point of view of retailer and the optimal inventory for the coalition. In all cases though, the difference is extremely small, ranging from 0 to 0.22% of the corresponding system optimal quantity. When the demand signals received by the two retailers are identical, then $Q_f^1 = Q_f^{CP}$, suggesting that the expected ratio of own demand to total demand is the same in case of shortage and in case of surplus (Lemma 4.1).

We perform the same analysis, when $p = 4$ and $c = 1$, resulting in a critical ratio of 0.75, and for $p = 2$ and $c = 1.5$, resulting in a critical ratio of 0.25. The difference between the optimal quantity for retailer 1 and the system optimal one are summarized in Table 4.3. We observe that the comparison gives the same directional results, independent of the critical ratio considered. We therefore continue our numerical analysis with the asymmetric information case, using cr=0.5 while the calculations for the base case under asymmetric information and critical ratios 0.75 and 0.25 are delegated into Appendix C.

	$\boldsymbol{Q_f^1 - Q_f^{CP}}$ when cr=0.75							$\boldsymbol{Q_f^1 - Q_f^{CP}}$ when cr=0.25						
	θ_2							θ_2						
θ_1	-3	-2	-1	0	1	2	3	-3	-2	-1	0	1	2	3
-3	0.000	0.009	0.017	0.024	0.030	0.036	0.041	0.000	0.010	0.019	0.027	0.034	0.040	0.046
-2	-0.008	8.000	0.007	0.013	0.019	0.024	0.029	-0.009	0.000	0.008	0.015	0.021	0.027	0.032
-1	-0.013	-0.006	0.000	0.006	0.011	0.015	0.019	-0.015	-0.007	0.000	0.006	0.012	0.017	0.021
0	-0.017	-0.011	-0.005	0.000	0.005	0.009	0.013	-0.019	-0.012	-0.006	0.000	0.005	0.010	0.014
1	-0.019	-0.014	-0.009	-0.004	0.000	0.004	0.007	-0.022	-0.015	-0.010	-0.005	0.000	0.004	0.008
2	-0.021	-0.016	-0.011	-0.007	-0.003	0.000	0.003	-0.023	-0.018	-0.013	-0.008	- 0.004	0.000	0.003
3	-0.025	-0.020	-0.016	-0.013	-0.009	-0.006	0.000	-0.024	-0.019	-0.015	-0.011	- 0.007	- 0.003	0.000

Table 4.3: $Q_f^1 - Q_f^{CP}$, as a function of θ_1 and θ_2, for cr=0.75 and cr=0.25

4.3.2 Asymmetric information

We calculate the optimal quantity for retailer 1 (Q^1), as a function of the market signal he receives, when he does not observe the signal received by the other retailer, but he knows only its distribution. In order to study retailer's incentives in the information sharing game, we compute the expected quantity an automated inventory system would set, assuming that $\hat{\theta}_i = \theta_i$, for $i = 1, 2$, and we compare it to Q^1.

We note that $Q^{CP}(\hat{\theta}_1, \hat{\theta}_2) = K + \hat{\theta}_1 + \hat{\theta}_2$, where $K = \mu_1 + \mu_2 + (G_1 \circ G_2)^{-1}(\frac{p-c}{p})$. Retailer 1, does not know the reporting rule of retailer 2 and therefore we cannot calculate $\mathbb{E}_{\hat{\theta}_2}[Q^{CP}]$ as a function only of $\hat{\theta}_1$. But it is reasonable to assume that retailer's 1 belief about retailer's 2 reporting rule does not depend on θ_1 and therefore $\mathbb{E}[\hat{\theta}_2] = \lambda$, where λ is a constant. Hence, $\mathbb{E}[Q^{CP}(\hat{\theta}_1)] = K+\lambda+\hat{\theta}_1$. For our calculations of $Q^1 - \mathbb{E}[Q^{CP}]$, we assume $\lambda = 0$ (or $\hat{\theta}_2 = \theta_2$) and $\hat{\theta}_1 = \theta_1$, but it is straightforward to compute how the difference would change for different beliefs about λ or different reporting rules of retailer 1.

In the base case, the average market size at both locations (before demand signals are realized) is the same ($\mu_1 = \mu_2 = 10$), the demand signals are independent, normally distributed, with mean 0 and variance 1 ($\theta_i \sim N(0, 1)$) and market uncertainty at each location follows as well the standard normal distribution ($\epsilon_i \sim N(0, 1)$).

For the base case, the optimal quantity for retailer 1 and the expected system quantity had the two retailers reported the true signals are presented in Table 4.4. It is interesting to note that in all instances $Q^1 < \mathbb{E}[Q^{CP}]$ and the absolute difference is increasing in the signal received by retailer 1. Retailer's 1 perception about total demand uncertainty is higher because he does not know the demand signal of retailer 2. The CP instead, in these numerical experiments, takes the received signals as given. But since the critical fractile is 0.5, this difference in total demand uncertainty does not influence the optimal total inventory level from the point of view of each player. The fact that retailer 1 always prefers a smaller quantity is in line with the result that "for linear holding/shortage costs the optimal probability of shortage under random yield is no smaller than the probability of shortage under certain yield" [27].

We continue by studying the effect of private information variability, the effect of market uncertainty and that of the magnitude of local demands. Figure 4-2a) shows

θ_1	Q^1	$\mathbb{E}[Q^{CP}]$	$Q^1 - \mathbb{E}[Q^{CP}]$
-3	16.966	17	-0.034
-2	17.958	18	-0.043
-1	18.952	19	-0.048
0	19.949	20	-0.051
1	20.947	21	-0.053
2	21.946	22	-0.054
3	22.946	23	-0.054

Table 4.4: Comparison of the optimal inventory for retailer 1 and the expected inventory in the system under truth-telling, as a function of θ_1

($Q^1 - \mathbb{E}[Q^{CP}]$) for the base case, for high signal variability ($\theta_2 \sim N(0, 3)$) and for very small market uncertainty ($\epsilon_i \sim N(0, 0.5)$) while Figure 4-2b) shows the results when the average market size of retailer 1 doubles ($\mu_1 = 20$) and when instead the average market size of retailer 2 doubles ($\mu_2 = 20$).

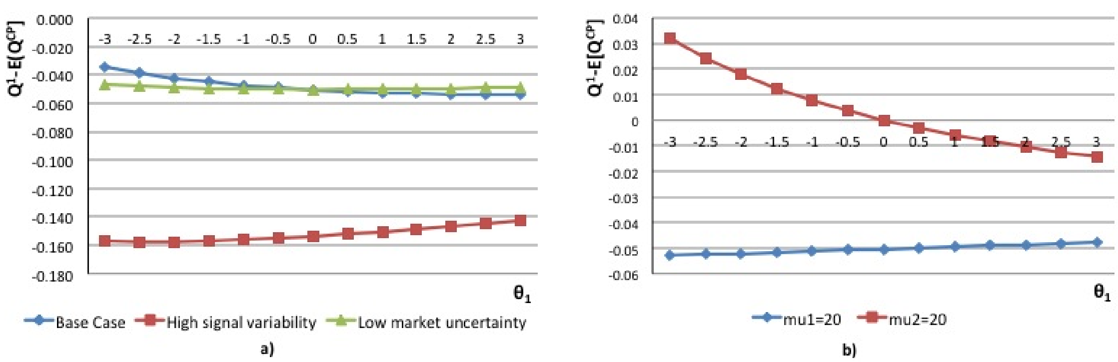

Figure 4-2: ($Q^1 - \mathbb{E}[Q^{CP}]$) as a function of the signal received, for high signal variability, small market uncertainty and asymmetric market sizes

Results suggest that when market signal variability increases, the difference between retailer's and system's preferred inventory level increases in magnitude. On the contrary, market uncertainty does not have a strong effect on the difference of preferred quantities. This is because both the retailer and the central system have the same information about total market uncertainty, while the automated inventory system treats the reported signals as true constants. In addition, results suggest that the relative size of the retailer plays an important role in his incentives. In our example, we observe that when $\mu_1 = 20$, retailer 1 always prefers for total inventory a quantity smaller than the system optimal inventory (independent of θ_1). On the other hand, when his market is relatively small ($\mu_2 = 20$), he may have an incentive to inflate his forecast when his market signal is low and under-report his signal when it is high.

To summarize, numerical analysis suggests that the market size of a retailer compared to the other member of the coalition plays an important role in his preferred total inventory level compared to the system optimal one. A retailer that expects lower demand than the other retailer has a higher expected allocation ratio in case of shortage than in case of surplus and this drives up his preferred inventory level. On the other hand, uncertainty in the allocated quantity, which increases with signal variability, drives the locally optimal inventory level down compared to the system optimal level.

4.4 Hypotheses

When there is only one retailer sharing information with a benevolent CP, game theory provides a definitive prediction of the result of the information sharing game: any equilibrium is totally informative. This result is the opposite to the case of one supplier-one manufacturer, where the only theoretical equilibrium in demand forecast sharing for capacity investment is uninformative [42]. In that case, players' incentives are opposed as the manufacturer always prefers a higher capacity. In the setting we consider, a single retailer's and CP's incentives coincide (Observation 4.1). Thus, truth-telling, in this "cheap talk" setting, is an equilibrium even if it may not be unique. The retailer truthfully communicates his demand forecast to the CP who, in turn, updates accordingly her belief about demand when setting the inventory. Based on this equilibrium, we formulate the following hypothesis:

Hypothesis 1: In the information sharing game between a retailer and a CP, (a) the retailer is fully trustworthy ($\hat{\theta} = \theta$) and (b) the CP is fully trusting ($Q(\hat{\theta}) = \mu + \hat{\theta} + G^{-1}(\frac{p-c}{p})$).

In the case of two or more retailers, they compete for common inventory. Theorem 4.3 predicts that in the forecast communication and inventory allocation game, retailers being fully trustworthy when reporting their demand forecasts and the CP fully trusting them do not constitute a Bayesian Equilibrium. Retailers that want to maximize their pecuniary payoffs have an incentive to distort the information they send and thus a rational CP will not consider their forecasts as credible. But in this strategic information transmission setting, even though there are no exogenous signaling costs when reporting a forecast, there are *endogenous signaling costs* resulting from the fact that the total inventory quantity held centrally by the CP will be allocated to the retailers after local demands are realized. Retailers are held responsible for both system understocking and overstocking costs. If the central inventory quantity is very high, individual expected overstocking costs increase while if the quantity is very small, individual expected under-stocking cost increase. Motivated by these results and the work of Ozer et al. (2011) [42] that suggest that in reality a continuum of trust exists when supply chain parties share forecast information (in contrary to the all-or-nothing view adopted by the extant literature), we formulate the following two hypotheses:

Hypothesis 2: In the information sharing and allocation game, retailers' reports $\hat{\theta}_i$'s are informative about their private forecasts θ_i's. More specifically, $\hat{\theta}_i$ is positively correlated with θ_i (retailers are are partially trustworthy).

Hypothesis 3: The CP relies on $\hat{\theta}_i$'s to determine inventory that is held centrally. More specifically, Q is positively correlated with $\sum_{i=1}^{n} \hat{\theta}_i$ (the CP is partially trusting).

The second factor in the supply chain environment that we study is the impact of market uncertainty. When demand uncertainty is resolved before information transmission, i.e. there is no market uncertainty after retailers receive their demand signals, Theorem 4.5 predicts that the resulting equilibrium is totally informative. Motivated by this prediction, we examine the following hypothesis.

Hypothesis 4: When demand uncertainty is resolved before information sharing (i.e. retailers' private forecasts are perfectly accurate), (a) retailers are fully trustworthy ($\hat{\theta}_i = \theta_i$) and (b) the CP is fully trusting ($Q(\hat{\theta}) = \sum_{i=1}^{n}[\mu_i + \hat{\theta}_i]$).

To summarize, theory predicts reliable information transmission only in the case of a single retailer and in the case of known demand by retailers when communication takes place. On the other hand, when more than one retailers compete for common inventory and demand is uncertain, truth-telling and trust do not form an equilibrium. We therefore hypothesize, that competition for common inventory and demand uncertainty harm reliable information sharing. To be specific, compared to the other two cases, we expect that the signal sent is less informative (lower correlation between $\hat{\theta}_i$ and θ_i) and the CP relies less on the information received to set inventory (lower correlation between $\sum_{i=1}^{n} \hat{\theta}_i$ and Q).

4.5 Experiments

4.5.1 Experimental design and procedures

We conducted a series of human-subject controlled laboratory experiments[3] to investigate the aforementioned hypotheses. We conducted three treatments / experiments

[3] All experiments were conducted at the Social and Behavioral Sciences Laboratory, at the University of Minnesota.

as summarized in Table 4.5. Each treatment is labeled as R_iD_j, where $i = 1, 2$ denotes the number of retailers and $j \in \{U, K\}$ stands for "uncertain" and "known" demand, respectively. In case R_1D_U, only one retailer interacts with a CP. In cases R_2D_U and R_2D_K, the number of retailers increases to two and the issue of inventory allocation becomes relevant. Under case R_2D_K, demand uncertainty is resolved before demand information transmission. Under two treatments retailers compete for common inventory (cases R_2D_U and R_2D_K) and under two treatments there is demand uncertainty when information is transmitted (cases R_1D_U and R_2D_U).

All other supply chain parameters are kept constant across the different treatments. We fix average demand, signal variability and market uncertainty (for the cases that is relevant).[4] For the demand parameters used in our experiments, we calculate a retailer's optimal inventory level, for various demand signal realizations, and we compare it to the system optimal had the CP known the real signals (i.e., the expected inventory level under the assumption that the other retailer reports his true signal and the CP trusts the reported information). The results, analogous to Table 4.4 but for discrete uniform demand distributions, are presented in Appendix C. We note that in all cases, the difference between the quantity that retailer 1 prefers and the system optimal inventory level is less than 1%.

We use revenue and cost parameters that result in critical ratio of 0.5. In that case, the optimal order quantity for the CP is equal to the mean of the demand. By doing so, we avoid experimental results to be influenced by the "pull-to-center" effect when setting the inventory quantity; the phenomenon of systematically ordering too little when the cost of underage is high and ordering too much when the cost of overage is high. This phenomenon may be explained by some well-known decision biases, e.g., anchoring and insufficient adjustment, observation bias, reference-dependent preferences and is well documented in the literature in single newsvendor [5, 7, 47] and

[4]For average demand and signal distribution we adopted the values used for the information sharing experiments reported in Ozer et al (2011). Regarding market uncertainty, we use the average of the aforementioned experiments.

multilocation newsvendor models [28]. The number of participants in each treatment was decided based on the supply chain group synthesis in each case (one vs. two retailers), taking into account the probability that a specific group will play more than once and the average number of rounds a given group is expected to play.

Treatment	No of retailers (n)	Demand Uncertainty	No. of participants	No. of rounds
Case R_1D_U	n=1	Yes	10	30
Case R_2D_U	n=2	Yes	15	30
Case R_2D_K	n=2	No	12	30

Notes. In all treatments, $\mu_i = 250$, $\theta_i \sim U[-150, 150]$, $\epsilon_i \sim U[-50, 50]$ (all discrete), $p = 2$ and $c = 1$.

Table 4.5: Experimental Design

We used a between-subjects design; i.e., each treatment involves an independent group of participants. We recruited students of the University of Minnesota for the experiments, through the Carlson School of Management Subject Pool [5]. Participants in each treatment were randomly assigned the role of retailer or central planner; this role assignment remained unchanged for all rounds. At the beginning of each round, all participants were randomly and anonymously assigned to a supply chain group. Players were informed that they would not be assigned to the same supply chain group in consecutive rounds. The experimenter also stressed out that rounds are independent, both in terms of group assignment and demand realizations, to avoid reputation effects and demand chasing.

Each subject received a detailed sheet of experimental instructions and a short summary sheet with the important info about supply chain parameters, the sequence

[5] Carlson School of Management (CSOM) Subject Pool was launched on September 26, 2008 and consists of a database of individuals who have previously agreed to be part of the pool. CSOM Subject Pool members are able to go online, view studies available for their participation, and sign up to participate in any study they are interested in, provided they meet the filtering criteria set by the researcher.

of events and a reminder about how profits are calculated. The detailed instructions to the CP and to the retailers for case R_2D_U can be found in Appendix D. Each participant was informed about the task performed by the other role, including the information available and profit objective. After the instructions were read, subjects were allowed to ask questions and then directed to the computers' room where the experiment was implemented. The experiment was programmed and conducted with the software z-Tree [20]. Participants were not allowed to talk to each other from the time they entered the laboratory until the time they left. During the experiment, participants interacted with each other only through computer terminals and did not know the identity of the person with whom they were playing.

Each treatment consisted of 35 rounds. The first 5 rounds served as a trial period and did not affect the final profit; they served as training for the participants to better understand the dynamics of the game. Participants were not informed how many rounds they would play after this trial period[6] to avoid end-of-treatment effects. Participants played the information sharing inventory game specified in section 4.2.1. Briefly, each round consisted of 3 periods. In period 1, each retailer i observed his private information (the exact value of $\mu + \theta_i$) and was asked to submit a report to the CP. In period 2, the CP, after observing the report(s) sent by the retailer(s), decided on the (total) order quantity to place. The software indicated to the CP the quantity that maximized her expected profit, if the reports sent are considered to be true. After the decisions were made, in period 3, demands were revealed and profits were calculated. At the end of each round, participants observed realized demand(s) (and total demand in case of two retailers), (total) order quantity, own inventory allocation and own profits from the round. We provide sample snapshots of the CP and retailer screens in Appendix E.

In the experiment, the demand values generated (θ_i and ϵ_i) varied between the two retailers in the same supply chain group (for case R_2D_U and case R_2D_K) and across

[6]Participants were told they are going to make decisions for up to 40 independent rounds (5 trial and up to 35 game rounds).

rounds. However, they remained the same across treatments whenever that was applicable (e.g. θ_1 in round 1 was the same across all cases, for all groups, while ϵ_1 in round 1 was the same between case R_1D_U and case R_2D_U, for all groups). This design feature is particularly useful because it allows to compare decisions across treatments while controlling for individual specific effects and to compare profits across treatments (in some cases), while controlling for the impact of demand realizations on actual profits.

At the beginning of each treatment, participants were asked for basic demographic information, including major, class year, specific courses and level of experience in supply chain management. At the end of each treatment, and after all rounds were played, participants were required to complete a post-game survey. The survey contained questions asking to comment on the choices they made during the study (report of private information or order quantity placement), on how their strategy changed throughout the session and on their degree of trust towards the other players in their supply chain. The level of trust they placed in the other supply chain group members was measured, for each role, on a Likert-type scale from 1 to 5, where 1 denotes "No trust at all" and 5 represents "Absolute trust". Finally, every participant received payment proportional to the total experimental dollars he or she earned, with a minimum participation fee of \$10 and maximum potential earnings of \$20.[7]

[7]Experimental earnings that a player would earn under full information, common knowledge and optimal decisions (i.e., had all supply chain members known the actual demand and ordered inventory quantity equal to demand) were calculated. Given the demand stream each player faced, a corresponding rate from experimental to USD dollars was calculated so that the maximum profit a player could make was \$10. Experimental profits of each player were then divided by his/her corresponding rate to calculate his/her earnings in USD. These dollar earnings were rounded up to the highest integer and added to the player's \$10 participation fee.

4.5.2 Experimental results

A continuum of trust exists both when pecuniary incentives are aligned and misaligned

Table 4.6 and Table 4.7 present, for all treatments, the summary statistics for retailers' reported information and CP's inventory decisions, respectively. First, we observe that the average forecast distortion is relatively higher in case R_2D_U, compared to the other treatments, as expected from the model predictions. More interestingly though, in cases R_2D_U and R_2D_K where there is competition for common inventory, in the majority of cases we observe forecast inflation (in 61.0% and 72.3% of observations for case R_2D_U and R_2D_K respectively).[8] Despite the fact that under R_2D_K retailers inflate their reports more often, the average magnitude of inflation is lower by almost 10 units (4.6 versus 14.4). In the single retailer case (case R_1D_U) we observe both forecast inflation and deflation, with the latter more pronounced (participants deflated twice as often as they inflated their forecasts while the average forecast distortion is negative, -1.51). In addition, a single retailer truthfully reported his/her received signal in almost half of the instances. Interestingly, in all cases, the magnitude or the direction of forecast distortion does not seem to depend on the magnitude of the signal received. In other words, we observe no pattern of $\hat{\theta}_i - \theta_i$ as a function of θ_i, for $i = 1, 2$.

We also observe that the CP, in all treatments, places an Order Quantity that is lower, on average, than the one suggested by the system as optimal, if the reported demand information was accurate. Similarly, the deviation from the suggested optimal quantity is more pronounced in case R_2D_U, as expected. When there is competition for inventory, the CP, in the majority of cases, sets a quantity lower than the suggested as optimal. Under case R_2D_U though, even if the CP orders a quantity lower than the suggested one, it is on average about 15 units above what would be optimal under common knowledge. This suggests that in this case the CP "discounts" the

[8] the percentages when taking the average per round are even higher; in 93.3% and 86.6% of rounds, respectively, we observe forecast inflation

	Averages per round								Individual observations			
Treatment	No of obs	Avg	[med]	(s.d.)	Inflate	Deflate	Truth	Slope $\hat{\theta}$ on θ	No of obs	Inflate	Deflate	Truth
Case R_1D_U	30	2.3	[2.8]	(1.7)	33.3%	63.3%	3.3%	0.997	150	16.7 %	36.0%	47.3%
Case R_2D_U	60	15.0	[12.2]	(11.6)	93.3%	6.7%	-	0.983	300	61.0%	23.0 %	16.0 %
Case R_2D_K	60	5.8	[4.0]	(4.5)	86.7%	13.3%	-	0.989	240	72.3%	21.7%	7.1%

Notes. "Obs", "Avg", "med" and "s.d." stand for observations, average, median and standard deviation of $|\hat{\theta} - \theta|$, respectively. Regression: $\hat{\theta} = a\theta + b$; all slopes are significant at 0.1% level.

Table 4.6: Summary Statistics on Reported $|\hat{\theta} - \theta|$

received signals, but not sufficiently. In other words, the CP shows more trust to the received information than she should, assuming that she correctly calculates the overage / underage cost trade-off. It is also interesting to note that under R_2D_K the percentage of time the CP trusts the reported information is the lowest among all cases. On average though, she discounts the reported signal less than in the case R_2D_U. This behavior is consistent with the observation that when the number of retailers is two and demand is constant (case R_2D_K), retailers inflate as often (per round) or even more often (individual observations), but not as much compared to when demand is uncertain (case R_2D_U).

Treatment	No of obs	$\lvert Q - Q(\hat{\theta})\rvert$			$Q - Q(\hat{\theta})$	$Q > Q(\hat{\theta})$	$Q = Q(\hat{\theta})$	$Q < Q(\hat{\theta})$	$Q - Q_f^{CP}$	$Q > Q_f^{CP}$	$Q = Q_f^{CP}$	$Q < Q_f^{CP}$	Q on $Q(\hat{\theta})$
		Avg	[med]	(s.d.)	Avg				Avg				
Case R_1D_U	150	14.84	[11]	(14.57)	-0.11	41.3%	16.7%	42.0%	-1.61	40.7%	5.3%	54.0%	0.926
Case R_2D_U	150	24.93	[13]	(31.55)	-17.74	22.0%	25.3%	52.7%	14.47	60.0%	2.0%	38.0%	0.938
Case R_2D_K	120	8.87	[6.5]	(9.87)	-8.43	22.0%	6.7%	71.7%	0.79	45.0%	8.3%	46.7%	0.968

Notes. "Obs", "Avg", "med" and "s.d." stand for observations, average, median and standard deviation, respectively. $Q(\hat{\theta})$ is the Q suggested as the optimal order quantity if the CP trusted the reported demand information. Q_f^{CP} is the optimal Q had the CP known the realized θ_i's (common knowledge) Regression: $Q = aQ(\hat{\theta}) + b$; all slopes are significant at 0.1% level.

Table 4.7: Summary Statistics on Total Order Quantity Q

Figures 4-3 and 4-4 present graphically the experimental data. In Figure 4-3, the reported demand signals (forecasts) are plotted against the real signals, for each of the three treatments. If retailers had been always reporting the truth, all observations would fall on the 45° line drawn in the plots. It is visually evident that in cases R_2D_U and R_2D_K retailers more often inflate their reports than they deflate them (more observations are above the diagonal line). Further, we observe that in case R_2D_U the variability and the magnitude of information distortion are higher than in the other two cases (consistent with the descriptive statistics).

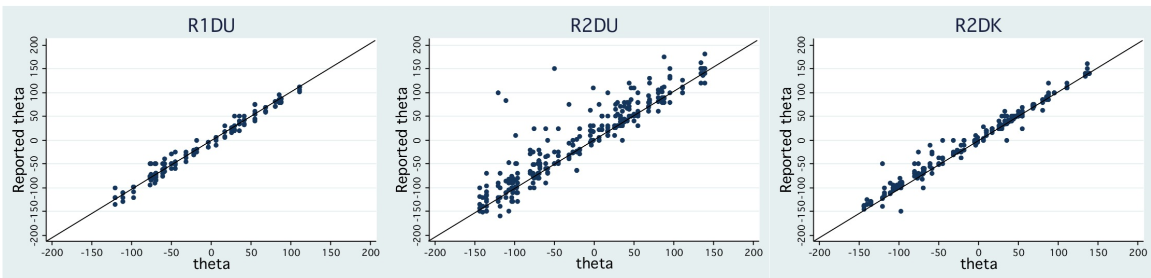

Figure 4-3: Retailers' reported versus actual demand signal

Figure 4-4 plots the inventory quantity set by the CP versus the forecasts she received (mean demand plus reported signals). As aforementioned, optimal inventory quantity, if the reported demand information was accurate, equals the sum of the reported forecasts. If the CP had trusted the received information, her optimal inventory choices would fall on the the 45° line in these graphs. We observe, instead, deviations under all treatments, less pronounced though under case R_2D_K. The horizontal line represents the optimal inventory if the CP ignores the received demand information. These graphs suggest that although the CP usually does not fully trust the reports, she takes the information sent by the retailers into consideration when setting the common inventory. We continue by formally testing the Hypotheses laid out in the previous section.

To test Hypothesis 1(a), we first observe that in 71 out of 150 instances retailers reported the true signal but the average reported signal is the same as the true signal only in 1 of the 30 rounds played (3.3%). Because players both inflate and deflate the received signal, we cannot reject the hypothesis that the average $\hat{\theta}$ is different than

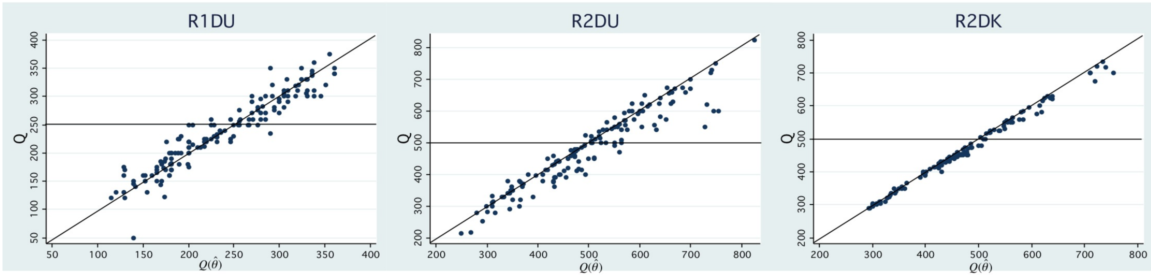

Figure 4-4: CP's inventory choice versus reported forecasts ($Q(\hat{\theta})$)

θ (two-sided Wilcoxon signed rank test, $p > 0.01$). In case R_2D_K, only in 17 out of 240 instances retailers reported the true signal (7.1%) and the average reported signal was never the same as the true signal. The two-sided Wilcoxon signed rank test shows that the average $\hat{\theta}_i$ is significantly different from θ_i ($p < 0.01$) and we can reject Hypothesis 4 (a). To test Hypothesis 2 we regress $\hat{\theta}_i$ per round on θ_i for case R_2D_U. The slope is significantly positive ($p < 0.01$) suggesting a strong positive correlation between $\hat{\theta}_i$ and θ_i. Therefore, we find evidence supporting Hypothesis 2. We also regress $\hat{\theta}_i$ per round on θ_i for the other two cases and the resulting slopes are both significantly positive as well ($p < 0.01$) and higher than that of case R_2D_U (Table 4.6). Experimental data refute the extreme theoretical cases of fully trustworthy ($\hat{\theta}_i = \theta_i$) and fully not trustworthy ($\hat{\theta}_i$ and θ_i are uncorrelated) retailers based on the supply chain environment. Retailers, in all cases, tend to cooperate and be partially trustworthy; in none of the cases they credibly share their forecasts but in all cases their reported signals are informative (strong positive correlation with the true signals). On the other hand, data suggest a directional shift of demand forecast reporting consistent with theory, which is further explored in the following section.

A similar analysis shows that when we regress Q on $Q(\hat{\theta})$, the resulting slopes are significantly positive ($p < 0.01$) for all treatments (Table 4.7). Therefore, Q and $\sum_i \hat{\theta}_i$ are positively correlated[9]. Further, the two-sided Wilcoxon signed rank tests show that Q is not significantly different from $Q(\hat{\theta})$ in case R_1D_U ($p > 0.05$) but the difference is statistically significant in the other 2 cases ($p < 0.01$). Hence, we find

[9]in all experiments the critical ratio is 0.5 and thus $Q(\hat{\theta}) = \sum_i \mu_i + \hat{\theta}_i$

evidence supporting Hypotheses 1(b) and 3 but Hypothesis 4(b) is rejected. However, even under case R_1D_U, the CP fully trusted the reported forecasts in less than 17% of the cases. Combining our findings, we have evidence that under all treatments participants tend to trust and cooperate. However, this does not imply neither that participants fully trust each other nor that the level of trust and cooperation is invariant under all treatments. The latter is the topic of the subsequent analysis.

Impact of inventory competition and forecast uncertainty on trust and cooperation

We investigate the impact of inventory competition and forecast uncertainty on the efficacy of forecast sharing. We use the following linear models to test the treatment effects regarding retailers' reports and CP's order quantity decisions:

$$\hat{\theta}_{it} - \theta_{it} = \beta_0^R + \beta_c^R \cdot C + \beta_u^R \cdot U + \beta_T^R \cdot t + \epsilon_{it} \tag{4.6}$$

$$Q_{jt} - Q(\hat{\theta})_{jt} = \beta_0^{CP} + \beta_c^{CP} \cdot C + \beta_u^{CP} \cdot U + \beta_T^{CP} \cdot t + \omega_j + e_{jt} \tag{4.7}$$

$$E_{kt} = \beta_0^E + \beta_c^E \cdot C + \beta_u^E \cdot U + \beta_T^E \cdot t + \beta_D^E \cdot D_{kt} + v_{kt} \tag{4.8}$$

The subscript i in Equation (4.6) denotes retailer i, where $i = 1, 2$, and $\hat{\theta}_{it}$ is the average report of all retailers i in round t (these retailers received the same signal). With this model, we average out individual-specific effects. The subscript j in Equation (4.7) is the index for a participant that was assigned the role of the CP ($j = 1, 2, .., l$, where l equals to the number of groups in each treatment). Even if the same "type" retailers in each group received the same demand signal in a given round, each CP received different reported forecast(s). In model 2, to control for individual heterogeneity and possible correlation in the decisions made by the same individual, CP, we introduce an individual specific error term, ω_j (see, e.g., [42]). The resulting equation is a random effects general linear model (GLM) [25].

We compare also the resulting profits with the profits of a supply chain under *common knowledge*, i.e., when all parties have the same demand information. In this case, we define system efficiency (E) to be equal to $\Pi(Q(\hat{\theta}))/\Pi(Q(\theta))$ (realized system profits to system profits had the CP known the retailers' demand signals and set inventory optimally). Model 3 (4.8), investigates average channel efficiency per round as a function of competition and demand uncertainty when information is transmitted. We control for demand realizations under each treatment and for group specific effects, by taking the average channel efficiency of all groups per round in a given treatment (recall that all groups in the same treatment see the same demand in a given round). Table 4.8 summarizes the definitions of the variables and the error terms. We include also t in all three models, to account for possible learning effects and time trends in the experiments.

Table 4.9 summarizes the regression results. The coefficients for C and U describe the changes in the dependent variables due to competition for common inventory and due to market uncertainty when the forecast information is transmitted, respectively. We observe that in the first model both coefficients are positive and statistically significant ($p < 0.01$), suggesting that both competition for common inventory and market uncertainty lead to higher forecast distortion and more specifically to forecast inflation. We also note that the effect of competition for common inventory is much more pronounced than that of market uncertainty (the coefficient of C is more than 50% higher than that of U). We observe that the impact of these factors is negative but not statistically significant in the second model. Because the difference $Q_{jt} - Q(\hat{\theta})_{jt}$ is in the majority of cases negative (the CP sets inventory less than the suggested as optimal had she believed the signals sent are true), negative coefficient of the treatment effects imply that the CP discounts more the reported information when there is competition for common inventory and when demand uncertainty is present. Even though a negative directional result suggests that the CP trusts less the reported information for determining system inventory when there is competition for common inventory and market uncertainty, we find no statistical support for this.

Variable	Definition
Dependent Variables	
$\hat{\theta}_{it} - \theta_{it}$	Difference between the average reported signal and the signal observed by retailer type i in round t ($i = 1, 2$)
$Q_{jt} - Q(\hat{\theta})_{jt}$	Difference between the (Total) Order Quantity placed by CP j in round t and the quantity suggested as optimal, if the CP believed the reported signals.
E_{kt}	System efficiency under Case k in round t, $k \in \{R_1D_U, R_2D_U, R_2D_K\}$
Treatment dummies	
C	Indicator variable for competition for common inventory; C=1 if the data are from Case R_2D_U or Case R_2D_K and 0 otherwise
U	Indicator variable for demand uncertainty when information is transmitted; U=1 if the data are from Case R_1D_U or Case R_2D_U and 0 otherwise
Other Independent variables	
t	Round, t=1,2...30
D_{kt}	Total Demand in round t under Case k, $k \in \{R_1D_U, R_2D_U, R_2D_K\}$
Error terms	
ϵ_{it}	Independent error across retailer "types" and periods
ω_j	Individual specific error for CPs
e_{jt}	Independent error across (Total) Order Quantity decisions
v_{kt}	Independent error across cases (treatments) and periods

Table 4.8: Variable Definition in Equations (4.6)—(4.8)

We note from the responses of the subjects to the post-experiment survey that the average CP's level of trust to the reported information was similar in all cases (4, 3.6 and 3.8 for cases R_1D_U, R_2D_U and R_2D_K respectively, on a scale from 1 to 5).

	Estimate (Standard error)					
Variable	$\hat{\theta}_{it} - \theta_{it}$		$Q_{jt} - Q(\hat{\theta})_{jt}$		Efficiency (%)	
Intercept	-8.226^{**}	(2.547)	14.470	(13.197)	0.908^{**}	(0.026)
C	15.920^{**}	(1.920)	-17.633	(10.25)	-0.089^{**}	(0.020)
U	9.801^{**}	(1.567)	-9.307	(10.872)	-0.035^{*}	(0.014)
t	-0.199^{*}	(0.081)	-0.340^{**}	(0.119)	0.002^{*}	(0.001)
D	-		-		0.0003^{**}	(0.00006)

$^{*}p < 0.05$ and $^{**}p < 0.01$.

Table 4.9: Impact of competition and market uncertainty on truth-telling, trusting and efficiency

Next, we examine the impact of treatment effects on cooperation and system efficiency. We find statistically significant evidence that system efficiency decreases with competition (almost by 1%) and market uncertainty (less than 0.5%) on average. This is expected as both these two factors lead to an average increase in forecast distortion, while the CP fails to account adequately for it, but keeps his level of trust almost unchanged (according to post-experiment survey answers). System profits decrease because the decision to set the common inventory is based on less accurate information. It is worth mentioning though, that average efficiency in all cases is very high (above 93.5%). Last, we note the significant coefficients for t in all regressions indicate that retailers tend to inflate less their forecasts over time, the CP discounts more the received signals and sets a quantity closer to the system optimal and system efficiency increases over time.

What is the value of information transmission and does the benefit of risk pooling exist under information asymmetry?

In this section we first investigate what is the value of communication among the supply chain parties, even if information transmitted is not fully reliable. For this reason, we compare system profits under each case to profits had the CP ignored the signal(s) sent by the retailer(s). We compute the system profits had the CP set the inventory optimally based only on her knowledge about the mean demand in each region and the distribution of signals and market uncertainty, if any. The realized average experimental profits under each treatment were 15.6%, 27.0% and 23.1% higher than the computed profits with no information transmission, for case R_1D_U, R_2D_U and R_2D_K, respectively. We have evidence that there is significant value in supply chain communication, even when parties do not fully cooperate. This observation is consistent with the finding that the demand signal transmitted is partially informative and the CP that sets inventory partially incorporates it in her decision to set common inventory.

We also examine if the risk pooling benefits of centralization survive information asymmetry when information transmission is not fully reliable. For this reason, we compare the average profits of retailer 1 under case R_1D_U ($n = 1$) and under case R_2D_U ($n = 2$). Under both these cases, by design, retailer 1 sees the same demand streams while market uncertainty after information transmission suggests that there may be benefits from demand risk pooling when $n = 2$ and common inventory is held centrally. We observe, instead, that the average, over all rounds, experimental profits of retailer 1 are lower under case R_2D_U than that under case R_1D_U ($6,353 versus $6,362). Additionally, we compare the average profits per round of retailer 1 under case R_1D_U and case R_2D_U and we cannot reject that they are not different (two-sided Wilcoxon rank test). This is consistent with the finding that system efficiency decreases when there is competition for common inventory. Even if inventory pooling increases expected profits under common knowledge (e.g., from $6,539 to $6,684 for retailer 1) the actual benefits under information asymmetry may be inexistent or even negative due to a decrease in cooperation and quality of information the inventory decision is based on.

In short, our experimental results suggest that there is value in information transmission, in all cases studied, even when theory predicts completely uninformative communication between the parties (case R_2D_U). Other things being equal (i.e., number of retailers, market uncertainty) it is always beneficial to allow communication between the more informed retailer(s) and the CP who sets inventory. On the other hand, we observe that the increase in information distortion, when we move from a single retailer to two retailers setting, out-weights the benefit of demand uncertainty pooling. Our results suggest that when it comes to deciding on inventory pooling in practice, supply chain parties should be aware that pooling may not always be a beneficial strategy under information asymmetry among players.

4.6 Concluding remarks

In this chapter, we study whether demand forecast sharing between retailers, who are better informed about local demand, and the central planner, who sets common system inventory, is reliable both in theory and in practice. We want to further investigate the influence of supply chain environmental factors on trust and performance. These factors include the number of retailers, market uncertainty, and level of automation.

We find that in the communication game between a single retailer and a benevolent CP, truthful information sharing and full trust is an equilibrium. However, when we consider multiple retailers, players' incentives do not coincide and therefore a truthful information sharing equilibrium is not sustainable, unless market uncertainty is resolved before information transmission (i.e., demand uncertainty is zero after local market signal is received by the retailer). The difference between individually and system-wide optimal inventory quantity is attributed to two factors: a) a retailer has partial information about total demand and therefore different belief about its distribution and b) the expected ratio of a retailer's own demand to total demand in case of shortage and in case of surplus may not be equal. Furthermore, we find that as the size of the coalition approaches infinity, truth-telling becomes sustainable again. Most importantly though, the results of our extensive numerical analysis show that the difference between the optimal inventory quantity from the point of view of a retailer and that of the system is extremely small. This implies that players' incentives are not far apart and information distortion could be in practice minimal.

Our experimental results suggest that a continuum of trust exists both when pecuniary incentives are aligned and misaligned. Experimental data refutes the extreme theoretical cases of fully trustworthy or fully not trustworthy retailers, but it suggests a directional shift of information reliability consistent with theory. To be more specific, both competition for common inventory and forecast uncertainty harm truth-telling, trust and cooperation (measured by the resulting system efficiency). Despite the fact that information is not fully reliable, in all our cases, the value of communication was significant. On the other hand, we observed that actual inventory pooling benefits may be inexistent or even negative due to a decrease in cooperation and quality of information the inventory decision is based on.

Chapter 5

Forecast Information Sharing and the Order Quantity Decision: Impact of Inventory Ownership

5.1 Introduction

In the previous chapter, we show that when retailers compete for common inventory they will have, in general, an incentive to misreport private information about their local demands to the decision maker who sets total inventory, as long as demand forecasts are not perfectly accurate (part of demand uncertainty remains unresolved when the communication takes place). In this chapter, adopting a similar information structure setting, we study the role of inventory ownership on individual players' incentives. We compare the resulting inventory levels in the system under local and central decision making, when each unit of inventory in the system, even if held centrally, belongs to a specific retailer. We also study, if in such a setting of allocation "guarantees", through dedicated inventories, credible information sharing between players forms an equilibrium.

To do so, we model a profit-maximizing firm that sells its product in two horizontal markets (e.g., geographical regions) that are subject to demand uncertainty. The firm has to decide on its inventory structure that determines who makes the inventory quantity decisions for each market. In all cases, inventory for both regions is held centrally. But for each region, a separate dedicated inventory quantity is held. In other words, each unit of inventory in the system belongs to a specific market, even before demand is known. After local demands are realized and before inventory is sent to each market, we allow for change of inventory ownership between regions. In such a setting, a minimum quantity of inventory allocation is guaranteed for each market in case of shortage, and a maximum quantity is guaranteed in case of surplus.

Compared to the proportional rule studied in chapter 4, this allocation rule reduces uncertainty for locations in regards to their final allocation. With the proportional to realized demands allocation there is neither maximum nor minimum (other than zero) quantity that a location can be allocated. This will depend not only on final demand realizations but also on the total inventory set by the CP based on her beliefs about demand (that in turn may depend on overoptimistic or very pessimistic forecasts from other locations).

As in the previous chapter, our model differentiates between demand information that is available to all (e.g., past sales data) and local knowledge that is available only to the regional managers (e.g., "feel" about the market, knowledge about trends in colors, styles, sizes, etc.). Local knowledge can be communicated efficiently to the central decision maker (without cost) but maybe untruthfully if the regional managers have an incentive to do so. We assume that each region is a separate business unit (local profit maximizer).

We focus on the impact of who holds the right to manage inventory (to determine the inventories to be held centrally) on supply chain players' incentives and strategic interaction. We employ a game theoretical model of information sharing to explore (a) how the placement of inventory decision rights (central versus regional) influences the total inventory level, and (b) whether a transfer pricing mechanism can be designed that incentivizes regional managers to truthfully share their local knowledge about demand in their regions.

5.2 The setting

We consider two distinct locations (geographical regions), indexed $i, j = 1, 2$, that face stochastic demands and hold stocks at a central location (newsvendor type problem). Each regional manager gains more knowledge about his local demand as time unfolds, which is captured through a demand signal. The demand at each region is given by $d_i = \mu_i + \theta_i + \epsilon_i$ (as described in section 3.2).

We consider two inventory arrangements: *locally managed inventory (LMI)* where inventory decision making is done locally and *central planer managed inventory (CPMI)*, where the right to decide the level of inventory is transferred to the central planer. Under both arrangements, inventory is held at a central location until regional demands are realized; we separate the ownership (with decision rights) and the location of inventories in the system. We use the notion of claims (see Anupindi et al, 2001) [2] that establish ownership for each unit of inventory in the system, regardless of its location. The claims give the units' owner ex ante decision rights regarding its level and ex post decision rights regarding its usage.

Under LMI, each regional manager (RM) determines and owns the quantity to be held centrally for him (dedicated inventory). In other words, the regional manager has a claim to the units of inventory held for him in the central warehouse. After demands are realized, a region's local demand that exceeds its available stock at the central location is satisfied using excess stocks, if any, belonging to the other region. Under CPMI inventory arrangement, the central planner (CP) solicits local demand information from the regional managers (RMs) and sets the inventory quantity that maximizes total system profit. We consider the case where the CP holds for each region a *dedicated* inventory quantity. The sum of these quantities maximizes total expected profit for the system, given the belief of the CP about regional demands.

Regional managers know the way their inventory is calculated as a function of the CP's beliefs. To be more specific, to study the information sharing game between RMs and the CP under CPMI, we assume a certain partition of the expected profit maximizing total inventory quantity into dedicated regional inventories. This allows us to study how the behavior of RMs changes with inventory "guarantees". Inventory is shipped to each region after local demands are realized. We allow for "change of ownership" of regional stocks. In other words, each retailer is guaranteed its dedicated inventory quantity if needed, but change of inventory ownership is allowed, after demand uncertainty is resolved, to balance supply and demand in different regions (exactly as in RM case).

We note that the inventory arrangement and the inventory claims regime are all agreed *ex ante*, before RMs receive their private information. It is a "contract" signed under symmetric information. However, the decision on total inventory is taken under asymmetric information. This is the *interim* stage of the game. Under LMI, when RMs determine their inventories they have incomplete information. Each RM has received his demand signal but he does not know the precise characteristics (demand signal received) of the other player. Market uncertainty is still unresolved for both players (ϵ_i's are random). Similarly, under CPMI, the CP has incomplete information when setting inventories, as she does not observe local demand signals. The allocation of inventory (final shipments to regional markets) is done *ex post*, under perfect information; market uncertainty is realized and local demands become common knowledge to all players. Figure 5-1 presents the events and their timing under the two inventory decision arrangements.

For comparison purposes, we use the two extreme cases (benchmarks): pure decentralized inventory structure where regional managers do not cooperate in managing their inventories and central decision making with complete information (demand signals are common knowledge to all players). In all cases, the CP is *benevolent* in the sense of total system profit maximizer.

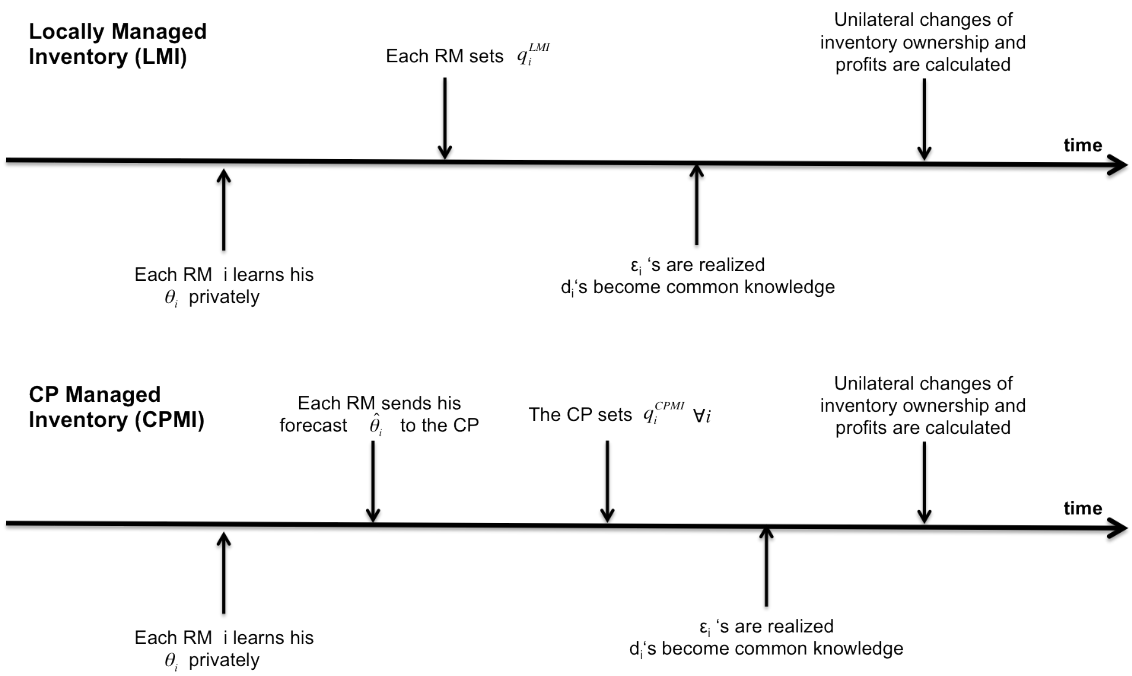

Figure 5-1: The timing of events in the information sharing game with common inventory

We compare the resulting inventory levels and the profit levels of the regions and system as a whole under each inventory arrangement, taking into account players' strategic interaction.

5.3 Locally Managed Inventory (LMI)

A setting lying in between decentralized inventory structure and complete inventory pooling is the one where inventory decisions are made separately by retailers but the allocation of excess stocks is postponed to after the local demands are realized.

Under LMI, RMs, after receiving their demand signal (θ_i), place an initial order to the CP q_i (we denote q_i^a as q_i to simplify the notation here, given that regions place no second order in the central warehouse after demands are realized). The inventory quantity that is held in the central warehouse (Q) is the sum of the individual orders (sourced from an un-capacitated external supplier), as in chapter 3. After demands are realized (d_i's), the CP sends inventory to the regions (α_i) based on the realized local demands. If for one region the inventory held is lower than its realized demand, excess inventory held for the other region, if any, can be sent to satisfy its extra demand. This is the case of decentralized inventory decision making with more accurate information about local demands and shipping (allocation) postponement to after demands are realized. Compared to the pure decentralized inventory structure, the optimal inventory orders from each region, taken unilaterally, are affected by the possibility of claims transfer after regional demands become known.

Define $x^+ = \max\{x, 0\}$. Then, the transfer of inventory claims from region i to region j, τ_{ij}, is given by the lower of two magnitudes, the excess need $(d_j - q_j)^+$ at region j, and the excess stock $(q_i - d_i)^+$ under the ownership of location i.

$$\tau_{ij} = \min[(d_j - q_j)^+, (q_i - d_i)^+] \tag{5.1}$$

Then, allocation that each location gets after demand realizations are observed is given by:

$$\alpha_i = q_i + \tau_{ji} - \tau_{ij} \tag{5.2}$$

Given the allocation, sales at region i are,

$$s_i = \min(\alpha_i, d_i) = \min(q_i, d_i) + \tau_{ji} \tag{5.3}$$

while the unsold stock owned by location i is

$$o_i = (\alpha_i - d_i)^+ = (q_i - d_i - \tau_{ij})^+ \tag{5.4}$$

and the unmet demand is

$$u_i = (d_i - \alpha_i)^+ = (d_i - q_i - \tau_{ji})^+ \tag{5.5}$$

Please note that transfer of inventory claims happens at a unit cost c. Both locations get every unit of inventory at a wholesale price c, analogous to the settings analyzed in previous chapters. We therefore assume, for the time being, that the profit generated by each unit sale is earned by the region where the sale takes place (in section 5.6 we relax this assumption).

Consider a location i that receives a signal θ_i. Following the same notation as in the previous chapter, we denote by $\tilde{d}_i$ the demand at location i after θ_i has been realized; a random variable with mean $\mu_i + \theta_i$. Let $q_i : [\underline{\theta_i}, \overline{\theta_i}] \rightarrow [0, \overline{q}]$ be the inventory choice function and $q_i(\theta_i)$ the inventory choice of location i under the realization θ_i. We assume an upper bound $\overline{q} = \mu_i + \overline{\theta_i} + \overline{\epsilon_i} + \mu_j + \overline{\theta_j} + \overline{\epsilon_j}$ since an inventory choice above the upper bound of the total demand at both locations increases cost at location i without changing expected revenue from sales. Furthermore, inventory choices are non-negative.

The expected profit at location i, under realization θ_j, is:

$$\begin{aligned} \pi_i^{LMI}(q_i, q_j, \theta_i, \theta_j) &= \mathbb{E}_{\epsilon_i, \epsilon_j}[p \min(\tilde{d}_i, \alpha_i) - c\alpha_i] \\ &= \mathbb{E}_{\epsilon_i, \epsilon_j}[ps_i - c(\tau_{ji} - \tau_{ij})] - cq_i \end{aligned} \tag{5.6}$$

We take the first-order derivative of the expected profit at location i, for a specific realization of θ_j, and by rearranging the terms we get:

$$\begin{aligned} \frac{\partial \pi_i^{LMI}(q_i, q_j, \theta_i, \theta_j)}{\partial q_i} &= p(1 - \Pr[\tilde{d}_i < q_i] - \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j]) \\ &\quad - c(1 - \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j] - \Pr[q_i + q_j - \tilde{d}_j < \tilde{d}_i < q_i]) \end{aligned} \tag{5.7}$$

The intuition behind is that one additional inventory unit held for location i will generate one more unit of sales at location i, unless there is local inventory surplus ($d_i < q_i$), or the sale would have been realized despite local inventory shortage, through a transfer of claim from location j (when $q_i < d_i < q_i + q_j - d_j$). This additional sale will generate a revenue p. To hold an additional unit of inventory for location i, an additional cost of c (wholesale price) is incurred unless this unit would have been a claim transfer from location j to i and a cost of c would have been incurred in anycase (whenever $q_i < d_i < q_i + q_j - d_j$). Also, the cost c for this additional unit is not incurred by location i whenever this extra unit is not needed at location i but needed at location j and the claim is transferred (which happens whenever demand exceeds effective supply at location j, i.e., $q_i + q_j - d_j < d_i < q_i$). The difference between the marginal revenue and the marginal cost of an additional unit held for location i gives its net marginal benefit.

Following the analysis of Rudi et al. (2001) [46] we define:
$\eta_i(q_i, \theta_i) = \Pr[\tilde{d}_i < q_i]$
$\beta_i(q_i, q_j, \theta_i, \theta_j) = \Pr[q_i + q_j - \tilde{d}_j < \tilde{d}_i < q_i]$
$\gamma_i(q_i, q_j, \theta_i, \theta_j) = \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j]$

If the joint distribution over demand is continuously differentiable, the above probability functions will be continuous as well. The profit maximizing inventory choice of location i, as a function of the received signals θ_i and θ_j and the quantity set by location j, is given by

$$\eta_i(q_i, \theta_i) + \frac{p-c}{p}\, \gamma_i(q_i, q_j, \theta_i, \theta_j) - \frac{c}{p}\, \beta_i(q_i, q_j, \theta_i, \theta_j) = \frac{p-c}{p} \tag{5.8}$$

Regional manager i knows though only his received signal θ_i and the distribution of θ_j. Therefore, his expected profit, over all possible types θ_j, is

$$\pi_i^{LMI}(q_i, q_j, \theta_i) = \mathbb{E}_{\epsilon_i, \epsilon_j, \theta_j}[p s_i - c(\tau_{ji} - \tau_{ij})] - c q_i \tag{5.9}$$

His best response, as a function of $q_j(\theta_j)$ and his signal θ_i is denoted by ρ_i^{LMI} and given by

$$\rho_i^{LMI}(\theta_i, q_j) = \arg\max_{q_i} \int_{\theta_j} (\mathbb{E}_{\epsilon_i,\epsilon_j}[ps_i - c(\tau_{ji} - \tau_{ij})] - cq_i) dF_j(\theta_j) \tag{5.10}$$

The first order derivative of the expected profit function is given by

$$\begin{aligned}
\frac{\partial \pi_i^{LMI}(q_i, q_j, \theta_i)}{\partial q_i} &= \frac{\partial}{\partial q_i} \int_{\theta_j} (\mathbb{E}_{\epsilon_i,\epsilon_j}[ps_i - c(\tau_{ji} - \tau_{ij})] - cq_i) dF_j(\theta_j) \\
&= \int_{\theta_j} \frac{\partial}{\partial q_i} (\mathbb{E}_{\epsilon_i,\epsilon_j}[ps_i - c(\tau_{ji} - \tau_{ij})] - cq_i) dF_j(\theta_j) \\
&= \int_{\theta_j} [p(1 - \Pr[\tilde{d}_i < q_i] - \Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j]) \\
&\quad - c(1 - \Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j]) \\
&\quad + c \Pr[q_i + q_j(\theta_j) - \tilde{d}_j < \tilde{d}_i < q_i])] dF_j(\theta_j) \\
&= p(1 - \Pr[\tilde{d}_i < q_i] - \mathbb{E}_{\theta_j} \Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j]) \\
&\quad - c(1 - \mathbb{E}_{\theta_j} \Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j] - \mathbb{E}_{\theta_j} \Pr[q_i + q_j(\theta_j) - \tilde{d}_j < \tilde{d}_i < q_i])
\end{aligned} \tag{5.11}$$

The first equation holds because the function under the expectation over θ_j is integrable and has a bounded derivative. So, it satisfies the Lipschitz condition of order one, and hence the expectation and the derivative can be interchanged (see Glasserman, 1994) [24].

The second derivative of the expected profit function of region i is

$$\frac{\partial^2 \pi_i^{LMI}(q_i, q_j, \theta_i)}{\partial q_i^2} = -\int_{\theta_j} [pu_i + (p - c)(g_{ij}^2 - g_{ij}^1) + c(b_{ij}^2 - b_{ij}^1)] dF_j(\theta_j) \tag{5.12}$$

where, letting ϕ_x denote the probability density function associated with the random variable x, we define

$u_i = \phi_{\tilde{d}_i}(q_i)$

$g_{ij}^2 = \Pr[\tilde{d}_i > q_i] \phi_{\tilde{d}_i + \tilde{d}_j | \tilde{d}_i > q_i}(q_i + q_j)$

$g_{ij}^1 = \Pr[\tilde{d}_i + \tilde{d}_j < q_i + q_j | \tilde{d}_i > q_i] \phi_{\tilde{d}_i}(q_i)$
$b_{ij}^2 = Pr[\tilde{d}_i < q_i] \phi_{\tilde{d}_i + \tilde{d}_j | \tilde{d}_i < q_i}(q_i + q_j)$
$b_{ij}^1 = \Pr[\tilde{d}_i + \tilde{d}_j > q_i + q_j | \tilde{d}_i < q_i] \phi_{\tilde{d}_i}(q_i)$
The integrand in (5.12) is a non negative function (see Rudi et al, 2001) [46], leading to the concavity of the expected profit function of region i in its own inventory choice.

Defining
$\beta_i(q_i, q_j, \theta_i) = \mathbb{E}_{\theta_j} \Pr[q_i + q_j(\theta_j) - \tilde{d}_j < \tilde{d}_i < q_i]$
$\gamma_i(q_i, q_j, \theta_i) = \mathbb{E}_{\theta_j} \Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j]$
the profit maximizing inventory choice for retailer i satisfies the equation

$$\eta_i(q_i, \theta_i) + \frac{p-c}{p} \gamma_i(q_i, q_j, \theta_i) - \frac{c}{p} \beta_i(q_i, q_j, \theta_i) = \frac{p-c}{p} \tag{5.13}$$

This defines a reaction function $\rho_i^{LMI}(\theta_i, q_j)$ for location's i optimal inventory, given the received signal θ_i and the inventory choice function of location j, $q_j(\theta_j)$. Concavity of the expected profit function also ensures that the solution of this equation is unique. This means that each location has a unique best response (ρ_i^{LMI}) to the other location's inventory choice, given its type. Moreover, the next Proposition states that this best response function is increasing in the demand signal received (type of location i).

Proposition 5.1: Each location i's best response to the quantity set by the other location is increasing in its signal θ_i (i.e., its type).
All proofs are in the Appendix B.3.

The next theorem shows that in the case of information asymmetry and local decision making, a pure strategy Bayesian Nash equilibrium exists. A Bayesian equilibrium is a pair of strategies $(q_i^{LMI}(\cdot), q_j^{LMI}(\cdot))$ such that, for each player i and every possible value of θ_i, strategy $q_i^{LMI}(\theta_i)$ maximizes $\pi_i^{LMI}(q_i, g_j^{LMI}(\theta_j), \theta_i)$.

Theorem 5.1: For $0 < c < p$, a pure strategy Bayesian Nash Equilibrium exists. Any pair of functions $(q_i^{LMI}(\theta_i), q_j^{LMI}(\theta_j))$ that satisfy (5.14) for $i = 1, 2$ is a Bayesian Nash equilibrium:

$$\eta_i(q_i^{LMI}(\theta_i), \theta_i) + \frac{p-c}{p}\mathbb{E}_{\theta_j}[\gamma_i(q_i^{LMI}(\theta_i), q_j^{LMI}(\theta_j), \theta_i, \theta_j)] - \frac{c}{p}\mathbb{E}_{\theta_j}[\beta_i(q_i^{LMI}(\theta_i), q_j^{LMI}(\theta_j), \theta_i, \theta_j)] = \frac{p-c}{p} \quad (5.14)$$

According to Theorem 5.1, finding an equilibrium in this game requires solving 2 integral equations simultaneously with variable limits. Solving the system of equations to obtain a closed form solution is not possible. Using Proposition 5.1, we know that the optimal inventory choice is monotonic as well. We continue by characterizing the functional form of the optimal inventory choice, as a function of the received signal.

Proposition 5.2: In all equilibria, the inventory choice function of location i is linear in its own received signal and more specifically of the form $q_i^{LMI}(\theta_i) = \mu_i + \theta_i + \delta_i$.

It is interesting to note that if the received signal increases by x units, the optimal inventory choice at location i will increase by the same amount. The received signal determines the mean demand at location i and it is incorporated in location's optimal inventory choice as it is (one-to-one relationship), no matter its magnitude (e.g., regardless of whether it is low or high).

In addition, multiplying equilibrium conditions by $f_i(\theta_i)$ and integrating both sides over θ_i we get that

$$\mathbb{E}[\eta_i(q_i, \theta_i)] + \frac{p-c}{p}\,\mathbb{E}[\gamma_i(q_i, q_j, \theta_i, \theta_j)] - \frac{c}{p}\,\mathbb{E}[\beta_i(q_i, q_j, \theta_i, \theta_j)] = \frac{p-c}{p} \quad (5.15)$$

From (5.15) we see that under information asymmetry, the optimality conditions for a Bayesian Nash equilibrium is the same as the one under full information where the corresponding probabilities of inventory shortage and surplus are calculated over all possible demand signal realizations.

Note that the equilibrium functions described in Theorem 5.1 may result in multiple equilibria that involve different inventory choice functions for each player. In our setting, the two locations have the same revenue and price parameters. If we additionally assume that demand at each location, before the signal transmission, follows the same distribution (i.e., the average market size at each location is the same ($\mu_i = \mu_j$) and signal and market uncertainty at each location follow the same distribution ($f_i(\cdot) = f_j(\cdot)$, $g_i(\cdot) = g_j(\cdot)$), then we can focus on symmetric equilibria, defined as follows:

Corollary 5.1: The symmetric Bayesian Nash equilibrium $(q^{LMI}(\theta_1), q^{LMI}(\theta_2))$ satisfies condition (5.16) for all θ_i.

$$\eta_i(q^{LMI}(\theta_i), \theta_i) + \frac{p-c}{p}\mathbb{E}_{\theta_j}[\gamma_i(q^{LMI}(\theta_i), q^{LMI}(\theta_j), \theta_i, \theta_j)] - \frac{c}{p}\mathbb{E}_{\theta_j}[\beta_i(q^{LMI}(\theta_i), q^{LMI}(\theta_j), \theta_i, \theta_j)] = \frac{p-c}{p} \quad (5.16)$$

5.4 Central Planner Managed Inventory (CPMI)

Under CPMI, inventory decisions are made centrally to maximize total profits. RMs receive their demand signal and send a report to the CP (communicate a demand forecast). The demand signal each RM receives for his region at the beginning of the selling season is private knowledge. Therefore, RMs can choose to communicate their signal truthfully or not. The CP in turn may use this information to determine inventory levels (q_1, q_2) to maximize the sum of profits across locations, based on her beliefs about local demands. This is a case of centralized decision making with demand information asymmetry and information sharing (maybe non-credible) in the form of demand forecasts.

The allocation dynamics are similar to the LMI case. Inventory held in the system ($Q = q_1 + q_2$) is in the form of dedicated inventory for each region but held centrally. After the CP makes the inventory decision, demands are realized (d_i's) and inventory is allocated according to the established policy. Each RM i is guaranteed q_i if $d_i \geq q_i$ and cannot be allocated more than q_i if $d_i < q_i$. To be more specific, the final allocation that region i gets (α_i) is given by equation (5.2), where τ_{ij} is defined as in (5.1). Hence, it satisfies the following relationships; $\alpha_i \leq q_i$ if $d_i < q_i$, $\alpha_i \geq q_i$ if $d_i > q_i$, $\alpha_i = q_i$ if $d_i = q_i$ and $\alpha_i + \alpha_j = q_i + q_j$.

The expected value of total profits across the two locations, denoted as Π^{CPMI}, for given θ_i's, is given by:

$$\begin{aligned}
\Pi^{CPMI}(q_i, q_j, \theta_i, \theta_j) &= p\mathbb{E}_{\epsilon_i,\epsilon_j} \min[(\tilde{d}_i + \tilde{d}_j), (q_i + q_j)] - c(q_i + q_j) \\
&= \mathbb{E}_{\epsilon_i,\epsilon_j}[p \min(d_i, \alpha_i) - c\alpha_i + p \min(d_j, \alpha_j) - c\alpha_j] \\
&= \mathbb{E}_{\epsilon_i,\epsilon_j}[ps_i - c(\tau_{ji} - \tau_{ij})] - cq_i + \mathbb{E}_{\epsilon_i,\epsilon_j}[ps_j - c(\tau_{ij} - \tau_{ji})] - cq_j \\
&= p\mathbb{E}_{\epsilon_i,\epsilon_j}(s_i + s_j) - c(q_i + q_j)
\end{aligned} \tag{5.17}$$

By taking the derivative of the CP's profit function and collecting the terms we have.

$$\begin{aligned}
\frac{\partial \Pi^{CPMI}(q_i, q_j, \theta_i, \theta_j)}{\partial q_i} &= p(1 - \Pr[\tilde{d}_i < q_i] - \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j] \\
&\quad + \Pr[q_i + q_j - \tilde{d}_j < \tilde{d}_i < q_i]) - c
\end{aligned} \tag{5.18}$$

The intuition parallels that of the Locally Managed Inventory with the main difference that the marginal unit generates revenue p when it is used to cover either excess demand at location i (that happens with probability $(1 - \Pr[\tilde{d}_i < q_i] - \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j])$ as in the LMI case) or excess demand at j (which happens when there is excess stock in i and shortage in j, i.e., when $q_i + q_j - \tilde{d}_j < \tilde{d}_i < q_i$). Furthermore, the marginal cost is c with probability 1 (no matter whether or where it is sold).

By setting (5.18) to zero we get the first order necessary optimality condition for q_i, $i = 1, 2$. Therefore, the profit-maximizing inventory choices needs to satisfy, for $i = 1, 2$, the condition (we later show that this condition is also sufficient)

$$\eta_i(q_i, \theta_i) + \gamma_i(q_i, q_j, \theta_i, \theta_j) - \beta_i(q_i, q_j, \theta_i, \theta_j) = \frac{p-c}{p} \tag{5.19}$$

We note that of the left side condition (5.19) can be re-written as:

$$\begin{aligned}
\eta_i + \gamma_i - \beta_i &= \Pr[\tilde{d}_i < q_i] + \Pr[q_i < \tilde{d}_i \cap q_i + q_j > \tilde{d}_i + \tilde{d}_j] - \Pr[\tilde{d}_i < q_i \cap q_i + q_j < \tilde{d}_i + \tilde{d}_j] \\
&= \Pr[\tilde{d}_i < q_i \cap q_i + q_j > \tilde{d}_i + \tilde{d}_j] + \Pr[q_i < \tilde{d}_i \cap q_i + q_j > \tilde{d}_i + \tilde{d}_j] \\
&= \Pr[\tilde{d}_i + \tilde{d}_j < q_i + q_j]
\end{aligned} \tag{5.20}$$

In other words, condition (5.19) for i and j, boils down to the single optimality condition

$$\Pr[\tilde{d}_i + \tilde{d}_j < Q] = \frac{p-c}{p} \tag{5.21}$$

where $Q = q_1 + q_2$. This is the well known newsvendor critical fractile optimality condition, with demand being the sum of local demands. Because the expected profit is concave in Q, it is also concave in q_i (composition of a concave function with an affine function). Therefore, first-order conditions (5.21) and (5.19) are sufficient for optimality. In the case of CPMI, what matters is the total inventory held centrally because (a) retailers are identical in their cost/revenue parameters and (b) inventory is held centrally and sent to retailers after demands are realized (there are no transshipment costs). Q^{CPMI} that maximizes system profits satisfies the optimality condition (5.21); $Q^{CPMI} = \mu_i + \mu_j + \theta_i + \theta_j + (G_i \circ G_j)^{-1}(\frac{p-c}{p})$. Therefore, the inventory choices (q_i^{CPMI}, q_j^{CPMI}) that the CP can make to maximize total profits, are infinite. Any combination of q_i, q_j that satisfies $q_i + q_j = Q^{CPMI}$ is an optimal solution from the system's perspective.

For expositional purposes, we plot total profits as a function of q_1 and q_2, for the case where $p = 2$, $c = 1$ (i.e., critical fractile is 0.5), the average demand in each market is 10 units, $\theta_1 + \theta_2 = 0$ and market uncertainty in each region is independent and distributed normally, with mean 0 and variance 2 (Figure 5-2). In this case it is optimal from a total profit perspective to set $Q^{CPMI} = 20$. As we observe from the figure, any combination q_i and q_j that adds up to 20 units maximizes the expected profit function.

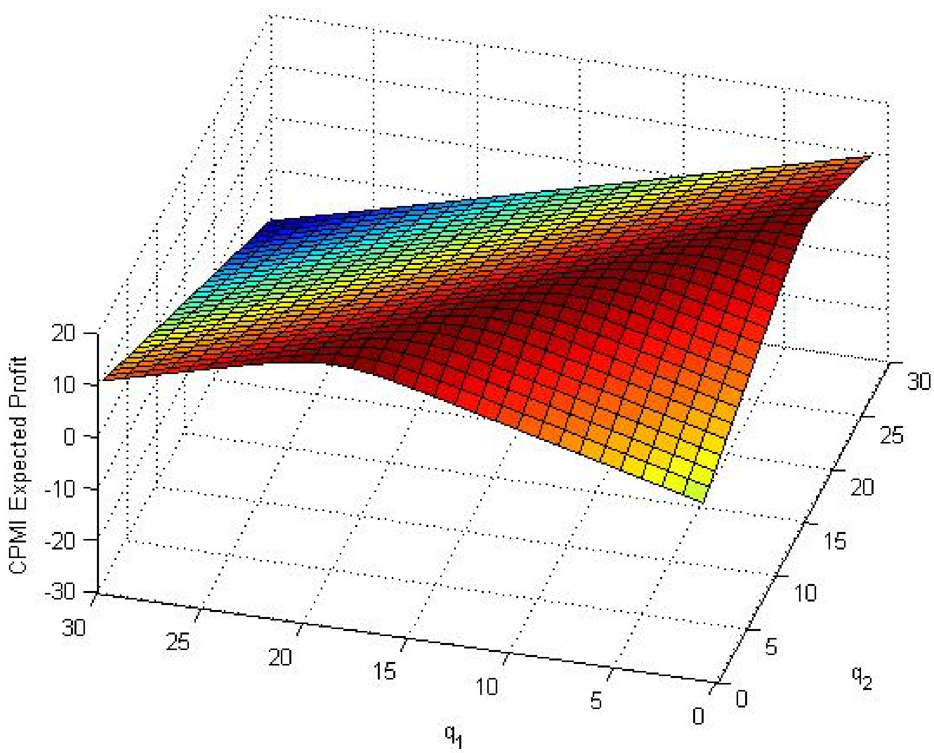

Figure 5-2: Total Expected Profit as a function of q_1 and q_2

For given θ_i's, q_i^{CPMI} can take any value in the interval $[0, Q^{CPMI}]$. This implies, that before inventory is decided by the CP, and most importantly when information transmission takes place, there are no minimum or maximum inventory guarantees for retailer i. Since there is no unique optimal solution (q_i^{CPMI}, q_j^{CPMI}), retailers cannot infer the inventory choices of the CP, given her beliefs about local demands.

To overcome this issue, we need to specify how the CP decides inventory ownership of Q^{CPMI} units between the two locations. For the rest of the chapter, where we study the information sharing game between locations and the CP, we assume that $q_i^{CPMI} = \mu_i + b(\theta_i) + s$, $\forall i$, where $b(\theta_i)$ denotes the belief of the CP about θ_i and $s = \frac{(G_i \circ G_j)^{-1}(\frac{p-c}{p})}{2}$. Please note that s can be either positive or negative, depending on the critical ratio.

5.4.1 Is truth-telling an equilibrium?

When all parties have the same information about demand (common knowledge setting), in the case of CPMI where inventory decisions are centrally coordinated, by definition the aggregate profits across locations are maximized. But under information asymmetry, is this still the case? When the CP solicits private demand information from the RMs and then she sets the inventory quantities (q_i^{CPMI}, q_j^{CPMI}) to maximize system profits, would RMs have an incentive to truthfully communicate their demand information?

To answer this question, we first study whether the incentives of a RM and the CP coincide under *common knowledge* (all players know the demand signal realization in both regions). In other words, given θ_i and θ_j, we compare $q_i^{LMI}(q_j)$ to $q_i^{CPMI}(q_j)$. For notational convenience, for the rest of the section we denote $\beta_i(q_i, q_j, \theta_i, \theta_j)$ and $\gamma_i(q_i, q_j, \theta_i, \theta_j)$ by β_i and γ_i, respectively.

Proposition 5.3: For given θ_i, θ_j, $q_i^{CPMI} \neq q_i^{LMI}$, $\forall$ q_j and $i = 1, 2$, unless $\frac{\gamma_i}{\beta_i} = \frac{p-c}{c}$ for some q_j, where the probabilities β_i and γ_i are evaluated at (q_i^{CPMI}, q_j) and $q_i^{CPMI} = Q^{CPMI} - q_j$.

Proposition 5.3 states that locally and centrally optimal inventory choice for region i, given the inventory of region j, will be different, unless the probability that excess demand at region i can be covered by excess inventory at region j over the probability that excess supply at region i can be used to satisfy excess demand at region j equals the markup (profit margin as a percentage of the cost). Both these probabilities are evaluated at the inventory choice q_i that maximizes aggregate profits, given inventory for region j.

It is therefore implied that if each regional manager had the right to set his own inventory, knowing both his demand signal and that of the other location, he would have chosen a different quantity than the one the CP will set for him. Keeping q_j fixed, q_i^{CPMI} does not maximize RM_i's profit as by definition that quantity is q_i^{LMI} given by equation (5.8). This result is similar to the case of proportional allocation under common knowledge (Lemma 4.1). In both cases, when demand signals are common knowledge, the quantity that a retailer / regional manager prefers is different to the one that maximizes system profits, unless special conditions happen to hold.

Corollary 5.3: Under common knowledge, if RMs had the right to decide their inventories, none of the centrally chosen inventory choices (q_i^{CPMI}, q_j^{CPMI}) that maximize system profits would form an equilibrium, unless $\frac{\gamma_i}{\beta_i} = \frac{\gamma_j}{\beta_j} = \frac{p-c}{c}$ evaluated at some (q_i^{CPMI}, q_j^{CPMI}).

To see this, let us assume that $q_j^{CPMI}(q_i) = q_j^{LMI}(q_i)$ for some q_i. That would mean that the CP chooses an inventory quantity for region j so that its regional expected profit is maximized, given the inventory choice for location i. The inventory choice of the CP for region i, in turn, needs to be such that the aggregate profits are maximized ($q_i^{CPMI} = Q^{CPMI} - q_j^{LMI}$). But according to proposition 5.3, for any given q_j (and thus also $q_j = q_j^{LMI}$), the quantity that maximizes the expected profit of region i is different than the quantity that maximizes system profits ($q_i^{LMI} \neq q_i^{CPMI}$) unless $\frac{\gamma_i}{\beta_i} = \frac{p-c}{c}$. The same argument holds for the second region. In short, central optimal inventory choices will coincide with local choice equilibrium quantities under common knowledge *if and only if* cost and demand parameters are such that equate left-hand sites of Equations (5.8) and (5.19) for $i = 1, 2$.

But when regional manager i reports his local demand information (sends a signal $\hat{\theta}_i$), he knows θ_i but not θ_j. Therefore he cannot directly compare q_i^{LMI} to q_i^{CPMI}, had both players had the same information. Regional manager i does not know q_i^{CPMI}, no matter the reporting strategy of the other regional manager and the CP's beliefs given the reporting strategy. To study the information sharing game between a regional manager and the central planner, we assume that among the infinitely many optimal inventory choices given her demand beliefs, the CP sets $q_i^{CPMI} = \mu_i + b(\theta_i) + s$, for $i = 1, 2$.

We next study whether truth-telling and trusting can form an equilibrium. Suppose that location i sends a signal $\hat{\theta}_i$, which may or may not be the same as his true signal θ_i, but location j transmits his true demand signal and the central planner believes that both locations transmit their true demand signals. As the next proposition shows, retailer i will have, in general, an incentive to report his demand signal falsely. Proposition 5.4 holds for any q_i^{CPMI} that depends on $b(\theta_i)$.

Proposition 5.4: When regional manager j truthfully reports his demand signal and the central planner trusts the received information, regional manager i has an incentive to falsely report his demand signal, unless $\frac{\gamma_i(q_i,\theta_i)}{\beta(q_i,\theta_i)} = \frac{p-c}{p}$ evaluated at $q_i^{CPMI}(\theta_i)$.

Having analyzed regional managers' inventory choices given their information versus inventories that maximize aggregate profits, we show that truthful information sharing of private demand information will not maximize, in general, the expected profits of regional managers (unless special conditions happen to hold). Therefore, reliable information sharing will not be a sustainable equilibrium in this setting as the next theorem formally states.

Theorem 5.2: Let $\phi(\hat{\theta}_i|\theta_i)$ denote regional manager's i reporting strategy given θ_i, $(q_i(\hat{\theta}_i), q_j(\hat{\theta}_j))$ the inventory choices (to be held centrally) of the CP for regions i and j respectively and $b(\theta|\hat{\theta})$ the central planner's posterior belief about θ after observing $\hat{\theta}$. Then,

- $\phi(\hat{\theta}|\theta) = \theta$

- $(q_i(\hat{\theta}_i, \hat{\theta}_j), q_j(\hat{\theta}_j, \hat{\theta}_j)) = (q_i^{CPMI}(\hat{\theta}_i), q_j^{CPMI}(\hat{\theta}_i))$

- $b(\theta|\hat{\theta}) = \hat{\theta}$

do not constitute a perfect Bayesian equilibrium. A perfectly informative Bayesian equilibrium of forecast sharing where regional managers share truthfully their private demand information and the CP trusts the information received and exactly incorporates it in her inventories choices does not exist.

On the other extreme, it is easy to show that an uninformative (babbling) equilibrium exists (as in all cheap talk games). In this equilibrium, each regional manager's reported forecast $\hat{\theta}_i$ is independent of θ_i. The CP does not update her beliefs about local demands after receiving the forecasts. She determines the system optimal quantities based on her initial knowledge about the distribution of θ_i's.

5.5 Numerical analysis

To complement our analytical findings, we continue by conducting numerical analysis to compare optimal inventory choices, and the corresponding expected profits, under LMI and CPMI when there is information asymmetry and unreliable demand forecast sharing. We then compare the results to the case where the CP has the same knowledge with RMs about local demands (she knows local demand signals) and to the pure decentralized case (two separate newsvendors that keep their inventories separately and no transshipment policy is in place).

Numerically solving for the equilibria defined by the integral equations (5.15) requires discretizing the distributions of θ_i's to n points and solving $2 \cdot n$ equations simultaneously. If we restrict ourselves to symmetric equilibria, the system of equations reduces to n but computing the equilibria remains computationally challenging. Furthermore, when we move to the discrete case, as the number of types increases, the problem of estimating the probability of their realization also increases. For these reasons, in inventory management games with incomplete information it is advisable to limit the number of types to a small number in order to successfully solve the problem and obtain useful insights [58].

Thus, for the numerical analysis, we just consider two types of possible demand signals, “Low” and “High”. We assume that θ_i can take only two values, θ_H with probability λ and θ_L with probability $1-\lambda$ and that they are independent across locations. Because of independence $\Pr[\theta_j = \theta_L | \theta_i = \theta_L] = \Pr[\theta_j = \theta_L | \theta_i = \theta_H] = 1-\lambda$ and $\Pr[\theta_j = \theta_H | \theta_i = \theta_L] = \Pr[\theta_j = \theta_H | \theta_i = \theta_H] = \lambda$. We denote the inventory choice of player i that receives θ_L by $q_i(\theta_L) = q_{iL}$, for $i = 1, 2$. Similarly, $q_i(\theta_H) = q_{iH}$. The expected payoff for each player as a function of his type is:

$$\begin{aligned}
\pi_{1L}^{LMI}(q_{1L}, q_2(\theta_2)) &= (1-\lambda)\pi_{1L}^{LMI}(q_{1L}, q_{2L}; \theta_{2L}) + \lambda\pi_{1L}^{LMI}(q_{1L}, q_{2H}; \theta_{2H}) \\
\pi_{1H}^{LMI}(q_{1H}, q_2(\theta_2)) &= (1-\lambda)\pi_{1H}^{LMI}(q_{1H}, q_{2L}; \theta_{2L}) + \lambda\pi_{1H}^{LMI}(q_{1H}, q_{2H}; \theta_{2H}) \\
\pi_{2L}^{LMI}(q_1(\theta_1), q_{2L}) &= (1-\lambda)\pi_{2L}^{LMI}(q_{1L}, q_{2L}; \theta_{1L}) + \lambda\pi_{2L}^{LMI}(q_{1H}, q_{2L}; \theta_{1H}) \\
\pi_{2H}^{LMI}(q_1(\theta_1), q_{2H}) &= (1-\lambda)\pi_{2H}^{LMI}(q_{1L}, q_{2H}; \theta_{1L}) + \lambda\pi_{2H}^{LMI}(q_{1H}, q_{2H}; \theta_{1H})
\end{aligned} \tag{5.22}$$

To determine the Bayesian Nash equilibrium under LMI in this case, we need to solve the following system of four nonlinear equations with four unknowns:

$$\begin{aligned}
\tfrac{\partial}{\partial q_{1L}}\pi_{1L}^{LMI}(q_{1L}, q_2(\theta_2)) &= (1-\lambda)\tfrac{\partial}{\partial q_{1L}}\pi_{1L}^{LMI}(q_{1L}, q_{2L}; \theta_{2L}) + \lambda\tfrac{\partial}{\partial q_{1L}}\pi_{1L}^{LMI}(q_{1L}, q_{2H}; \theta_{2H}) \\
\tfrac{\partial}{\partial q_{1H}}\pi_{1H}^{LMI}(q_{1H}, q_2(\theta_2)) &= (1-\lambda)\tfrac{\partial}{\partial q_{1H}}\pi_{1H}^{LMI}(q_{1H}, q_{2L}; \theta_{2L}) + \lambda\tfrac{\partial}{\partial q_{1H}}\pi_{1H}^{LMI}(q_{1H}, q_{2H}; \theta_{2H}) \\
\tfrac{\partial}{\partial q_{2L}}\pi_{2L}^{LMI}(q_1(\theta_1), q_{2L}) &= (1-\lambda)\tfrac{\partial}{\partial q_{2L}}\pi_{2L}^{LMI}(q_{1L}, q_{2L}; \theta_{1L}) + \lambda\tfrac{\partial}{\partial q_{2L}}\pi_{2L}^{LMI}(q_{1H}, q_{2L}; \theta_{1H}) \\
\tfrac{\partial}{\partial q_{2H}}\pi_{2H}^{LMI}(q_1(\theta_1), q_{2H}) &= (1-\lambda)\tfrac{\partial}{\partial q_{2H}}\pi_{2H}^{LMI}(q_{1L}, q_{2H}; \theta_{1L}) + \lambda\tfrac{\partial}{\partial q_{2H}}\pi_{2H}^{LMI}(q_{1H}, q_{2H}; \theta_{1H})
\end{aligned} \tag{5.23}$$

where $\frac{\partial}{\partial q_{iL}}\pi_{iL}^{LMI}(q_{iL}, q_{jL}; \theta_{jL})$ is given by (5.7) where $\theta_i = \theta_L$, $\theta_j = \theta_L$, $q_i = q_{iL}$ and $q_j = q_{jL}$ and $\frac{\partial}{\partial q_{iL}}\pi_{iL}^{LMI}(q_{iL}, q_{jH}; \theta_{jH})$ is given by (5.7) where $\theta_i = \theta_L$, $\theta_j = \theta_H$, $q_i = q_{iL}$ and $q_j = q_{jH}$. Similarly, for $\frac{\partial}{\partial q_{iH}}\pi_{iH}^{LMI}(q_{iH}, q_{jL}; \theta_{jL})$ and $\frac{\partial}{\partial q_{iH}}\pi_{iH}^{LMI}(q_{iH}, q_{jH}; \theta_{jH})$, with the difference that $\theta_i = \theta_H$, for $i = 1, 2$.

We compute explicit solutions for the case where demand at the two locations is distributed independently or $\mathrm{Cov}(\theta_i, \theta_j) = 0$, $\mathrm{Cov}(\epsilon_i, \epsilon_j) = 0$, for $i, j = 1, 2$ and $i \neq j$ and $\mathrm{Cov}(\theta_i, \epsilon_j) = 0$, for $i, j = 1, 2$. It is also assumed that retailers are symmetric in their cost parameters. Suppose average demand at each location, before the market signal, is 10 units ($\mu_1 = \mu_2 = 10$). Demand signals at location i and j are independent draws from the same discrete distribution: $\theta_L = -2$, $\theta_H = 2$, with $\lambda = 0.5$. Market uncertainty at each location is distributed normally, with mean 0 and standard deviation 2. In other words, each region, before the local demand signal is revealed, faces demand $d_i \sim N(8, 2^2)$ with probability 0.5 and $d_i \sim N(12, 2^2)$ with probability 0.5. After receiving a demand signal at the beginning of the selling season, RM i knows that the demand at his region is normally distributed with mean $10 + \theta_i$ and standard deviation 2 ($\tilde{d}_i \sim N(10 + \theta_i, 2^2)$). The belief of RM i about demand at region j does not change, as it is common knowledge to all players that demands in the two regions are independent.

The resulting equilibrium inventories under LMI and the expected profits are computed for different pairs of cost / price parameters (the corresponding critical ratios are 0.25, 0.5 and 0.75). We report the values for location i, as the resulting equilibrium is symmetric (Table 5.1). In the same table we report the separate newsvendors optimal quantities - without unilateral transfer rights - (q_i^{NV}) and the expected profit in that case (π_i^{NV}), for comparison purposes.

	$\mathbf{q_i^{LMI}}$		$\pi_{\mathbf{i}}^{\mathbf{LMI}}$		$\mathbf{q_i^{NV}}$		$\pi_{\mathbf{i}}^{\mathbf{NV}}$	
Critical ratio	θ_L	θ_H	θ_L	θ_H	θ_L	θ_H	θ_L	θ_H
p=2, c=1.5 *(cr=0.25)*	7.15	11.15	3.10	5.10	6.65	10.65	2.73	4.73
p=2, c=1 *(cr=0.50)*	8.00	12.00	6.87	10.87	8.00	12.00	6.40	10.40
p=2, c=0.5 *(cr=0.75)*	8.85	12.85	11.10	17.10	9.35	13.35	10.69	16.69

Table 5.1: Optimal inventory choices under LMI: ($q_i^{LMI}(\theta_i) = (q_{iL}, q_{iH})$)

Equilibrium inventory quantities are increasing in the received signal and the difference in the optimal inventory choice equals the difference in the received signal (in all three cases $q_i^{LMI}(\theta_H) - q_i^{LMI}(\theta_L) = \theta_H - \theta_L$). This is consistent with Proposition 5.2 for the continuous case. Expected profit is higher when the received signal is "High" because average demand at location i is higher while demand variability remains unchanged. Compared to the simple newsvendor inventory quantities, the possibility of inventory rebalancing after demands are realized may lead to lower (when the cr=0.75) or higher inventory (when cr=0.25), but in all cases increases region's expected profit. The result is analogous to the case of locally held inventories with the option of transshipments after demands are revealed.

We compute (q_{iL}, q_{iH}) for different values of λ, and we observe that the optimal strategy $q_i^{LMI}(\theta_i)$ (for given critical fractile) does not depend on the probability that the other location will face low or high demand; it remains unchanged for any $\lambda \in (0, 1)$. This is because of the one-to-one relationship between the optimal inventory choice and the received signal. More specifically, for the continuous case we have shown that the optimal strategy of location i is of the form $q_i^{LMI}(\theta_i) = \mu_i + \theta_i + \delta_i$ (Proposition 5.2). Equilibrium results in Table 5.1 suggest that this holds also for the discrete case with two types. When $q_{jH} - q_{jL} = \theta_H - \theta_L$, it is easy to show that $\frac{\partial}{\partial q_{iL}}\pi_{iL}^{LMI}(q_{iL}, q_{jL}; \theta_{jL}) = \frac{\partial}{\partial q_{iL}}\pi_{iL}^{LMI}(q_{iL}, q_{jH}; \theta_{jH})$ and $\frac{\partial}{\partial q_{iH}}\pi_{iH}^{LMI}(q_{iH}, q_{jL}; \theta_{jL}) = \frac{\partial}{\partial q_{iH}}\pi_{iH}^{LMI}(q_{iH}, q_{jH}; \theta_{jH})$. Hence, λ plays no role for the solution of the system of equations (5.23) that determines the optimal inventory choice function of location i.

We proceed by studying the boundary cases where $\lambda = 0$ and $\lambda = 1$. In these cases, each regional manager knows with certainty the type of the other player. In other words, there is common knowledge about regional demand distributions and the setting is similar to the one considered in Rudi et al. (2001) [46]. Table 5.2 presents the optimal inventory choice of location i $q_{if}^{LMI}(\theta_i, \theta_j) = (q_{iL}^f(\theta_j), q_{iH}^f(\theta_j))$, as a function of its received signal and the demand (signal) of location j.

It is interesting to notice that location i's optimal inventory choice does not depend on the signal received at location j ($q_{iL}^f(\theta_j = L) = q_{iL}^f(\theta_j = H)$ and $q_{iH}^f(\theta_j = L) = q_{iH}^f(\theta_j = H)$). The signal at location j determines its mean demand and the

	q^f_{iL}		q^f_{iH}	
Critical ratio	$\theta_j = L$	$\theta_j = H$	$\theta_j = L$	$\theta_j = H$
p=2, c=1.5 *(cr=0.25)*	7.15	7.15	11.15	11.15
p=2, c=1 *(cr=0.50)*	8.00	8.00	12.00	12.00
p=2, c=0.5 *(cr=0.75)*	8.85	8.85	12.85	12.85

Table 5.2: Optimal inventory choices under LMI and common knowledge

regional manager at j adequately accounts for it by adjusting his dedicated inventory. According to our analytical and numerical results, the probability that location j will have high / low demand signal does not affect the optimal inventory strategy of location i. Furthermore, the optimal quantity at location i, as a function of its demand signal, is the same under common knowledge and information asymmetry ($q_i^{LMI}(\theta_i) = q_{if}^{LMI}(\theta_i)$). This means that, under LMI, information asymmetry does not create additional system inefficiency. Each regional manager has the necessary information to decide optimally on his inventory level.

The quantity δ_i depends on both locations' market uncertainty and the critical fractile but not on the average demand at location j. This is in contrast to the numerical results obtained under common inventory and proportional allocation. In chapter 4 we showed that the reporting strategy of a retailer may change based on his received signal when the average market size of the other player is high (Figure 4-2 b). Figure 5-3 shows (q_{iL}, q_{iH}) and their difference for various levels of market uncertainty ($\sigma_{\epsilon_i} = \sigma_{\epsilon_j}$), when cr=0.75.

Next, we compare expected total inventory in the system under LMI ($\mathbb{E}[Q^{LMI}] = \mathbb{E}[q_1^{LMI} + q_2^{LMI}]$) against the inventory that the CP would set (CPMI) when she does not update her belief about local demands based on the forecasts sent by RMs (Q^{CPMI}). The CP ignores the received information in the uninformative commu-

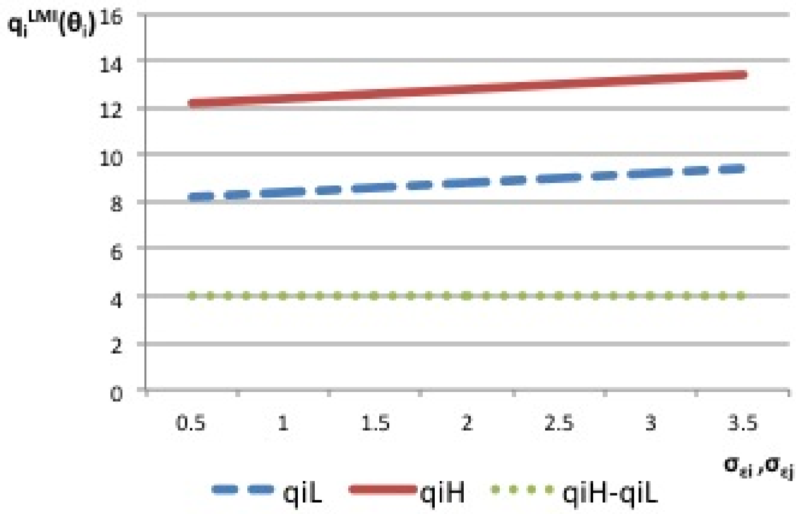

Figure 5-3: $q_i^{LMI}(\theta_H) - q_i^{LMI}(\theta_L)$ as a function of market uncertainty

nication equilibrium detected in section 5.4.1 and bases her inventory decision on her prior information about local demands; $d_i \sim N(8, 2^2)$, with probability 0.5 and $d_i \sim N(12, 2^2)$ with probability 0.5, for $i = 1, 2$. In addition, we use as benchmark the case where the CP knows local demand signal realizations when she sets inventories ($\mathbb{E}[Q_f^{CPMI}]$) (this case also represents the truth-telling / trusting situation). We compute average total inventory and expected profits. The results are presented in Table 5.3. In Table 5.4 we report inventory choices for each location, based on the true or reported signal, under these three inventory decision making regimes.

	LMI		**CPMI**		**CPMI full info**	
Critical ratio	$\mathbb{E}[Q^{LMI}]$	$\mathbb{E}[\Pi^{LMI}]$	Q^{CPMI}	$\mathbb{E}[\Pi^{CPMI}]$	$\mathbb{E}[Q_f^{CPMI}]$	$\mathbb{E}[\Pi_f^{CPMI}]$
p=2, c=1.5 *(cr=0.25)*	18.30	8.20	17.73	7.42	18.09	9.25
p=2, c=1 *(cr=0.50)*	20.00	17.74	20.00	16.77	20.00	17.74
p=2, c=0.5 *(cr=0.75)*	21.70	28.20	22.26	27.42	21.91	28.21

Table 5.3: Comparison of total inventory and expected profits under LMI, CPMI and CPMI with full information

	q_i^{LMI}		q_i^{CPMI}		q_{if}^{CPMI}	
Critical ratio	θ_L	θ_H	$\hat{\theta}_L$	$\hat{\theta}_H$	θ_L	θ_H
p=2, c=1.5 *(cr=0.25)*	7.15	11.15	8.86	8.86	7.05	11.05
p=2, c=1 *(cr=0.50)*	8.00	12.00	10.00	10.00	8.00	12.00
p=2, c=0.5 *(cr=0.75)*	8.85	12.85	11.13	11.13	8.95	12.95

Table 5.4: Comparison of local inventory under LMI, CPMI and CPMI with full information

When the CP can see local demand forecasts with the same level of detail and accuracy as RMs, local decision making results, on average, in higher than optimal inventories in the system when the critical fractile is 0.25 and lower than optimal when the critical fractile is 0.75. On the other hand, average inventory when each regional manager behaves as an independent newsvendor, is higher than the system optimal (if inventory ordering decisions had been coordinated by a central planner and unit transfers had been allowed) for critical ratios above 0.5 (22.70 versus 21.91 for cr=0.75). So, $\mathbb{E}[q_i^{NV} + q_j^{NV}] > \mathbb{E}[Q_f^{CP}] > \mathbb{E}[q_i^{LMI} + q_j^{LMI}]$ for cr> 0.5. The reverse ordering holds for critical fractile below 0.5 ($\mathbb{E}[q_i^{NV} + q_j^{NV}] < \mathbb{E}[Q_f^{CP}] < \mathbb{E}[q_i^{LMI} + q_j^{LMI}]$ for cr<0.5). The possibility to cover excess demand by inventory from another location or to send excess stock at another location, incentivizes a regional manager to order a lower than centrally optimal quantity when the cost of the product is high and a higher than centrally optimal quantity when the cost is low.

We continue by comparing the resulting inventories when the CP has less information about local demands (she bases her inventory decision on her original beliefs about demands (demand signals at the beginning of the period are ignored). In our numerical examples, in that case, the resulting inventory is lower than that under complete information when the critical ratio is low and higher when the critical ratio is high. This is due to higher uncertainty that the CP faces when she does not know the signal realizations but only their distribution and it is the reason why in all cases expected profits are lower than those under common knowledge.

Comparing local decision making with better local demand information to central (coordinated) decision making with less accurate information, we observe that the directional results of inventory comparisons depend on the critical fractile. For low critical fractiles, LMI will lead to higher inventories, on average, than CPMI. For high critical fractile, the reverse is true; inventories under LMI are lower than those under CPMI. We expect this directional comparison to be true in general, because of the reference point CPMI with complete information.

In all three cases, average inventory is 20 units (mean demand) when the critical ratio is 0.5, but, as expected, the average profits differ under each scenario. The highest total expected profit is achieved under central decision making and complete demand information. In our examples, local decision making based on better local demand information performs better than central decision making with less accurate demand knowledge. This need not always to be the case: the directional results when comparing expected profits under LMI and CPMI, independent of the optimal critical fractile, will depend on the amount and accuracy of information that is locally available (variability of the distribution of θ_i's). In our examples, the value of local information (that is lost when we move to CPMI) is higher than the value of inventory coordination achieved under central decision making.

It is interesting to notice than when the critical ratio is 0.5, locally optimal, under partial information, and system optimal, under complete information, inventory choices coincide, not only in expectation but under all instances. This is because when the critical ratio is half, demand uncertainty does not play a role in inventory decisions; it is always optimal to order the average demand (even when players are asymmetric). Each regional manager knows his signal when he sets inventory for his location and the CP, who sets inventory for both locations, knows both signals.

Hence, their optimal inventory choices coincide. This result sheds light to when there is misalignment of incentives between regional managers and the CP; when the critical fractile is different than 0.5 and RMs have private information about their regional demands, the difference between locally and centrally optimal inventory quantities may result in unreliable information sharing. The next question that we would like to answer is whether a transfer pricing mechanism exists that leads to an alignment of incentives and therefore induces reliable information sharing between the players. That is the topic of the next section.

5.6 Coordinating transfer prices

In the previous section, change of inventory ownership, after demands are realized, costs the region that gets the additional unit c and generates revenue c respectively to the region that gives the unit. It is therefore implicitly assumed that a profit margin $(p - c)$ is earned by the market where the sale took place, no matter if the product unit sold was initially owned by the other location. In addition, the location that gives out one unit of excess inventory to the market that is stocked out recovers the full procurement cost c. In this section we study the case where when a unit of excess inventory owned by location i is used to cover excess demand at location j, the additional revenue from the sale p is split arbitrarily between the two locations. Is there a way to split the additional revenue so that truthful demand information sharing is induced?

To model this situation, we denote by c_{ij} the price location i charges location j for a unit sold at j that was owned initially by location i. Location j finds it profitable to use an excess unit initially owned by i to generate a sale at market j when the latter is stocked out, only when $c_{ij} < p$. After demands are realized, region i is willing to "sell" to region j any excess units as long as $c_{ij} > 0$. In other words, change of ownership from market i to market j is mutually profitable only whenever there is excess demand at location j and excess supply at location i. In this case, location i gains from the additional sale c_{ij} and location j gains $p - c_{ij}$.

In such a setting, *under LMI*, the expected profit of region i, as a function of inventory quantities, for given θ_i and θ_j, is given by the expression:

$$\pi_i^{LMI}(q_i, q_j, \theta_i, \theta_j) = \mathbb{E}_{\epsilon_i, \epsilon_j}[ps_i + c_{ij}\tau_{ij} - c_{ji}\tau_{ji}] - cq_i \tag{5.24}$$

By taking the derivative with respect to q_i we get:

$$\begin{aligned} \frac{\partial \pi_i^{LMI}(q_i, q_j, \theta_i, \theta_j)}{\partial q_i} &= p(1 - \Pr[\tilde{d}_i < q_i] - \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j]) \\ &+ c_{ji} \Pr[q_i < \tilde{d}_i < q_i + q_j - \tilde{d}_j] \\ &+ c_{ij} \Pr[q_i + q_j - \tilde{d}_j < \tilde{d}_i < q_i] \\ &- c \end{aligned} \tag{5.25}$$

The first term is the expected additional revenue from having one more inventory unit available at site i, as in expression (5.7). The difference with (5.7) is that now an additional cost c is incurred with probability 1, while the additional second and third term denote the savings from not transferring one unit from market j (that happens if there is excess demand at i and inventory surplus at j) and the revenue c_{ij} from selling one unit to region j (that happens when there is inventory surplus at i while shortage at j).

Following the same notation, the first-order sufficient optimality condition, for region i, after demand signals are realized, becomes:

$$\eta_i(q_i, \theta_i) + \frac{p - c_{ji}}{p}\, \gamma_i(q_i, q_j, \theta_i, \theta_j) - \frac{c_{ij}}{p}\, \beta_i(q_i, q_j, \theta_i, \theta_j) = \frac{p - c}{p} \tag{5.26}$$

Under CPMI, optimality conditions (5.24) do not change. From a central planner's perspective, it does not matter how profit is split between the two regions.

Does it exist a pair of transfer prices between the two markets that could induce truthful sharing of private demand information between RMs and the central planner that sets inventories? To answer this, we are looking for transfer prices that align individual and system incentives; which make the optimal quantities set by the CP to simultaneously satisfy the individual optimality conditions of the RMs. In other words, we are looking for transfer prices that under LMI would result in inventory equilibrium quantities equal to the quantities chosen by the CP as total profit maximizing.

Theorem 5.3: For a given pair of received signals (θ_i, θ_j), there exist transfer prices c_{ij}, for $i, j = 1, 2$ that align individual locations' and company's incentives. These are given by

$$c_{ij} = (\frac{\beta_j \gamma_i - \beta_i \beta_j}{\gamma_i \gamma_j - \beta_i \beta_j})p \tag{5.27}$$

where these probabilities are evaluated at a chosen (q_i^{CPMI}, q_j^{CPMI}).

Theorem 5.3 implies that there are infinitely many pairs of transfer prices that achieve alignment of incentives. For a any pair (q_i^{CPMI}, q_j^{CPMI}) chosen by the CP, such that $q_i^{CPMI} + q_j^{CPMI} = Q^{CPM}$, a unique set of transfer prices can be calculated. These transfer prices make system optimal inventory choices also optimal for individual locations. However, the CP does not know ex-ante Q^{CPM}, because she does not know the demand signal realizations. She can only credibly commit to inventory choices that maximize individual and system profits, given the transfer prices set. In other words, we model the case where individuals locations report their demand signals to the CP considering the direct effect on inventories and not the indirect one on transfer prices between locations. Inventory choices of the CP are system wide and locally optimal, given the transfer prices set after information transmission. But would locations have an incentive to misreport reported signals in order to influence the transfer prices set by the CP? That would depend on the mechanism used to split Q^{CPMI} ownership between the two locations (as a function of reported demand signals) which in turn would determine the corresponding transfer prices. We consider it an interesting extension for future work.

5.7 Concluding remarks

In this chapter we compare local to central inventory decision making, when the latter is based on less accurate demand information. We also study whether under central inventory coordination, truthful information sharing is expected between regional managers that have better information about their local demands and the central planner that sets inventory. Compared to chapter 4, we consider the case where a scheme with inventory guarantees is in place through dedicated inventories for each region.

Under LMI, each region determines its inventory that is held centrally before the beginning of the selling season. When we consider two inventory locations, a pure strategy Bayesian equilibrium exists. In all equilibria, the inventory choice of a location is increasing monotonically in its received signal (i.e., in its average demand). Under CPMI, what matters for the CP is the total inventory held for both locations and not how this is split between them. Hence, when each regional manager sends his demand information to the CP, he cannot anticipate her inventory choice for his region. We show though, that under common knowledge, the system optimal quantity cannot be split between the two regions (dedicated inventories) so that local and system incentives are aligned, unless special conditions happen to hold. Furthermore, we show that truth-telling and trusting do not, in general, constitute a Perfect Bayesian equilibrium under information asymmetry.

We proceed by doing numerical analysis to compare the resulting inventories and the expected profits under LMI and CPMI (considering the babbling equilibrium). In our numerical examples, when the critical fractile is below 0.5, LMI results in higher inventories, on average, than CPMI (and than system optimal inventory under common knowledge). When the critical fractile is above 0.5, the directional results are reverse. Also, we show that, unlike CPMI case, regional demand information asymmetry does not create any additional system inefficiency under LMI. Regarding expected profits, no inferences can be made; which decision making arrangement is preferred will depend on the relative value of local information compared to that of central coordination of inventory choices. Last, even if it is easy to show that a pair of prices, for unit ownership transfers between locations, exists that aligns central and local incentives under common knowledge, it remains an open question if or how such a system could be implemented by the CP, under CPMI and asymmetric information, to incentivize local managers to report truthfully their demand information.

Chapter 6

Conclusions and Discussion

In this thesis, we studied the issue of demand information sharing – forecasts and realized demands – within an inventory pooling coalition. Even though the value of information is a topic well-studied in operations and supply chain management literature, it is usually assumed that information sharing, when it happens, is done in a credible way. On the other hand, many well-documented failures in businesses are due to misreporting of private information, such as order exaggeration in anticipation of inventory shortage or overoptimistic soft orders that never materialize. We focus our work on whether reliable demand information sharing occurs between retailers and a benevolent central planner (CP) who coordinates ordering and inventory allocation within an inventory pooling coalition. Retailers do not compete for demand but they may compete for inventory. Each retailer has some private information about demand in his region due to his proximity to the market, that may transmit truthfully or not to the central planner.

First, we study analytically the impact of various allocation mechanisms on the ordering behavior of retailers, after total inventory quantity in the central warehouse is set. To do so, we first show that when allocation is based on realized demands (i.e., realized demands become common knowledge to all players), all allocation rules considered, i.e., proportional, linear and uniform are efficient (they exclude wastage) and Pareto optimal. But when realized demands in each region remain private knowledge to the retailers, the allocation rule employed plays an important role. Only under a uniform allocation rule, retailers will report their true needs by placing a final order equal to their realized demand. This result is analogous to the typical capacity rationing case, even if in the setting under consideration (a) the allocation is determined after demand uncertainty is resolved and (b) each retailer may receive both below or above his final order. The main difference is that uniform allocation based on final orders is both truth inducing and Pareto optimal in our case. In the capacity rationing case, uniform allocation is not Pareto optimal, because it is not individually responsive, a necessary condition for Pareto optimality. We further propose a modified uniform allocation rule that not only is Pareto optimal and truth inducing, but also it guarantees each retailer a profit higher than what he would have earned in a pure decentralized system, under any demand realization.

Next, we proceed by studying analytically and experimentally demand forecast sharing between retailers and the CP who solicits this information to set total inventory. Realized demands become common knowledge when allocation takes place and we focus on the game of demand signal (forecast) reporting to influence the inventory held in the system. We consider the game where a proportional (to local realized demand) allocation mechanism is employed: a mechanism widely used in practice with several attractive properties. Using game theoretical models, we find that when there is unresolved demand uncertainty when communication takes place, truth-telling and trusting do not form a Perfect Bayesian equilibrium. In addition, under an automated inventory system that takes as input the forecasts reported by the retailers and orders the optimal inventory level for the whole coalition, a pure strategy Bayesian Nash equilibrium among retailers does not exist. We then study, in a controlled laboratory environment that simulates the supply chain setting into consideration, the impact of a) competition for common inventory and b) market uncertainty on information distortion, trust and supply chain efficiency. Our results suggest that a continuum of trust exists both when pecuniary incentives are aligned or misaligned, refuting the extreme theoretical cases of fully trustworthy or fully non-trustworthy retailers. Further, we find that both competition for common inventory and forecast uncertainty harm significantly truth-telling and cooperation among supply chain parties. That's why, inventory pooling under information asymmetry may have negative results despite demand risk aggregation. Even if information was not fully reliable, the value of communication was significant in all our experiments.

Last, we study the impact of inventory ownership on the incentives of the players to truthfully share their forecasts. We find that dedicated inventories do not induce truth-telling either. When we consider two separate locations that locally decide on their level of inventories and take into account the possibility of inventory ownership transfers after demands are realized, there is a unique Bayesian Nash equilibrium. The optimal inventory choice of a location is increasing in its received demand signal. When the CP makes the ordering decision, what matters is only the total inventory held. Unless special conditions happen to hold, it cannot be split between the two locations, before demand is realized, in a way that local and system incentives are aligned. We numerically compare resulting inventories and profits under local decision making with more accurate information versus central decision making where coordination of orders is achieved. We find that when the critical fractile is high, central decision making results in higher inventories while the reverse is true when the critical fractile is low. Expected profit directional comparisons depend on the value of local information that is lost when we move to the central decision making versus the additional value of inventory coordination (in the the babbling equilibrium).

This work has several limitations given the analytical complexity of the problem. We study separately the inventory allocation game when final demands are not known to the central planner and that of demand forecast information sharing to influence the inventory level of the coalition. It remains an open question what the interactions would be when both issues are considered together. For example, if the final allocation is tied to the forecast reported, how would the dynamics of forecast information sharing change? What would be the impact on retailers' ordering behavior and would the final allocation be efficient?

There are several interesting extensions of this thesis. To begin with, we could further explore the impact of behavioral factors in the demand forecast sharing game. When the inventory is common, it is interesting to investigate how the size of the pooling coalition and that of retailers affects trusting relationships. Furthermore, we would like to study whether the level of trustworthiness of retailers changes when there is a guarantee about how their forecasts are used for setting the common inventory level. We consider this an interesting question with potentially very relevant managerial implications. As future research, we are planning to run additional treatments where the number of retailers increases to 3 and 4, the CP is automated and retailers are non-identical. It would be also interesting to experimentally study how inventory ownership impacts retailers' forecast reporting strategy. Even if dedicated inventories do not align individual and system pecuniary incentives, would reduced uncertainty with regards to the final allocation increase retailers' trustworthiness and enhance cooperation?

A second topic of interest is to study through behavioral experiments the impact of allocation mechanisms on retailers' ordering behavior to shed light on potentially "missing" components in this interaction. Do equilibrium concepts, which assume that players are perfectly rational, substantially exaggerate retailers' tendency to strategically order more / less than what they need? Under which allocation mechanisms is order distortion more pronounced?

Another path of future research is to study demand information sharing in an inventory pooling coalition and focus on the behavioral implications when placing soft orders versus sharing demand forecasts. We consider again that each retailer has better demand information due to his proximity to the market and he shares it (maybe untruthfully) with the CP either with the form of forecast sharing (sending his demand signal) or with the form of a non-binding order before demand is realized. How inventory levels and profits compare under these two different ways of information sharing?

A different supply chain setting where demand forecast information sharing plays a crucial role is that of supplier –manufacturers / retailers context. When the upstream player is a separate business unit with her own profit margin, how do the dynamics of communication change? How the issues of trust and trustworthiness impact strategic information transmission and inventory (or capacity) choice and/or allocation, in a supplier – multiple retailers context is still an open question.

A final interesting extension for future work is to study, both analytically and experimentally, the dynamics of communication and information sharing - final demand or forecasts - in multi-period problems. When repeated interactions are considered, under what conditions may truth-telling/trust form a sustainable equilibrium? Do relational contracts (e.g., based on trigger strategies) induce cooperation? In multi-period games, many additional issues may also play an important role, e.g., feedback and learning, reputation considerations, possibility of punishment of untruthful behavior, trust in long-term relationships.

Appendix A

Notation

Symbol	Definition
Common parameters	
n	number of retailers
μ_i	Average demand at location i (a positive constant), $i = 1...n$
θ_i	Demand signal retailer i receives at the beginning of the period (private information about the demand within his region): a zero mean r.v.
$F_i(\cdot)$	Cumulative distribution of θ_i
$\underline{\theta_i}, \overline{\theta_i}$	Minimum and maximum possible values of θ_i respectively
ϵ_i	Market uncertainty within retailer's i region: a zero mean r.v.
$G_i(\cdot)$	Cumulative distribution of ϵ_i
$\underline{\epsilon_i}, \overline{\epsilon_i}$	Minimum and maximum possible values of ϵ_i respectively
d_i	Demand within retailer's i region: a r.v. defined as the sum of three components: $d_i = \mu_i + \theta_i + \epsilon_i$
p	Unit selling price
c	Unit procurement cost

Symbol	**Definition**

Chapter 3

q_i^a	Initial order of retailer i to the warehouse (before market uncertainty is resolved)
q_i^b	Fnal order of retailer i to the warehouse (after market uncertainty is resolved)
Q	Total inventory that is held centrally ($Q = \sum_i q_i^a$)
α	Allocation (final) to retailers vector
d_i	Demand (realization) at location i

Chapter 4

$\hat{\theta}_i$	Report of his signal retailer i sends to the CP (or inputs to the inventory system)
$\tilde{d}_i$	Demand within retailer's i region after he observes his signal θ_i: a r.v. with mean $\mu_i + \theta_i$
$L_i(\cdot)$	Cumulative distribution of $\tilde{d}_i$
$\hat{d}_i$	Demand within retailer's i region, according to the central system, when retailer i inputs $\hat{\theta}_i$: a r.v. with mean $\mu_i + \hat{\theta}_i$
D_i	Total demand from the point of view of retailer i after he receives his signal θ_i
$\hat{D}$	Total demand according to the central system after retailers input their signals, i.e. $\hat{D} = \sum \hat{d}_i$
Q	Total *common* inventory in the system
Q_f^i	Optimal common inventory from the point of view of retailer i when demand signals are common knowledge
Q_f^{CP}	System optimal common inventory when demand signals are common knowledge
$\alpha_i(\mathbf{d}, Q)$	The allocation of common inventory that retailer i gets after demands are realized: a r.v. defined as $\alpha_i(\mathbf{d}, Q) = \frac{d_i}{\sum_{i=1}^n d_i} Q$

Symbol	Definition
$\bar{\alpha}_i(\mathbf{d}, Q)$	The r.v. for the allocation that retailer i will get after demands are realized, after retailer i has received his demand signal, i.e.: $\bar{\alpha}_i(\mathbf{d}, Q) = \frac{\tilde{d}_i}{D_i} Q$

Chapter 5

τ_{ij}	Units claims transferred from region i to region j after demands are realized
q_i	Inventory held centrally that belongs to region i (region i's inventory claims)
$q_i^{LMI}(\theta_i)$	Profit maximizing inventory choice of location i as a function of its signal, under Locally Managed Inventory (LMI)
$\rho_i^{LMI}(\theta_i, q_j)$	Best response function of region i as a function of its signal and the inventory choice of region j
$Q^{CPMI}(\theta_i, \theta_j)$	Profit maximizing total inventory choice of the CP under Central Planner Managed Inventory (CPMI)
q_i^{CPMI}	Inventory for location i (decided by the Central Planner) under (CPMI)
c_{ij}	Price location i charges location j for one unit of inventory transfer after demand realizations

Table A.1: Summary of notation

Appendix B

Proofs

B.1 Proofs of chapter 3

Proof of Proposition 3.1

(a) By definition, a Pareto mechanism is the solution to the following problem:
$\max \sum_i^n \pi_i(\alpha_i, d_i)$
subject to: $\sum_{i=1}^n \alpha_i = Q$
The Lagrange function of the problem is: $\Lambda(\alpha_i, d_i, \lambda) = \sum_{i=1}^n \pi_i(\alpha_i, d_i) + \lambda(\sum_{i=1}^n \alpha_i - Q)$ and by taking the first derivatives, the solution needs to satisfy:
$\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} + \lambda = 0 \;\; \forall i$ and $\sum_{i=1}^n \alpha_i - Q = 0$

(b) First we prove that if an allocation is not efficient it is not Pareto optimal (efficiency is necessary condition for Pareto optimality):
Case 1: $\sum_{i=1}^n d_i \leq Q$ and $\alpha_i < d_i$ for some i. $\sum_{i=1}^n \alpha_i = Q$ implies that $\exists j$ such that $\alpha_j > d_j$. If we decrease α_j by $\delta > 0$ and increase α_i by the same amount, total system profit will increase by δp.
Case 2: $\sum_{i=1}^n d_i \geq Q$ and $\alpha_i > d_i$ for some i. $\sum_{i=1}^n \alpha_i = Q$ implies that $\exists j$ such that $\alpha_j < d_j$. If we decrease α_i by $\delta > 0$ and increase α_j by the same amount, total system profit will increase by δp.

Then, we show that efficiency implies Pareto optimality (sufficient condition): When $\sum_{i=1}^{n} d_i \leq Q$, efficiency implies that $\alpha_i \geq d_i \;\; \forall i$. Because the profit function of retailer i is not differentiable at $\alpha_i = d_i$, we define efficiency as $\alpha_i \geq d_i + \epsilon$ where ϵ is an infinitesimally small positive number. Then, $\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = -c \;\; \forall i \; \Rightarrow$ the allocation is Pareto optimal. When $\sum_{i=1}^{n} d_i > Q$, efficiency implies that $\alpha_i \leq d_i - \epsilon \;\; \forall i$. Then, $\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = p - c \;\; \forall i \; \Rightarrow$ the allocation is Pareto optimal. □

Proof of Proposition 3.2

For identical retailers, when $\sum_{i=1}^{n} d_i > Q$, and d_i increases, $\frac{\partial \Pi}{\partial \alpha_i} = 0$ subject to $\sum_{i=1}^{n} \alpha_i = Q$. This is because $\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = p - c$ remains constant and therefore the optimality condition $\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = \frac{\partial \pi_j(\alpha_j, d_j)}{\partial \alpha_j} \;\; \forall i, j$ still holds and so does the optimal allocation. Similarly, when $\sum_{i=1}^{n} d_i < Q$ and d_i decreases, $\frac{\partial \Pi}{\partial \alpha_i} = 0$ ($\frac{\partial \pi_i(\alpha_i, d_i)}{\partial \alpha_i} = -c$ is constant as well) and optimal allocation does not change. □

Proof of Proposition 3.3

The expected profit of retailer i is: $\pi_i^{CK}(q_i^a) = p\mathbb{E}_{\boldsymbol{\epsilon}_i, \epsilon_{-i}} \min[d_i, \alpha_i] - c\alpha_i$. Given the demand vector $\mathbf{d}$ and total inventory Q, it can be rewritten as:

$\pi_i^{CK}(q_i^a|\mathbf{d}, Q) = -cq_i^a + p\min[d_i, \alpha_i(d_i, q_i^a, Q)] - c[\alpha_i(d_i, q_i^a, Q) - q_i^a]^+ + c[q_i^a - \alpha_i(d_i, q_i^a, Q)]^+$

By definition $\min(d_i, q_i^a) \leq \alpha_i(d_i, q_i^a, Q) \leq \max(d_i, q_i^a)$. We consider two cases:

Case 1: $d_i \leq q_i^a \Rightarrow d_i \leq \alpha_i \leq q_i^a$. In this case:

$\pi_i^{DP}(q_i^a) = -cq_i^a + pd_i + c(q_i^a - \alpha_i) \geq -cq_i^a + pd_i = \pi_i^D(q_i^a)$ Case 2: $d_i > q_i^a \Rightarrow q_i^a \leq \alpha_i \leq d_i$. In this case:

$\pi_i^{DP}(q_i^a) = -cq_i^a + p\alpha_i - c(\alpha_i - q_i^a) = (p-c)q_i^a + (p-c)(\alpha_i - q_i^a) \geq (p-c)q_i^a = \pi_i^D(q_i^a)$.

The result of the proposition follows immediately. □

Proof of Proposition 3.4

For a given Q and a vector of demand realizations $\mathbf{d}$, either $\sum_{i=1}^{n} d_i > Q$ or $\sum_{i=1}^{n} d_i < Q$. Because demand distributions are continuous $\Pr[\sum_{i=1}^{n} d_i = Q] = 0$.

Let's consider the case where $\sum_{i=1}^{n} d_i > Q$. Then $\alpha_i(q_i^{b^*}(d_i), q_{-i}^{b^*}(d_{-i})) < d_i$ for some i. For $q^{b^*}(d)$ to be a dominant equilibrium, i must always maximize his profits by ordering $q_i^{b^*}(d_i)$. We show this is not the case for $q_i^b(d_i) = d_i$ under proportional or linear allocation rule. Under both these mechanisms, $\alpha_i(q_i^b, q_{-i}^b)$ is continuous and increasing in q_i^b. Thus there exists an $\epsilon > 0$ such that $\alpha_i(q_i^b(d_i), q_{-i}^b(d_{-i})) < \alpha_i(q_i^b(d_i) + \epsilon, q_{-i}^b(d_{-i})) < d_i$. As profit is increasing in α_i for $\alpha_i < d_i$, truth-telling cannot be optimal. Similar argument holds for the case where $\sum_{i=1}^{n} d_i < Q$. □

Proof of Proposition 3.5

Following Sprumont's (1991) general proof [52], we consider an arbitrary retailer $i \in n$ with demand d_i:

Case 0: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) = d_i + \sum_{j \neq i} q_j^b(d_j) = Q$ and the retailer orders his demand.

Case 1: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) > Q$. Retailer i might have an incentive to lie only if $d_i > \alpha_i(q_i^{b*}(d_i), q_{-i}^b(d_{-i})) = \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$.

- If he reports some $q_i^b \geq q_i^{b*}(d_i) = d_i$, he gets again $\frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$, so there is no improvement.

- if he reports $q_i^b < q_i^{b*}(d_i) = d_i$, there are two cases:

a) if $\sum_{i=1}^{n} q_i^b > Q$, he gets $\min\{q_i^b, \frac{1}{\hat{n}'}(Q - \sum_{j=\hat{n}'+1}^{n} q_j^b)\}$, where $\hat{n}'$ is the largest integer such that $\alpha_{\hat{n}'}(\mathbf{q}^b, \hat{n}') \leq q_{\hat{n}'}^b$ when retailer i orders q_i^b. If the retailer places $q_i^b > \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$, then $n' = n$ and he gets $\frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$, so there is no improvement. If $q_i^b < \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$, then $\frac{1}{\hat{n}'}(Q - \sum_{j=\hat{n}'+1}^{n} q_j^b) > \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$. But then $\min\{q_i^b, \frac{1}{\hat{n}'}(Q - \sum_{j=\hat{n}'+1}^{n} q_j^b)\} = q_i^b < \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)$ and the retailer is worse off.

b) if $\sum_{i=1}^{n} q_i^b < Q$, then he will receive $\max\{q_i^b, \frac{1}{n-\hat{n}'}(Q - \sum_{j=1}^{\hat{n}'} q_j^b)\}$. This cannot be larger than $\min\{q_i^{b*}, \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)\}$. Otherwise, the equations:

$\min\{q_i^{b*}(d_i), \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)\} + \sum_{j=1, j \neq i}^{n} \min\{q_j^b, \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)\} = Q$ and

$\max\{q_i^b, \frac{1}{n-\hat{n}'}(Q - \sum_{j=1}^{\hat{n}'} q_j^b)\} + \sum_{j=1, j \neq i}^{n} \max\{q_j^b, \frac{1}{n-\hat{n}'}(Q - \sum_{j=1}^{\hat{n}'} q_j^b)\} = Q$

imply that for some $j \neq i$, $\max\{q_j^b, \frac{1}{n-\hat{n}'}(Q - \sum_{j=1}^{\hat{n}} q_j^b)\} < \min\{q_j^b, \frac{1}{\hat{n}}(Q - \sum_{j=\hat{n}+1}^{n} q_j^b)\}$, hence $q_j^b < q_j^b$ (contradiction).

Case 3: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) < Q$. Similar argument as in Case 2 holds. □

Proof of Proposition 3.6

First we prove that truth telling (placing a final order equal to the realized demand) is a dominant strategy equilibrium under a modified uniform allocation rule (the proof is similar to the case of a uniform allocation rule). Consider an arbitrary retailer $i \in n$ with demand d_i:

Case 0: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) = d_i + \sum_{j \neq i} q_j^b(d_j) = Q$ and the retailer orders his demand.

Case 1: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) > Q$. Retailer i might have an incentive to lie only if $d_i > \alpha_i(q_i^{b*}(d_i), q_{-i}^b(d_{-i})) = \frac{1}{\hat{n}}(Q - \sum_{j=1}^{\hat{n}} q_j^a - \sum_{j=\hat{n}+1}^{n} q_j^b) = A_i$

- If he reports some $q_i^b \geq q_i^{b*}(d_i) = d_i$, he gets again A_i, so there is no improvement.

- if he reports $q_i^b < q_i^{b*}(d_i) = d_i$, there are two cases:

a) if $\sum_{i=1}^{n} q_i^b > Q$, he gets $\min\{q_i^b, A_i'\}$, where $A_i' = \frac{1}{\hat{n}'}(Q - \sum_{j=1}^{\hat{n}'} q_j^a - \sum_{j=\hat{n}'+1}^{n} q_j^b)$ and $\hat{n}'$ is the largest integer integer such that $A_i' \leq q_{\hat{n}'}^b$ when retailer i orders q_i^b. But in order to have $A_i' > A_i$, he should order $q_i^b < A_i$. He would thus be worse off.

b) if $\sum_{i=1}^{n} q_i^b < Q$, then he will receive $\max\{q_i^b, B_i'\}$, where $B_i' = q_i^a - \frac{1}{n-\hat{n}'}(\sum_{j=1}^{\hat{n}'} q_j^b + \sum_{j=\hat{n}'+1}^{n} q_j^a - Q)$ and $\hat{n}'$ is the largest integer integer such that $B_i' \leq q_{\hat{n}'}^b$ when retailer i orders q_i^b. This cannot be larger than $\min\{q_i^{b*}, A_i\}$. Otherwise, the equations:

$\min\{q_i^{b*}(d_i), A_i\} + \sum_{j=1, j \neq i}^{n} \min\{q_j^b, A_j\} = Q$ and

$\max\{q_i^b, B_i'\} + \sum_{j=1, j \neq i}^{n} \max\{q_j^b, B_j'\} = Q$

imply that for some $j \neq i$, $\max\{q_j^b, B_j'\} < \min\{q_j^b, A_j\}$. But by definition, $A_j \geq B_j'$ hence $q_j^b < q_j^b$ (contradiction).

Case 3: $q_i^{b*}(d_i) + \sum_{j \neq i} q_j^b(d_j) < Q$. Similar argument as in Case 2 holds. Then, it is easy to see that this allocation excludes wastage when all retailers order their realized demands and that also satisfies condition (3.7). So, the result follows directly from Propositions 3.2 and 3.3. □

B.2 Proofs of chapter 4

Proof of Lemma 4.1

When the CP knows the realized demand signals, she sets Q_f^{CP} (given by equation (4.5)) that minimizes the coalition's expected overage and underage cost, and satisfies the necessary and sufficient condition:

$$(p-c)\Pr[\sum_{i=1}^{n}\tilde{d}_i > Q_f^{CP}(\theta)] = c\Pr[\sum_{i=1}^{n}\tilde{d}_i < Q_f^{CP}(\theta)] \tag{B.1}$$

We define the random variable $\tilde{r}_i = \frac{\tilde{d}_i}{\sum_{i=1}^{n}\tilde{d}_i}$ with known cdf $\Phi_i(\cdot)$ and pdf $\phi_i(\cdot)$. Then, $\tilde{\alpha}_i = \tilde{r}_i Q$. Under symmetric information, the expected profit of each retailer i, can be rewritten as:

$$\begin{aligned}
\pi_i(Q,\boldsymbol{\theta}) &= \mathbb{E}[p\min[\tilde{d}_i,\tilde{\alpha}_i(\mathbf{d},Q)] - c\tilde{\alpha}_i(\mathbf{d},Q)] \\
&= p\mathbb{E}\min[\tilde{d}_i,\tilde{\alpha}_i(\mathbf{d},Q)] - c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q)] \\
&= p\mathbb{E}[\tilde{d}_i - (\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)^+] - c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q)] \\
&= p\mathbb{E}[\tilde{d}_i] - p\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ - c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q)] \\
&= p\mathbb{E}[\tilde{d}_i] - c\mathbb{E}[\tilde{d}_i] + c\mathbb{E}[\tilde{d}_i] - p\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ - c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q)] \\
&= (p-c)\mathbb{E}[\tilde{d}_i] + c\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)] - p\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ \\
&= (p-c)\mathbb{E}[\tilde{d}_i] + c\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ + c\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^- - p\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ \\
&= (p-c)\mathbb{E}[\tilde{d}_i] + c\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ - c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q) - \bar{d}_i]^+ - p\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ \\
&= (p-c)\mathbb{E}[\tilde{d}_i] - [(p-c)\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ + c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q) - \tilde{d}_i]^+]
\end{aligned} \tag{B.2}$$

Given that $(p-c)\mathbb{E}[\tilde{d}_i]$ is a constant, retailer i maximizes his expected profit when he minimizes his expected overage and underage cost $C_i(Q) = (p-c)\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+ + c\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q) - \tilde{d}_i]^+$. Taking the derivative of retailer's i cost function with respect to Q we get:

$$\begin{aligned}\frac{dC_i}{dQ} &= (p-c)\frac{d\mathbb{E}[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+}{dQ} + c\frac{d\mathbb{E}[\tilde{\alpha}_i(\mathbf{d},Q) - \tilde{d}_i]^+}{dQ} \\ &= (p-c)\mathbb{E}\frac{d[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+}{dQ} + c\mathbb{E}\frac{d[\tilde{\alpha}_i(\mathbf{d},Q) - \tilde{d}_i]^+}{dQ} \\ &= (p-c)\mathbb{E}[\frac{d[\tilde{d}_i - \tilde{\alpha}_i(\mathbf{d},Q)]^+}{d\tilde{\alpha}_i(\mathbf{d},Q)}\frac{d\tilde{\alpha}_i(\mathbf{d},Q)}{dQ}] + c\mathbb{E}[\frac{d[\tilde{\alpha}_i(\mathbf{d},Q) - d_i]^+}{d\tilde{\alpha}_i(\mathbf{d},Q)}\frac{d\tilde{\alpha}_i(\mathbf{d},Q)}{dQ}]\end{aligned} \tag{B.3}$$

The second equality holds because the function under the expectation is integrable and has a bounded derivative. So, it satisfies the Lipschitz condition of order one, and hence the expectation and the derivative can be interchanged (see Glasserman 1994). Let 1_ω be the indicator function of the event ω; i.e. $1_\omega = 1$ if the event ω is true and $1_\omega = 0$ otherwise. Using the indicator function we get

$$\begin{aligned}\frac{dC_i}{dQ} &= (p-c)\mathbb{E}[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)}(-1)\frac{d\tilde{\alpha}_i(\mathbf{d},Q)}{dQ}] + c\mathbb{E}[1_{\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d},Q)}1\frac{d\tilde{\alpha}_i(\mathbf{d},Q)}{dQ}] \\ &= -(p-c)\mathbb{E}[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}] + c\mathbb{E}[1_{\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}]\end{aligned} \tag{B.4}$$

But

$$\begin{aligned}\mathbb{E}[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}] &= \mathbb{E}[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)} = 1]\Pr[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)} = 1] \\ &+ \mathbb{E}[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)} = 0]\Pr[1_{\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)} = 0] \\ &= \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)]\Pr[\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d},Q)]\end{aligned} \tag{B.5}$$

Similarly,

$$\mathbb{E}[1_{\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d},Q)}\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}] = \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d},Q)]\Pr[\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d},Q)] \tag{B.6}$$

Therefore, the first-order necessary optimality condition of Q for retailer i is given by equating the derivative (B.4) to 0 and rearranging the terms:

$$(p-c)\Pr[\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d}, Q)]\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d}, Q)] = c\Pr[\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d}, Q)]\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d}, Q)] \quad \text{(B.7)}$$

Under the proportional allocation rule:

$$\Pr[\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d}, Q)] = \Pr[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i} Q < \tilde{d}_i] = \Pr[\sum_{i=1}^n \tilde{d}_i > Q] \quad \text{(B.8)}$$

It is easy to show that $\frac{d^2 C_i}{dQ^2} > 0$. The FOC can be rewritten as

$$\frac{dC_i}{dQ} = -(p-c)\int_{\tilde{r}_i Q}^{\mu_i+\theta_i+\overline{\epsilon_i}} \int_0^1 \tilde{r}_i \Lambda_i(\tilde{r}_i, \tilde{d}_i) d\tilde{r}_i d\tilde{d}_i + c\int_{\mu_i+\theta_i+\underline{\epsilon_i}}^{\tilde{r}_i Q} \int_0^1 \tilde{r}_i \Lambda_i(\tilde{r}_i, \tilde{d}_i) d\tilde{r}_i d\tilde{d}_i$$

where Λ_i denotes the joint pdf of $\phi(\tilde{r}_i)$ and $l_i(\tilde{d}_i)$.

The second order derivative is

$$\frac{d^2 C_i}{dQ^2} = p\int_0^1 \tilde{r}_i^2 \Lambda_i(\tilde{r}_i, \tilde{r}_i Q) d\tilde{r}_i > 0$$

Therefore, the optimal total inventory Q_f^i from the point of view of retailer i is the solution to the equation :

$$(p-c)\Pr[\sum_{i=1}^n \tilde{d}_i > Q_f^i]\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d}, Q^i)] = c\Pr[\sum_{i=1}^n \tilde{d}_i < Q_f^i]\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d}, Q^i)] \quad \text{(B.9)}$$

Conditions (B.1) and (B.9) are equivalent *iff* $\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i > \tilde{\alpha}_i(\mathbf{d}, Q^i)] = \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\tilde{d}_i < \tilde{\alpha}_i(\mathbf{d}, Q^i)]$ which is equivalent to $\mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\sum_{i=1}^n \tilde{d}_i > Q_f^i] = \mathbb{E}[\frac{\tilde{d}_i}{\sum_{i=1}^n \tilde{d}_i}|\sum_{i=1}^n \tilde{d}_i < Q_f^i]$ under the proportional allocation rule. This completes the proof. □

Proof of Theorem 4.1

We denote by $\hat{d}_i$ the random variable with mean $\mu_i+\hat{\theta}_i$, support $[\mu_i+\hat{\theta}_i+\underline{\epsilon_i}, \mu_i+\hat{\theta}_i+\overline{\epsilon_i}]$ and known density function. We define $\hat{D} = \sum_{i=1}^n \hat{d}_i$. The quantity set by the system is the unique solution of the equation:

$$(p-c)\Pr[\hat{D} > Q^{CP}(\hat{\boldsymbol{\theta}})] = c\Pr[\hat{D} < Q^{CP}(\hat{\boldsymbol{\theta}})] \tag{B.10}$$

We define the random variable $\bar{r}_i = \frac{\tilde{d}_i}{\tilde{d}_i + \sum_{j\neq i} d_j}$ with known cdf $\bar{\Phi}_i(\cdot)$ and pdf $\bar{\phi}_i(\cdot)$. Then, $\bar{\alpha}_i = \bar{r}_i Q$. The expected profit of each retailer i, can then be rewritten as:

$$\begin{aligned}\pi_i(Q,\boldsymbol{\theta}) &= \mathbb{E}[p\min[\tilde{d}_i, \bar{\alpha}_i(\mathbf{d},Q)] - c\bar{\alpha}_i(\mathbf{d},Q)] \\ &= (p-c)\mathbb{E}[\tilde{d}_i] - [(p-c)\mathbb{E}[\tilde{d}_i - \bar{\alpha}_i(\mathbf{d},Q)]^+ + c\mathbb{E}[\bar{\alpha}_i(\mathbf{d},Q) - \tilde{d}_i]^+]\end{aligned} \tag{B.11}$$

Following the same steps as in the proof of Lemma 4.1, the first-order necessary and sufficient optimality condition of Q for retailer i is:

$$(p-c)\Pr[\tilde{d}_i > \bar{\alpha}_i(\mathbf{d},Q)]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d},Q)] = c\Pr[\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)] \tag{B.12}$$

Under the proportional allocation rule:

$$\Pr[\tilde{d}_i > \bar{\alpha}_i(\mathbf{d},Q)] = \Pr[\frac{\tilde{d}_i}{D_i}Q < \tilde{d}_i] = \Pr[D_i > Q] \tag{B.13}$$

Therefore, the optimal total inventory Q^i from the point of view of retailer i is the solution to the equation :

$$(p-c)\Pr[D_i > Q^i]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d},Q^i)] = c\Pr[D_i < Q^i]\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q^i)] \tag{B.14}$$

Conditions (B.10) and (B.14) are not equivalent, and hence $Q^{CP} \neq Q^i$, unless $\int\int_{\tilde{d}_i > \bar{\alpha}_i} \bar{r}_i \Gamma(\bar{r}_i, \tilde{d}_i) d\bar{r}_i d\tilde{d}_i = \Pr[\hat{D} > Q^i]\mathbb{E}[\bar{r}_i]$. This completes the proof. □

Proof of Theorem 4.2

$$\int_{\theta_{-i}} \mathbb{E}_\epsilon[P_i(Q(\hat{\theta}_i, \hat{\theta}^*_{-i}(\theta_{-i})), (\theta_i, \theta_{-i})] f(\theta_{-i}) dd\theta_{-i} = \pi_i(Q(\hat{\theta}_i, \hat{\theta}^*_{-i}(\theta_{-i})), \theta_i) \tag{B.15}$$

From Theorem 4.1, there exists a unique $Q^i(\theta_i)$ that maximizes (4.3) (the solution of the equation (B.14)). In an automated inventory system, $Q^{CP}(\hat{\boldsymbol{\theta}}) = K + \hat{\theta}_i + \hat{\theta}_{-i}$, where $K = \sum_{i=1}^{n} \mu_i + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$. Because $\frac{\partial Q^{CP}}{\partial \hat{\theta}_i} = 1$, the best response function of retailer i is given by $\hat{\theta}_i^*(\theta_i, \hat{\theta}_{-i}) = Q^i(\theta_i) - K - \hat{\theta}_{-i}$. The best response function of $\hat{\theta}_i$, for given θ_i, is continuous and decreasing in $\hat{\theta}_{-i}$. But, players' best response functions do not cross because they don't satisfy the condition $|\frac{d\hat{\theta}_i^*}{d\hat{\theta}_{-i}}||\frac{d\hat{\theta}_{-i}^*}{d\hat{\theta}_i}| < 1$. For example, we consider the case of two retailers, i and j. The best response function of each retailer has slope -1 and a different intercept ($Q^i - K$ and $Q^j - K$). Hence, they are parallel; they will have a common point iff $\theta_i = \theta_j$ (leading to infinite number of equilibria). But since $\theta_i^{\prime}s$ are continuous, $\Pr[\theta_i = \theta_j] = 0$. Consequently, the best response functions do not cross and a Bayesian Nash equilibrium does not exist. □

Proof of Theorem 4.3

Given the belief that $b(\theta|\hat{\theta}) = \hat{\theta}$, the central planer's strategy to set $Q(\hat{\theta}) = \sum_{i=1}^{n}(\mu_i + \hat{\theta}_i) + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$ is best response because it is the quantity that maximizes her expected profit. When the CP is fully trusting, her ordering behavior is identical to that of an automated inventory system. From Theorem 1, the quantity that maximizes retailer's i expected profit may be different from $Q(\hat{\theta})$. In that case, retailer i will have an incentive to report $\hat{\theta}_i \neq \theta_i$ (even when $\phi(\hat{\theta}_{-i}|\theta_{-i}) = \theta_{-i}$) and influence $Q(\hat{\theta})$, because $\frac{dQ(\hat{\theta})}{d\hat{\theta}_i} = 1$. In other words, the reporting strategy $\phi(\hat{\theta}_i|\theta_i) = \theta_i$ is not optimal. Therefore, retailer i has an incentive to deviate. Also, given retailers' reporting strategy, the CP's updated belief $b(\theta|\hat{\theta}) = \hat{\theta}$ is not rational. □

Proof of Lemma 4.2

Let $D_{-i} = \sum_{j \neq i} d_j$, $F^{\bigotimes(n-1)}$ and $f^{\bigotimes(n-1)}$ the cdf and the pdf of D_{-i} respectively. Then,

$$\lim_{n \to +\infty} \mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d}, Q)] = \lim_{n \to +\infty} \frac{\int_0^Q \int_0^{Q-t} \frac{\tilde{d}_i}{\tilde{d}_i + t} l_i(\tilde{d}_i) f^{\bigotimes(n-1)}(t) dt d\tilde{d}_i}{\Pr[\tilde{d}_i < \bar{\alpha}_i(\mathbf{d}, Q)]} \tag{B.16}$$

where Q is a function of n ($Q^{CP}(n)$). But, $\lim_{n\to+\infty}\Pr[\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)] = \Pr[\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)] = C$, where C is a constant depending on the strategy of the CP and independent of n, e.g. for $Q^{CP}(\theta)$, $\Pr[\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)] = \frac{p}{c}$. But,

$$\begin{aligned}\int_0^Q \int_0^{Q-t} \frac{\tilde{d}_i}{\tilde{d}_i + t} l_i(\tilde{d}_i) f^{\bigotimes(n-1)}(t)dtd\tilde{d}_i &\leq \int_0^\infty \int_0^\infty \frac{\tilde{d}_i}{\tilde{d}_i + t} l_i(\tilde{d}_i) f^{\bigotimes(n-1)}(t)dtd\tilde{d}_i \\ &\leq \int_0^\infty \int_0^\infty \frac{\tilde{d}_i}{t} l_i(\tilde{d}_i) f^{\bigotimes(n-1)}(t)dtd\tilde{d}_i \\ &= \int_0^\infty \tilde{d}_i l_i(\tilde{d}_i)(\int_0^\infty \frac{f^{\bigotimes(n-1)}(t)}{t}dt)d\tilde{d}_i \\ &= \mathbb{E}[\tilde{d}_i]\int_0^\infty \frac{f^{\bigotimes(n-1)}(t)}{t}dt\end{aligned} \tag{B.17}$$

Hence,
$\lim_{n\to\infty}\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d},Q)] \leq \lim_{n\to\infty}\mathbb{E}[\tilde{d}_i]\int_0^\infty \frac{1}{t}f^{\bigotimes(n-1)}(t)dt = \mathbb{E}[\tilde{d}_i]\lim_{n\to\infty}\int_0^\infty \frac{1}{t}f^{\bigotimes(n-1)}(t)dt$. But, $\lim_{n\to\infty}\int_0^\infty \frac{1}{t}f^{\bigotimes(n-1)}(t)dt = \int_0^\infty \lim_{n\to\infty}\frac{1}{t}f^{\bigotimes(n-1)}(t)dt$, if (a) $\frac{1}{t}f^{\bigotimes(n-1)}(t)$ is Rieman integrable and (b) $f^{\bigotimes(n-1)}(t)$ converges uniformly to $\bar{f}$ as $n\to\infty$. $\frac{1}{t}f^{\bigotimes(n-1)}(t)$ is continuous and therefore Rieman integrable. It remains to show part (b) is true.

We assume that $\sup\{f'^{\bigotimes(n-1)}(t)|t \geq 0\} \to 0$ as $n\to\infty$, i.e. $f'^{\bigotimes(n-1)}(t)$ exists and $f^{\bigotimes(n-1)}$ becomes sufficiently flat when $n\to\infty$. Let $\xi_n : \mathbb{R}\to\mathbb{R}$ with $\xi_n(t) = \frac{1}{t}f^{\bigotimes(n-1)}(t)$. Then,

$$\frac{d}{dt}\xi_n(t) = \frac{\frac{d}{dt}f^{\bigotimes(n-1)}(t)t - f^{\bigotimes(n-1)}(t)}{t^2} \tag{B.18}$$

Setting the derivative to 0, we get: $f'^{\bigotimes(n-1)}(t^*)t^* - f^{\bigotimes(n-1)}(t^*) = 0$. Then, $\xi_n(t^*) = \frac{f^{\bigotimes(n-1)}(t^*)}{t^*} = f'^{\bigotimes(n-1)}(t^*)$. $\xi_{n-1}(\infty) = 0$ and by assumption $\xi_{n-1}(0) = 0$.Then, t^* is the value that yields the $\sup\{\xi_n(t^*)\}$. By definition,

$$\xi_n(t) \leq \sup\{\xi_n(t)|t\geq 0\} = f_n(t^*) = f'^{\bigotimes(n-1)}(t^*) \to 0 \ \ as \ \ n\to\infty \tag{B.19}$$

Hence, $\sup\{|\xi_n(t) - 0| \mid t \geq 0\} \to 0 \quad as \quad n \to \infty$ and therefore $\xi_n(t)$ converges uniformly to $\xi(t) = 0$. Therefore,

$$\lim_{n\to\infty} \int_0^\infty \frac{1}{t} f^{\bigotimes(n-1)}(t)dt = \int_0^\infty \lim_{n\to\infty} \frac{1}{t} f^{\bigotimes(n-1)}(t)dt = 0 \tag{B.20}$$

Similarly, we can prove that $\lim_{n\to+\infty} \mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d}, Q)] = 0$. This completes the proof. □

Proof of Theorem 4.4

a) When $n \to \infty$, from Lemma 4.2, $\mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i < \bar{\alpha}_i(\mathbf{d}, Q)] = \mathbb{E}[\frac{\tilde{d}_i}{D_i}|\tilde{d}_i > \bar{\alpha}_i(\mathbf{d}, Q)] = 0$. Hence, the optimality condition (B.14) of retailer i, $\forall i$, holds for any Q. If all retailers report $\phi(\hat{\theta}_i|\theta_i) = \theta_i$, the resulting system inventory is $Q^{CP} = \sum_{i=1}^n(\mu_i + \hat{\theta}_i) + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$ that satisfies (B.10). Therefore, none of the retailers has an incentive to deviate by misreporting his true signal.

b) Given that $b(\theta|\hat{\theta}) = \hat{\theta}$, the CP orders the system profit maximizing quantity: $Q = (\mu_i + \hat{\theta}_i) + (G_1 \circ G_2 ... \circ G_n)^{-1}(\frac{p-c}{p})$. This quantity satisfies the optimality condition (17) of any retailer i. Thus, retailer i has no incentive to deviate from reporting $\phi(\hat{\theta}_i|\theta_i) = \theta_i$. Given retailers' best reporting strategy, the CP's updated belief $b(\theta|\hat{\theta}) = \hat{\theta}$ is rational. □

Proof of Theorem 4.5

a) When $\phi(\hat{\theta}_i|\theta_i) = \theta_i \quad \forall i$, each retailer gets $\alpha_i = d_i$. This is the quantity that maximizes his profit and therefore retailer i has no incentive to deviate by distorting his reported signal.

b) Given that $b(\theta|\hat{\theta}) = \hat{\theta}$, the CP believes that retailer's i demand is $\mu_i + \hat{\theta}_i$. So, setting $Q = \sum_{i=1}^n d_i = \sum_{i=1}^n(\mu_i + \hat{\theta}_i)$ is best response from the CP to the reported signals $\hat{\theta}_i$'s because it maximizes her profit. When the central planner is fully trusting, and retailer i reports $\phi(\hat{\theta}_i|\theta_i) = \theta_i$, he will get an allocation $\alpha_i(Q) =$

$\frac{\mu_i+\theta_i}{\sum_{i=1}^{n}(\mu_i+\theta_i)}\sum_{i=1}^{n}(\mu_i+\hat{\theta}_i) = \frac{\mu_i+\theta_i}{\sum_{i=1}^{n}(\mu_i+\theta_i)}\sum_{i=1}^{n}(\mu_i+\theta_i) = \mu_i+\theta_i = d_i$ (assuming that the other retailers are following the same strategy). Therefore, retailer i has no incentive to unilaterally distort his reported signal. Given retailers' best response reporting strategy, the CP's updated belief $b(\theta|\hat{\theta}) = \hat{\theta}$ is rational. $\square$

B.3 Proofs of chapter 5

Proof of Proposition 5.1

Let the best response functions for location i at θ_i and at $\theta_i + \delta$ for a given $\delta > 0$ be $q_i = \rho_i^{LMI}(\theta_i, q_j)$ and $q_i^{'} = \rho_i^{LMI}(\theta_i + \delta, q_j)$, respectively. These are the solutions to the following equations:

$$\int_{\theta_j} p(1 - \Pr[\epsilon_i < q_i - \mu_i - \theta_i] - \Pr[q_i - \mu_i - \theta_i < \epsilon_i < q_i + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i]) \\ - c(1 - \Pr[q_i - \mu_i - \theta_i < \epsilon_i < q_i + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i] \\ - \Pr[q_i + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i < \epsilon_i < q_i])dF_j(\theta_j) = \frac{p-c}{p} \tag{B.21}$$

$$\int_{\theta_j} p(1 - \Pr[\epsilon_i < q_i^{'} - \mu_i - \theta_i - \delta] \\ - \Pr[q_i^{'} - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'} + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i - \delta]) \\ - c(1 - \Pr[q_i^{'} - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'} + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i - \delta] \\ - \Pr[q_i^{'} + q_j - \mu_j - \theta_j - \epsilon_j - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'}])dF_j(\theta_j) = \frac{p-c}{p} \tag{B.22}$$

Taking the difference between (B.21) and (B.22) (and substituting $\tilde{d}_j = \mu_j + \theta_j + \epsilon_j$) we get:

$$\int_{\theta_j} (p(-\Pr[\epsilon_i < q_i - \mu_i - \theta_i] + \Pr[\epsilon_i < q_i^{'} - \mu_i - \theta_i - \delta] \\ - \Pr[q_i - \mu_i - \theta_i < \epsilon_i < q_i + q_j - \tilde{d}_j - \mu_i - \theta_i] \\ + \Pr[q_i^{'} - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'} + q_j - \tilde{d}_j - \mu_i - \theta_i - \delta]) \\ - c(-\Pr[q_i - \mu_i - \theta_i < \epsilon_i < q_i + q_j - \tilde{d}_j - \mu_i - \theta_i] \\ + \Pr[q_i^{'} - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'} + q_j - \tilde{d}_j - \mu_i - \theta_i - \delta] \\ - \Pr[q_i + q_j - \tilde{d}_j - \mu_i - \theta_i < \epsilon_i < q_i] \\ + \Pr[q_i^{'} + q_j - \tilde{d}_j - \mu_i - \theta_i - \delta < \epsilon_i < q_i^{'}]))dF_j(\theta_j)) = 0 \tag{B.23}$$

We note that for $q_i^{'} = q_i + \delta$, equality (B.23) holds. Thus, there exists a best response function that is increasing in location i's signal. Combined with that each location has a unique best response to the other location's quantity, we conclude that each location i's best response to q_j is increasing in its type θ_i.□

Proof of Theorem 5.1

In Bayesian games with continuous strategy spaces and continuous types, a pure strategy Bayesian Nash equilibrium exists if (a) strategy sets and type sets are compact and (b) payoff functions are continuous and concave and in own strategies [21]. By assumption, strategy and type spaces consist of closed and bounded single intervals. Therefore, they are nonempty, convex and compact and condition (a) is satisfied. The expected profit at location i is continuous and twice differentiable in q_i and q_j (note that the functions $\eta_i(q_i, \theta_i)$, $\beta_i(q_i, q_j, \theta_i, \theta_j)$ and $\gamma_i(q_i, q_j, \theta_i, \theta_j)$ are continuous, monotonic and differentiable in q_i and q_j). The expected profit at i is also concave in q_i (Eq. 5.12) and thus condition (b) is satisfied as well. The second part of the theorem directly follows from the definition of a Bayesian Nash equilibrium. □

Proof of Proposition 5.2

The derivative of Equation (5.11) with respect to θ_i is given by:

$$\frac{\partial^2 \pi_i^{LMI}}{\partial q_i \partial \theta_i} = \frac{\partial}{\partial \theta_i} \int_{\theta_j} (p(1 - \Pr[\tilde{d}_i < q_i]) - (p-c)\Pr[q_i < \tilde{d}_i < q_i + q_j(\theta_j) - \tilde{d}_j]$$
$$+ c\Pr[q_i + q_j(\theta_j) + \tilde{d}_j < \tilde{d}_i < q_i] - c)dF_j(\theta_j)$$
$$= \int_{\theta_j} pg_i(q_i - \mu_i - \theta_i) - (p-c)(g_i(q_i - \mu_i - \theta_i)$$
$$\Pr[\epsilon_i + \epsilon_j < q_i + q_j(\theta_j) - \theta_i - \theta_j - \mu_i - \mu_j | \epsilon_i > q_i - \mu_i - \theta_i]$$
$$- \Pr[\epsilon_i > q_i - \mu_i - \theta_i]\phi_{\epsilon_i + \epsilon_j | \epsilon_i > q_i - \mu_i - \theta_i}(q_i + q_j(\theta_j) - \theta_i - \theta_j - \mu_i - \mu_j))$$
$$+ c(-g_i(q_i - \mu_i - \theta_i)\Pr[\epsilon_i + \epsilon_j > q_i + q_j(\theta_j) - \theta_i - \theta_j - \mu_i - \mu_j | \epsilon_i < q_i - \mu_i - \theta_i]$$
$$+ \Pr[\epsilon_i < q_i - \mu_i - \theta_i]\phi_{\epsilon_i + \epsilon_j | \epsilon_i < q_i - \mu_i - \theta_i}(q_i + q_j(\theta_j) - \theta_i - \theta_j - \mu_i - \mu_j)$$
$$= \int_{\theta_j} (pa_i + (p-c)(g_{ij}^2 - g_{ij}^1) + c(b_{ij}^2 - b_{ij}^1))dF_j(\theta_j) \tag{B.24}$$

Using the Implicit function theorem, for given q_j, we get:

$$\frac{\partial q_i^{LMI}(\theta_i)}{\partial \theta_i}) = -\frac{\frac{\partial^2 \pi_i^{LMI}(q_i, q_j, \theta_i)}{\partial q_i \partial \theta_i}}{\frac{\partial^2 \pi_i^{LMI}(q_i, q_j, \theta_i)}{\partial q_i^2}} = 1 \tag{B.25}$$

Therefore $q_i^{LMI}(\theta_i) = \theta_i + \delta_i'$ where δ_i' is a constant. We can re-write $q_i^{LMI}(\theta_i) = \mu_i + \theta_i + \delta_i$. $\square$

Proof of Proposition 5.3

It is sufficient to show that the optimality conditions (5.8) and (5.19) are not equivalent, unless $\frac{\gamma_i(q_i,q_j,\theta_i,\theta_j)}{\beta_i(q_i,q_j,\theta_i,\theta_j)} = \frac{p-c}{c}$ at (q_i^*, q_j). Let's assume that for some q_j, $q_i^{LMI} = q_i^{CPMI} = q_i^*$. That means that

$$\begin{aligned} \eta_i(q_i^*, \theta_i) + \frac{p-c}{p}\, \gamma_i(q_i^*, q_j, \theta_i, \theta_j) - \frac{c}{p}\, \beta_i(q_i^*, q_j, \theta_i, \theta_j) = \\ \eta_i(q_i^*, \theta_i) + \gamma_i(q_i^*, q_j, \theta_i, \theta_j) - \beta_i(q_i^*, q_j, \theta_i, \theta_j) \\ \Rightarrow \\ c \cdot \gamma_i(q_i^*, q_j, \theta_i, \theta_j) = (p-c) \cdot \beta_i(q_i^*, q_j, \theta_i, \theta_j) \end{aligned} \tag{B.26}$$

For $\gamma_i(q_i^*, q_j, \theta_i, \theta_j) \neq 0$ and $\beta_i(q_i^*, q_j, \theta_i, \theta_j) \neq 0$, condition (B.26) reduces to $\frac{\gamma_i(q_i,q_j,\theta_i,\theta_j)}{\beta_i(q_i,q_j,\theta_i,\theta_j)} = \frac{p-c}{c}$. Otherwise, equation (B.26) is infeasible, and $\nexists$ q_j such that $q_i^{CPMI} = q_i^{LMI}$. □

Proof of Proposition 5.4

Because retailer j follows his equilibrium strategy of reporting the truth, $\hat{\theta}_j = \theta_j$. Similarly, because the central planner follows her equilibrium strategy in response to the received signals, $b(\theta_i|\hat{\theta}_i) = \hat{\theta}_i$ and $b(\theta_j|\hat{\theta}_j) = \hat{\theta}_j = \theta_j$. Hence, total inventory is given by: $Q^{CPMI} = \mu_i + \mu_j + \hat{\theta}_i + \theta_j + (G_i \circ G_j)^{-1}(\frac{p-c}{p})$. More specifically, $q_i^{CPMI} = \mu_i + \hat{\theta}_i + s$ and $q_j^{CPMI} = \mu_j + \theta_j + s$. Probability $\eta_i(q_i, \theta_i)$ does not depend on θ_j. Substituting q_j^{CPMI} in probability $\beta_j(q_i, q_j, \theta_i, \theta_j)$, we observe that it does not depend on θ_j, when retailer j reports his true signal and the CP believes him.

$$\begin{aligned} \beta_i(q_i, q_j^{CPMI}, \theta_i, \hat{\theta}_j) &= \Pr[q_i, q_j^{CPMI} - d_j < \tilde{d}_i < q_i] \\ &= \Pr[q_i + \mu_j + \theta_j + s - \mu_j - \theta_j - \epsilon_j < \tilde{d}_i < q_i] \\ &= \Pr[q_i + s - \epsilon_j < \tilde{d}_i < q_i] \end{aligned} \tag{B.27}$$

Similarly ,

$$\gamma_i(q_i, q_j^{CPMI}, \theta_i, \hat{\theta}_j) = \Pr[q_i < \tilde{d}_i < q_i + s - \epsilon_j] \tag{B.28}$$

We note, that the inventory dynamics under CPMI and LMI are the same (only the decision rights and hence the information set the decision is based on differ). Therefore, the optimality condition (5.13) characterizes the inventory choice that maximizes regional manager's i expected profit under CPMI, given his information about demand. We see, that the condition is independent of θ_j when RM j tells the truth and the CP believes it and acts optimally upon it.

We define $\beta_i(q_i, q_j^{CPMI}, \theta_i, \hat{\theta}_j) = \beta_i(q_i, \theta_i)$ and $\gamma_i(q_i, q_j^{CPMI}, \theta_i, \hat{\theta}_j) = \gamma_i(q_i, \theta_i)$. Hence, RM i's optimal inventory choice coincides with that of the CP *iff*

$$\frac{\gamma_i(q_i, \theta_i)}{\beta_i(q_i, \theta_i)} = \frac{p - c}{c} \tag{B.29}$$

at $q_i^{CPMI} = \mu_i + \theta_i + s$. Because $\frac{\partial q_i^{CPMI}}{\hat{\theta}_i} = 1 \neq 0$, RM i will have an incentive to report $\hat{\theta}_i \neq \theta_i$ unless (B.29) holds. $\square$

Proof of Theorem 5.2

The central planer's strategy to choose $(q_i^{CPMI}(\hat{\theta}_i), q_j^{CPMI}(\hat{\theta}_j))$, as defined in proposition 5.4, given the belief $b(\theta|\hat{\theta}) = \hat{\theta}$ is best response because expected aggregate profits are maximized. If regional managers send the real demand signal and the CP trusts the information received and incorporates it in her ordering decisions, she will make inventory choices that will not in general maximize the expected profit of a regional manager (Proposition 5.4). Thus, regional managers may have an incentive to distort the information sent (unless $\frac{\gamma_i(q_i,\theta_i)}{\beta_i(q_i,\theta_i)} = \frac{\gamma_j(q_j,\theta_j)}{\beta_j(q_j,\theta_j)} = \frac{p-c}{c}$ evaluated at the inventory choices $(q_i^{CPMI}(\theta_i), q_j^{CPMI}(\theta_i))$. In that case, regional manager i will have an incentive to report $\hat{\theta}_i \neq \theta_i$, even when $\phi(\hat{\theta}_j) = \theta_j$ because $\frac{\partial q_i^{CPMI}}{\hat{\theta}_i} \neq 0$. In other words, the reporting strategy $\phi(\hat{\theta}_i) = \theta_i$ is not optimal. Therefore, regional manager i will have an incentive to deviate for certain realizations of θ_i. Also, given regional managers' reporting strategy, the CP's updated belief $b(\theta|\hat{\theta}) = \hat{\theta}$ is not rational. $\square$

Proof of Theorem 5.3

We look for c_{ij} and c_{ji} that equate conditions (5.19) and (5.13). We have a system of two equations with 2 unknowns:

$$\begin{aligned} \eta_i + \frac{p - c_{ji}}{p}\gamma_i - \frac{c_{ij}}{p}\beta_i = \eta_i + \gamma_i - \beta_i \\ \eta_j + \frac{p - c_{ij}}{p}\gamma_j - \frac{c_{ji}}{p}\beta_j = \eta_j + \gamma_j - \beta_j \end{aligned} \tag{B.30}$$

After some algebraic manipulation, we arrive at the solution (5.25) for $i, j = 1, 2$. To prove existence of solutions such that $0 \leq c_{ij} \leq p$, we first note that if $c_{ij} = p$ for $i, j = 1, 2$, $q_i^{LMI} \geq q_i^{CPMI}$ (comparing (5.19) and (5.13)). Similarly, for $c_{ij} = 0$, $q_i^{LMI} \leq q_i^{CPMI}$. By continuity and because q_i^{LMI} increases both in c_{ij} and c_{ji}, a unique pair of transfer prices that satisfies (B.30) exists for each pair of (q_i^{CPMI}, q_j^{CPMI}). $\square$

Appendix C

Additional Numerical Analysis

Chapter 4

Information asymmetry and cr=0.75 / 0.25 For the base case and asymmetric information, Table C.1 shows the optimal quantity for retailer 1 and the expected system quantity had the two retailers reported the true signals when cr=0.75 and cr=0.25. Similar to when the critical ratio is 0.5, in all instances $Q^1 < \mathbb{E}[Q^{CP}]$ and the absolute difference is increasing in the signal received by retailer 1.

	cr=0.75			**cr=0.25**		
θ_1	Q^1	$\mathbb{E}[Q^{CP}]$	$Q^1 - \mathbb{E}[Q^{CP}]$	Q^1	$\mathbb{E}[Q^{CP}]$	$Q^1 - \mathbb{E}[Q^{CP}]$
-3	18.136	18.168	-0.032	15.795	15.832	-0.037
-2	19.128	19.168	-0.040	16.787	16.832	-0.045
-1	20.123	20.168	-0.045	17.782	17.832	-0.050
0	21.120	21.168	-0.048	18.778	18.832	-0.053
1	22.118	22.168	-0.050	19.777	19.832	-0.055
2	23.117	23.168	-0.051	20.776	20.832	-0.056
3	24.116	24.168	-0.052	21.775	21.832	-0.056

Table C.1: Comparison of the optimal inventory for retailer 1 and the expected inventory in the system under truth-telling, as a function of θ_1, for cr=0.75 and cr=0.25

Discrete uniform demand By simulating the demand distributions used in experimental treatment R_2D_U (discrete uniform), we compute, for different demand signal realizations, the inventory level that maximizes the average profit of retailer 1, assuming that retailer 2 reports his true forecast and the CP trusts the received information and optimally sets the inventory level. We employ the revenue / cost parameters used in the experiment ($p = 2$ and $c = 1$). Table C.2 presents the results. In all cases the difference between the quantity that retailer 1 prefers and the system optimal is very small, less than 1%.

θ_1	Q^1	$\mathbb{E}[Q^{CP}]$	$Q^1 - \mathbb{E}[Q^{CP}]$
-150	352	350	2
-120	382	380	2
-100	402	400	2
-80	420	420	0
-60	440	440	0
-40	459	460	-1
-20	479	480	-1
0	499	500	-1
20	519	520	-1
40	539	540	-1
60	560	560	0
80	580	580	0
100	600	600	0
120	620	620	0
150	650	650	0

Table C.2: The inventory that maximizes average profit for retailer 1, for various values of θ_1, and the system optimal inventory level for discrete uniform demand

Appendix D

Experiment Instructions

D.1 Handout to regional managers: Case R_2D_U

Experiment Instructions

Treatment 1: 2 Regional Managers, a Central Planner, Market Uncertainty

Handout to regional managers

This is an information sharing inventory game

You are part of a supply chain of 3 people: **2 regional managers and 1 central planner**. You are assigned the role of one of the two regional managers. In this task, you will make decisions for up to bf 40 independent rounds (5 trial and up to 35 game rounds). In each round, you will play with 2 other participants in this room (one other regional manager and one central planner). At the beginning of each round all players are **randomly and anonymously** assigned to a supply chain group. Your role will always be a regional manager but you will not be assigned to the same supply chain group in consecutive rounds.

Setting

You are the manager of a regional store that sells seasonal products (e.g., holiday decorations). This store is part of a network of 2 regional stores operated by a parent company.

The company has a central planner who decides on the amount of inventory to purchase for each product in the stores' collection for a given season. For the purpose of this exercise we will focus on one product. To help inform the central planner's decision, each regional manager provides the central planner with a demand forecast several months before the selling season. The central planner sees this information and may use it to determine a total order quantity for the product. The total quantity of the product is held in the parent company's single warehouse and shipped to the regional stores after demand is realized.

The timing of events is as follows:

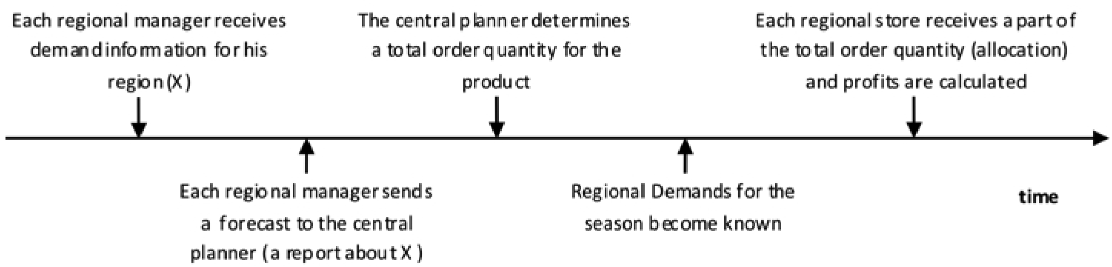

Regional Demand

The product's demand for a given season within one retail region is determined by the sum of two components:

Regional Demand = X + Y

X is randomly generated in each round for each retail region. Each regional manager observes the exact value of X for his (and only his) region. The other regional manager and the central planner only know that X for any retail region is equally likely to be **any integer value between 100 and 400.**

Y is also randomly generated in each round for each retail region. The regional managers and the central planner only know that Y for any retail region is equally likely to be **any integer value between -50 and 50.**

The values of X and Y in each round are independent. They are also independent across retail regions and rounds. Thus, if X and/or Y are large (or small) for a region in the current round, this will not affect whether they are large (or small) in another region or in future rounds.

Your decision

As described above, as a regional manager you have better information about your regional demand because you observe the exact value of your X.

Your decision is to send a report of your information to the central planner. In each round, you will send to the central planner **a report** of X for that season.

Please note that as long as your report is between 100 and 400 units, the central planner cannot verify whether you reported your private information truthfully or not.

Central planner's decision

The central planner sees the reports of Xs for the two regional managers and then chooses the **total order quantity** to purchase for the product (common inventory for the two regions is held centrally). The product costs $1 per unit for the central planner to purchase from a wholesaler and is sold by the regional managers for $2 per unit.

The objective of the central planner is to maximize the companys profit, defined as the revenue earned across all store regions minus the cost of the acquiring the product:
Companys Profit = $2 * min [Total Demand, Total Order Quantity] $1 * Total Order Quantity
Total Demand stands for the sum of product demand across all store regions.

Your profit

As a regional manager, you will be rewarded in proportion to the local profit of your store. For each unit you sell, your store earns $2. For each unit you receive from the company's warehouse, your store pays $1. Your store's profit is calculated as follows:

Local Profit = $2 * min [Regional Demand, Allocation] $1 * Allocation

Your Allocation is based on the *total order quantity* chosen by the central planner and the ratio of your regional demand versus total demand across all regions. It is given by:

Allocation = (Regional Demand / Total Demand)*Total Order Quantity

Please note that with this allocation rule it is possible to receive a quantity below/above your regional demand. No inventory is left at the central warehouse at the end of the season.

Example of the allocation rule

*Suppose your final regional demand is 270 units and the final demand at the other regional store is 330. If the total order quantity was 620 units, your allocation will be (270/(270+330))*620= 279 units (allocation is above your demand). If instead the total order quantity had been 580 units, your allocation would be 270/(270+330))*580= 261 units (allocation is below your demand).*

How you will be paid

If you follow these instructions and make good decisions you will earn a considerable amount of money that will be paid to you in cash at the end of the session.

At the end of the experiment, your stores profit from the game rounds (the first 5 trial rounds do not affect your profits) will be divided by 690 to determine your earnings from the experiment. These dollar earnings will be added to your $10 participation fee and displayed on your computer screen. The maximum amount of dollars you may receive is $20.

When you have completed all rounds of the game, you will be asked to fill out the Post-Game survey. Then the researcher will pass by your carrel desk and you will be paid your earnings in cash. After you have been paid, you will be free to leave.

Thank You

D.2 Handout to central planners: Case R_2D_U

Experiment Instructions

Treatment 1: 2 Regional Managers, a Central Planner, Market Uncertainty

Handout to the Central Planner

This is an information sharing inventory game

You are part of a supply chain of 3 people: **2 regional managers and 1 central planner**. You are assigned the role of the central planner. In this task, you will make decisions for up to **40 independent rounds (5 trial and up to 35 game rounds)**. In each round, you will play with 2 other participants in this room (2 regional managers). At the beginning of each round players will be **randomly and anonymously** assigned to a supply chain group. Your role will always be the central planner of the supply chain group but you will not be assigned to the same group in consecutive rounds.

Setting

You are the central planner of a company that owns a network of 2 regional stores. Each store sells the same seasonal products (e.g., holiday decorations). You are responsible for deciding on the amount of inventory to purchase for each product in the stores' collection for a given season. For the purpose of this exercise we will focus on one product. To help inform your decision, each store manager provides you with a demand forecast several months before the selling season. You may use this information to determine the companys total order quantity for the product. The total order quantity of the product is held in the company's single warehouse and shipped to the regional stores after demand is realized.

The timing of events is as follows:

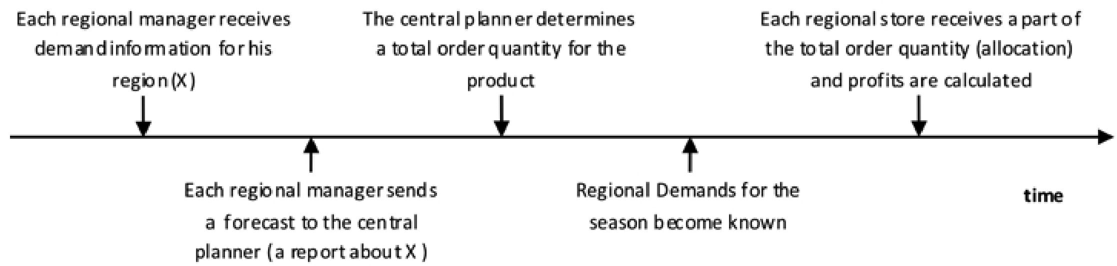

Regional Demand

The product's demand for a given season within one retail region is determined by the sum of two components:

Regional Demand = X + Y

X is randomly generated in each round for each retail region. Each regional manager observes the exact value of X for his (and only his) region. The other regional manager and you (the central planner) only know that X for any retail region is equally likely to be **any integer value between 100 and 400.**

Y is also randomly generated in each round for each retail region. The regional managers and the central planner only know that Y for any retail region is equally likely to be **any integer value between -50 and 50.**

The values of X and Y in each round are independent. They are also independent across retail regions and rounds. Thus, if X and/or Y are large (or small) for a region in the current round, this will not affect whether they are large (or small) in another region or in future rounds.

Regional managers' decision

As described above, each regional manager has better information about his regional demand because he observes the exact value of his X.

Each regional manager must decide what information to provide to you (the central planner), about X. This information takes the form of a **report of X** for the round.

Please note that as long as the report is between 100 and 400 units, you cannot verify whether a regional manager reported his private information truthfully or not.

The objective of each regional manager is to maximize **the local profit of his store**. For each unit a store sells, the store earns $2. For each unit a store receives from the companys warehouse, it pays $1. The profit at each store is calculated as follows:

Local Profit = $2 * min [Regional Demand, Allocation] $1 * Allocation

A store's allocation is based on the total order quantity (your decision) and the ratio of the store's regional demand versus total demand across all regions:

Allocation = (Regional Demand / Total Demand)*Total Order Quantity

Total Demand stands for the sum of the product demand across all store regions. Please note that with this allocation rule it is possible that a regional store receives a quantity below/above its demand. No inventory is left at the central warehouse at the end of the season.

Example of the allocation rule

*Suppose the final regional demands at the two stores are 270 and 330 units respectively. If your total order quantity was 620 units, stores will be allocated (270/(270+330))*620= 279 units, and (330/(270+330))*620= 341 units respectively (allocation is above their demand). If, instead, the total order quantity had been 580 units, their allocation would be 261 and 319 units respectively(allocation is below their demand).*

<u>Your decision</u>

Your task it to decide the **total order quantity** to purchase for the product. This quantity will be held in the company's single warehouse as common inventory for the 2 regions until regional demands are realized. The product costs $1 per unit to purchase from a wholesaler. No additional costs are incurred in holding or shipping the product.

Your profit

Being the central planner, **you will be rewarded based on the companys total** profit that is the sum of the local profits of each store. Companys profit is defined as the revenue earned across all store regions minus the cost of acquiring the product:

Companys Profit = $2 * min [Total Demand, Total Order Quantity] $1 * Total Order Quantity

How you will be paid

If you follow these instructions and make good decisions you will earn a considerable amount of money that will be paid to you in cash at the end of the session.

At the end of the experiment, your total earnings from the game rounds (the first 5 trial rounds do not affect your profits) will be divided by 1400 to determine your earnings from the experiment. These dollar earnings will be added to your $10 participation fee and displayed on your computer screen. The maximum amount of dollars you may receive is $20.

When you have completed all rounds of the game, you will be asked to fill out the Post-Game survey. Then the researcher will pass by your carrel desk and you will be paid your earnings in cash. After you have been paid, you will be free to leave.

Thank You

D.3 Summary sheets

Summary sheet for regional managers

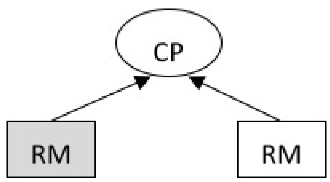

Summary

Period 0: Common Information

Regional Demand = X + Y

X is equally likely to be any integer value between 100 and 400.

Y is equally likely to be any integer value between -50 and 50.

Total Demand = Regional Demand 1 + Regional Demand 2

Timing of Events

Period 1: Each Regional Manager (RM) learns the exact value of his (and only his) X.

Your decision: send to the Central Planner a report of X for your region.

Period 2: The Central Planner (CP) sees the reports and places a Total Order Quantity.

Period 3: Ys are revealed (hence Regional Demands and Total Demand are known).

Allocations are determined. Your allocation is:

Allocation = (Your Regional Demand / Total Demand)*Total Order Quantity

Profits are calculated. Your profit (=Local Profit)

Local Profit = $2* min [Regional Demand, Allocation] $1 * Allocation

Summary sheet for central planners

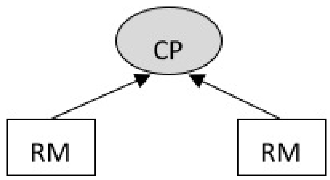

Summary

Period 0: Common Information

Regional Demand = X + Y

X is equally likely to be any integer value between 100 and 400.

Y is equally likely to be any integer value between -50 and 50.

Total Demand = Regional Demand 1 + Regional Demand 2

Timing of Events

Period 1: Each Regional Manager (RM) learns the exact value of his (and only his) X. Each RM sends to you, the Central Planner (CP), a report of X for his region.

Period 2: Your decision:You see these reports and place a Total Order Quantity.

Period 3: Ys are revealed (hence Regional Demands and Total Demand are known). Allocations are determined and profits are calculated.

Your profit (=Company's Profit)

Company's Profit = $2 * min [Total Demand, Total Order Quantity] $1 * Total Order Quantity

Appendix E

Snapshots from the Experiment Software

E.1 Snapshots of a retailer's computer screen: Case R_2D_U

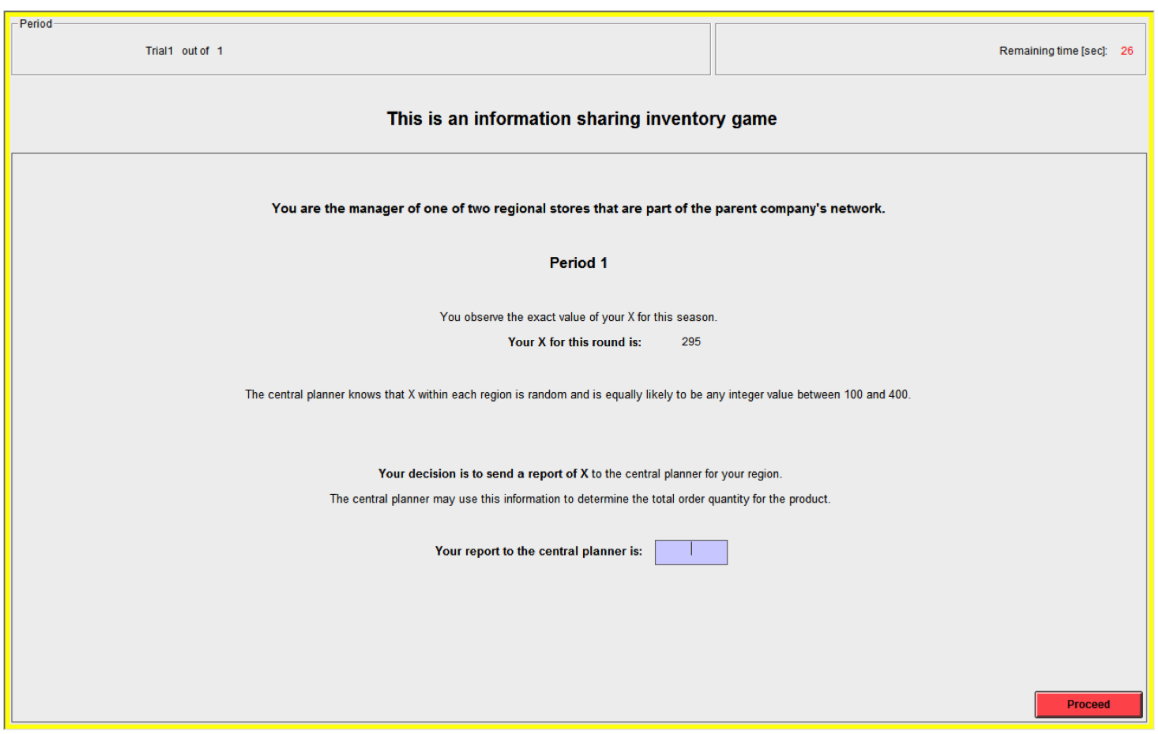

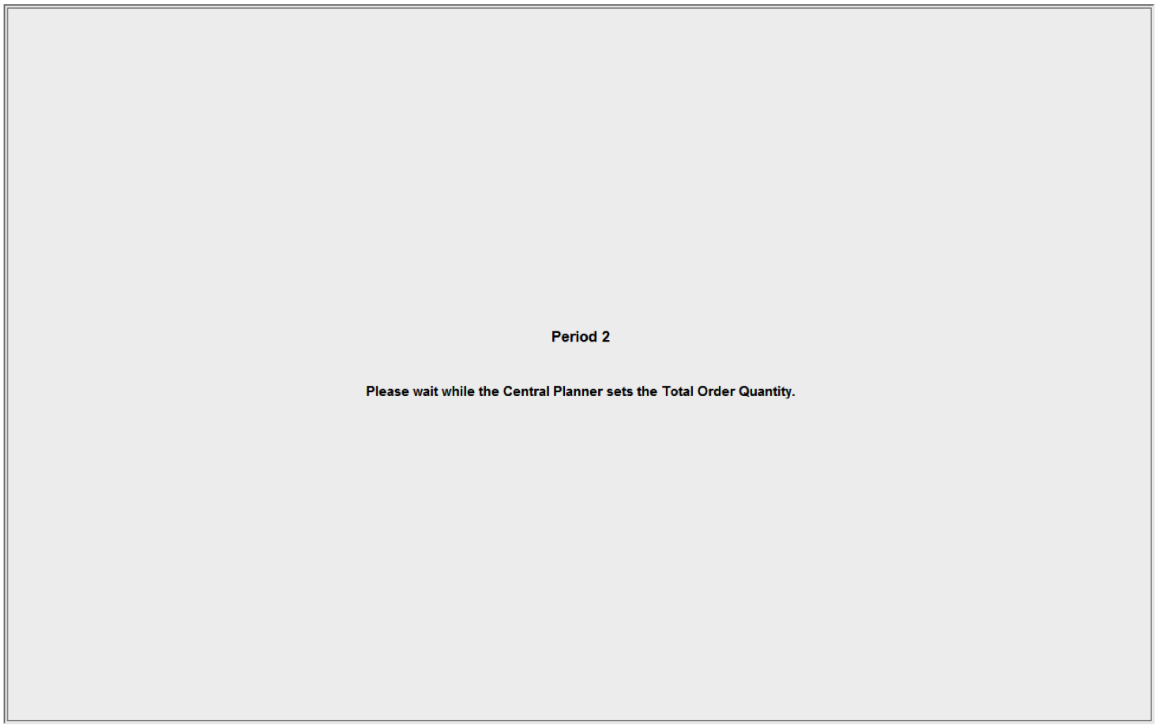
Period 2
Please wait while the Central Planner sets the Total Order Quantity.

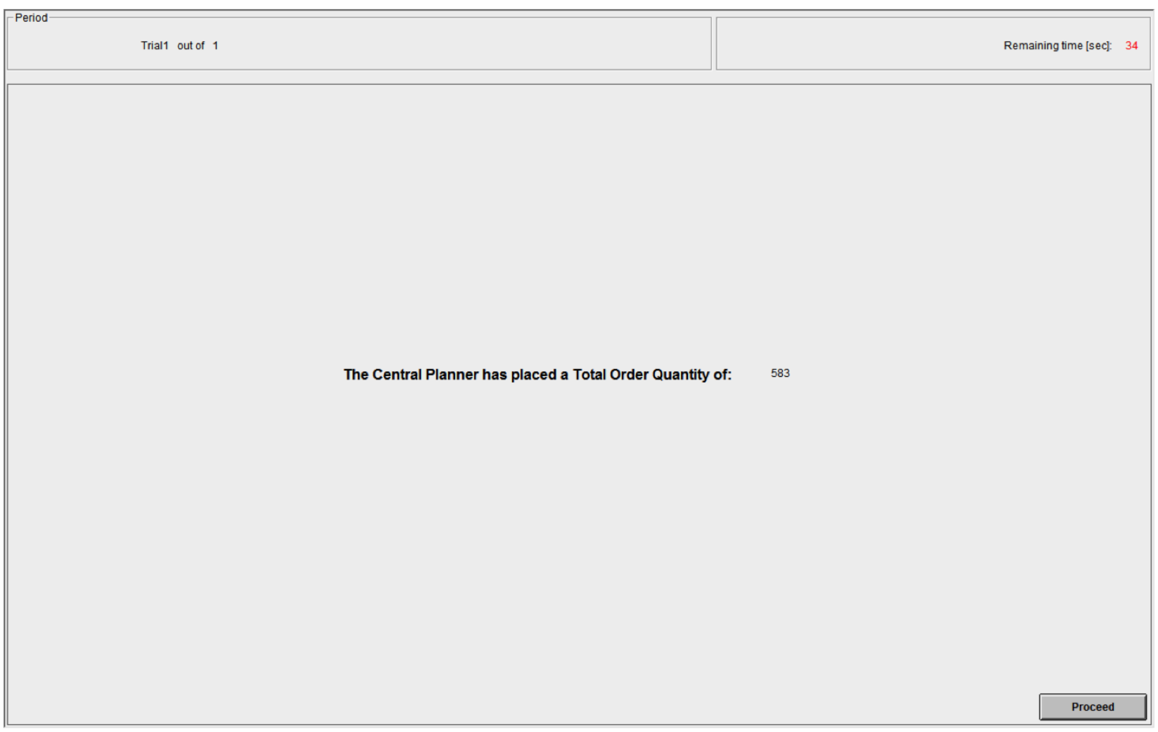
Period
Trial1 out of 1
Remaining time [sec]: 34
The Central Planner has placed a Total Order Quantity of: 583
Proceed

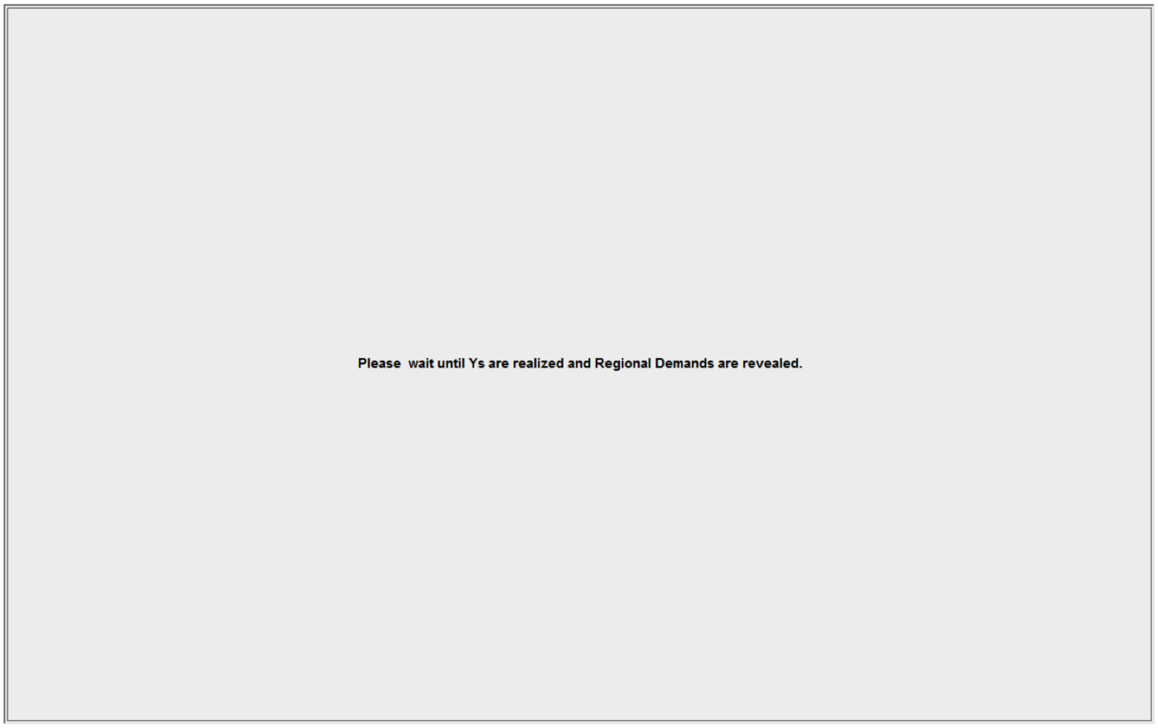
Please wait until Ys are realized and Regional Demands are revealed.

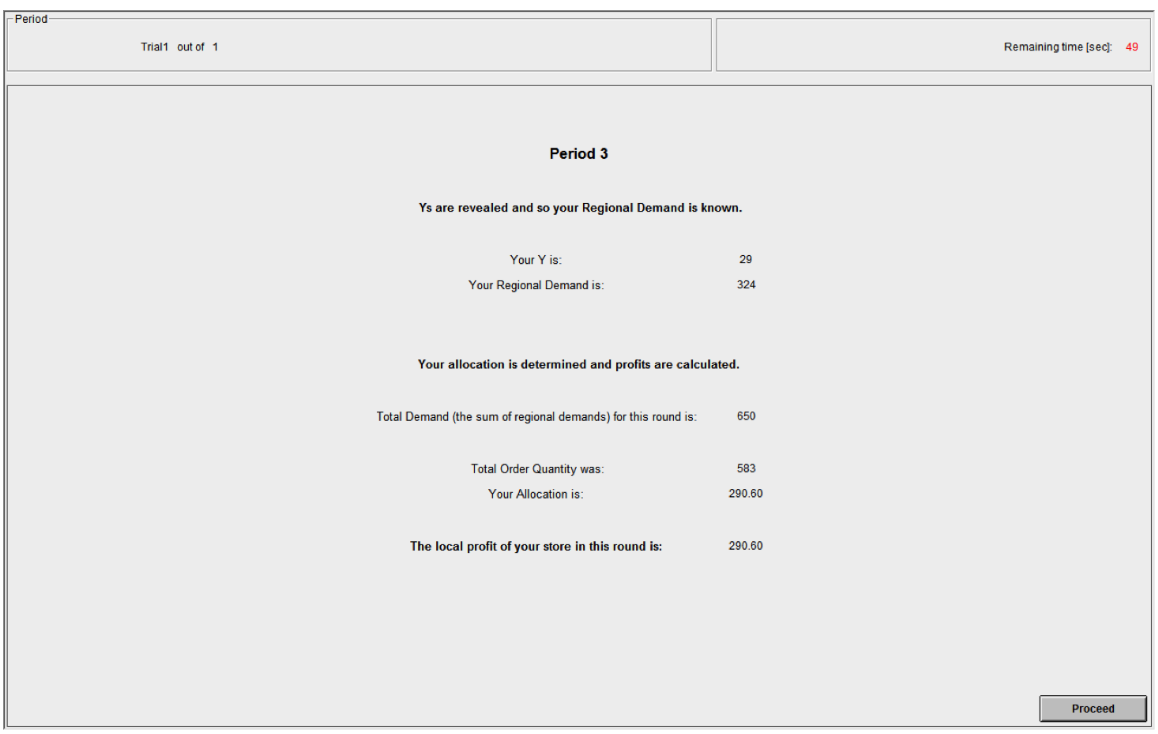
Period
Trial1 out of 1
Remaining time [sec]: 49
Period 3
Ys are revealed and so your Regional Demand is known.
Your Y is: 29
Your Regional Demand is: 324
Your allocation is determined and profits are calculated.
Total Demand (the sum of regional demands) for this round is: 650
Total Order Quantity was: 583
Your Allocation is: 290.60
The local profit of your store in this round is: 290.60
Proceed

E.2 Snapshots of a central planner's computer screen: Case R_2D_U

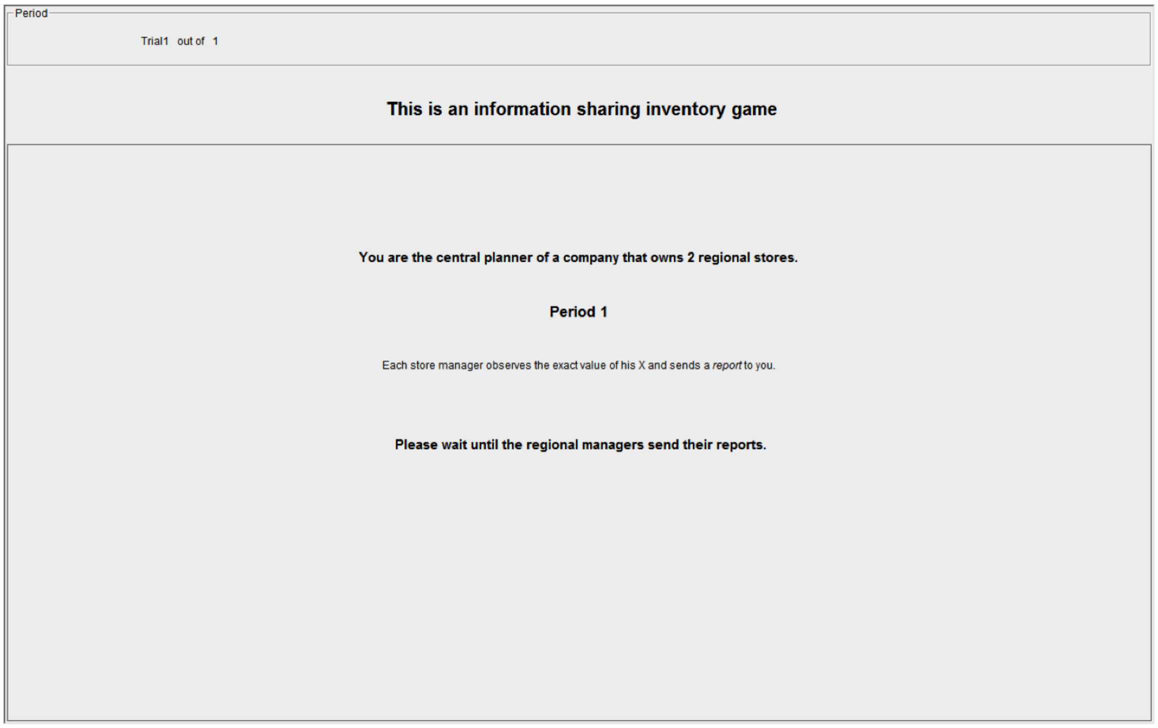


Period
Trial1 out of 1
Remaining time [sec]: 29
Period 2
Recall that X for each region is random and is equally likely to be any integer value between 100 and 400.
For this round, you receive the following reports of X for each of the two regional stores.
The report from the 1st regional manager is: 295
The report from the 2nd regional manager is: 288
Your task is to decide on the Total Order Quantity to place.
If the reports are the true Xs, the total order quantity that maximizes the expected profit of the company is: 583
Your Total Order Quantity is:
This quantity will be held at the company's warehouse as common inventory for the two regions, until regional demands are realized.
Proceed

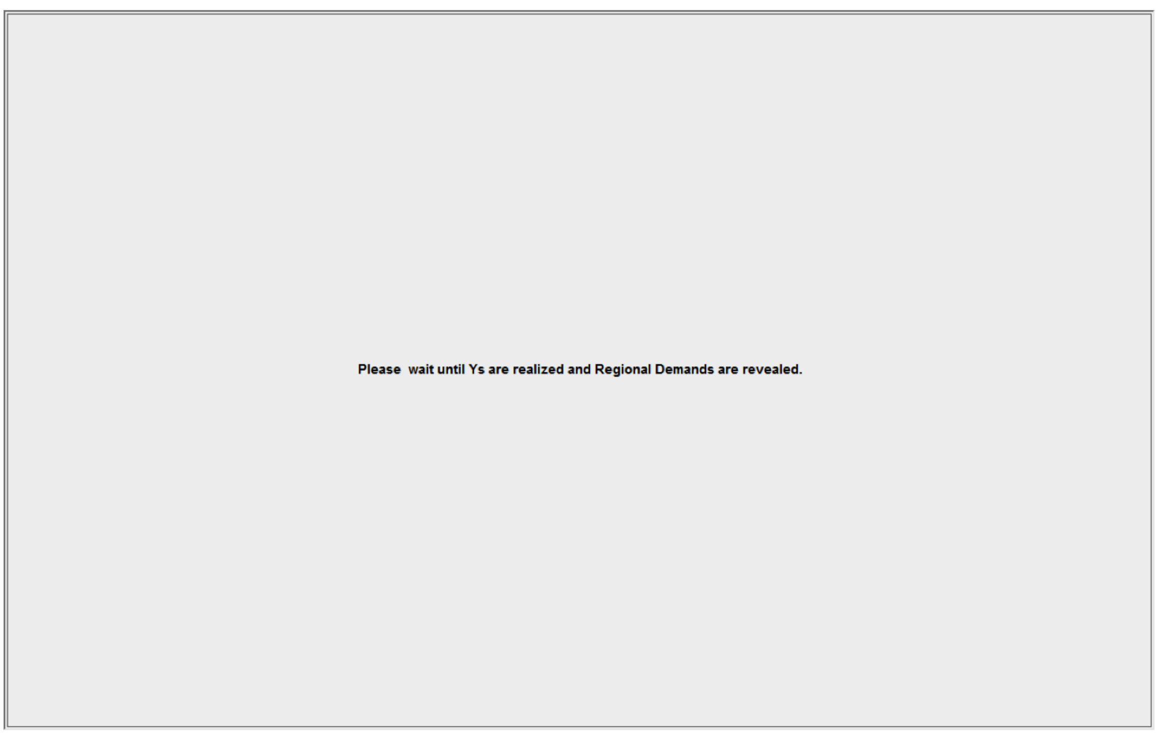
Please wait until Ys are realized and Regional Demands are revealed.

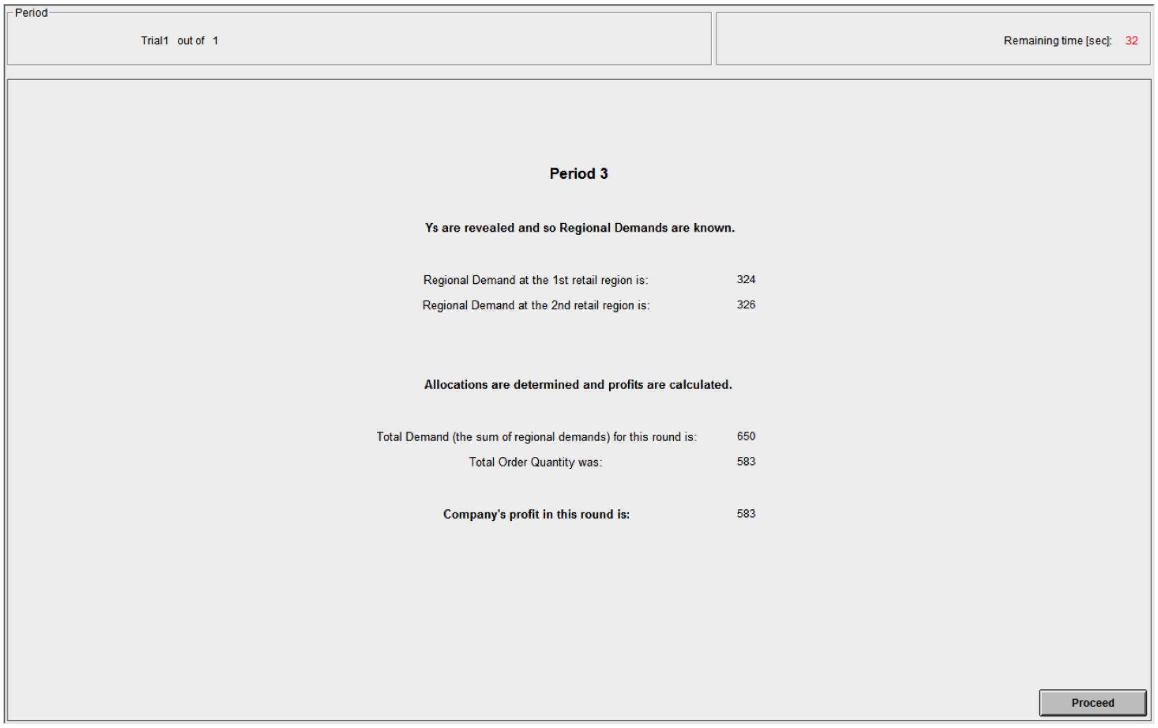
Period
Trial1 out of 1
Remaining time [sec]: 32
Period 3
Ys are revealed and so Regional Demands are known.
Regional Demand at the 1st retail region is: 324
Regional Demand at the 2nd retail region is: 326
Allocations are determined and profits are calculated.
Total Demand (the sum of regional demands) for this round is: 650
Total Order Quantity was: 583
Company's profit in this round is: 583
Proceed

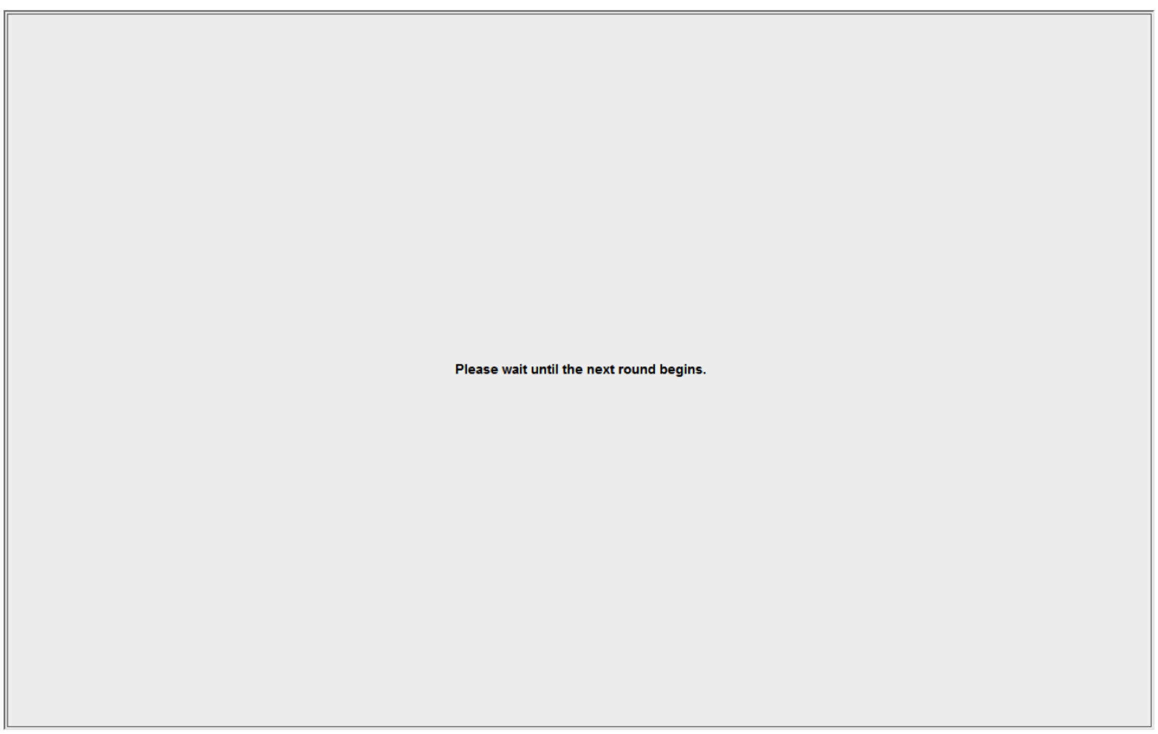
Please wait until the next round begins.

Bibliography

[1] K. S. Anand and H. Mendelson. Information and organization for horizontal multi-market coordination. *Management Science*, 43(12):1609–1627, 1997.

[2] R. Anupindi. A General Framework for the Study of Decentralized Distribution Systems. *Manufacturing & Service Operations Management*, 9(2):309–368, Sept. 2001.

[3] E. Bendoly, R. Croson, P. Goncalves, and K. Schultz. Bodies of knowledge for research in Behavioral Operations. *Proceedings of the Behavioral and Quantitative Game Theory: Conference on Future Directions*, 19(4):434–452, 2010.

[4] U. Benzion, Y. Cohen, R. Peled, and T. Shavit. Decision-making and the newsvendor problem: an experimental study. *Journal of the Operational Research Society*, 59(9):1281–1287, July 2007.

[5] G. E. Bolton and E. Katok. Learning by Doing in the Newsvendor Problem: A Laboratory Investigation of the Role of Experience and Feedback. *Manufacturing & Service Operations Management*, 10(3):519–538, 2008.

[6] A. A. Bostian, C. A. Holt, S. Jain, and K. Ramdas. Is Transshipment a Behaviorally-Robust Risk-Pooling Strategy? 2012.

[7] A. A. Bostian, C. A. Holt, and A. M. Smith. Newsvendor "Pull-to-Center" Effect: Adaptive Learning in a Laboratory Experiment. *Manufacturing & Service Operations Management*, 10(4):590–608, Sept. 2008.

[8] G. P. Cachon and M. A. Lariviere. An equilibrium analysis of linear, proportional and uniform allocation of scarce capacity. *IEE Transactions*, 31:835–849, 1999.

[9] G. P. Cachon and M. A. Lariviere. Capacity Allocation Using Past Sales: When to Turn-and-Earn. *Management Science*, 45(5):685–703, May 1999.

[10] G. P. Cachon and M. A. Lariviere. Capacity choice and allocation: strategic behavior and supply chain performance. *Management Science*, 45(8):1091–1108, 1999.

[11] G. P. Cachon and M. A. Lariviere. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science*, 47(5):629–646, 2001.

[12] F. Chen. Information sharing and supply chain coordination. *Handbooks in operations research and management science*, 11:341–421, 2003.

[13] Y. Chen, X. Su, and X. Zhao. Modeling Bounded Rationality in Capacity Allocation Games with the Quantal Response Equilibrium. *Management Science*, 58(10):1952–1962, June 2012.

[14] R. B. Cialdini and N. J. Goldstein. Social influence: compliance and conformity. *Annual review of psychology*, 55(1974):591–621, Jan. 2004.

[15] C. J. Corbett and K. Rajaram. A Generalization of the Inventory Pooling Effect to Nonnormal Dependent Demand. *Manufacturing Service Operations Management*, 8(4):351–358, Jan. 2006.

[16] V. P. Crawford and J. Sobel. Strategic Information Transmission. *Econometrica*, 50(6):1431–1451, 1982.

[17] K. Donohue and E. Siemsen. Behavioral Operations: Applications in Supply Chain Management. In *Wiley Encyclopedia of Operations Research and Management Science*, chapter Behavioral. 2010.

[18] G. D. Eppen. Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem. *Management Science*, 25(5):498–501, May 1979.

[19] G. D. Eppen and L. Schrage. Centralized ordering policies in a multi-warehouse system with lead times and random demand. *Multi-level production/inventory control systems: Theory and practice*, 16:51–67, 1981.

[20] U. Fischbacher. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, June 2007.

[21] D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

[22] Y. Gerchak and D. Gupta. On apportioning costs to customers in centralized continuous review inventory systems. *Journal of Operations Management*, 10(4):546–551, 1991.

[23] Y. Gerchak, R. G. Vickson, and M. Parlar. Periodic Review Production Models With Variable Yield And Uncertain Demand. *IIE transactions*, 20(2):144–150, 1988.

[24] P. Glasserman. Perturbation analysis of production networks. In *Stochastic Modeling and Analysis of Manufacturing Systems*, pages 233–280. Springer, 1994.

[25] W. H. Greene. *Econometric analysis, 5th edition*. Prentice Hall, Upper Saddle River, NJ, 2003.

[26] B. C. Hartman, M. Dror, and M. Shaked. Cores of Inventory Centralization Games. *Games and Economic Behavior*, 31(1):26–49, Apr. 2000.

[27] M. Henig and Y. Gerchak. The structure of periodic review policies in the presence of random yield. *Operations Research*, 38(4):634–643, 1990.

[28] T.-H. Ho, N. Lim, and T. H. Cui. Reference Dependence in Multilocation Newsvendor Models: A Structural Analysis. *Management Science*, 56(11):1891–1910, Oct. 2010.

[29] K. Inderfurth, A. Sadrieh, and G. Voigt. The Impact of Information Sharing on Supply Chain Performance under Asymmetric Information. *Production and Operations Management*, 22(2):410–425, Mar. 2013.

[30] H. Jiang, S. Netessine, and S. Savin. Robust newsvendor competition under asymmetric information. *Operations research*, 59(1):254–261, 2010.

[31] E. Kemahlioglu-Ziya and J. J. I. Bartholdi. Centralizing Inventory in Supply Chains by Using Shapley Value to Allocate the Profits. *Manufacturing & Service Operations Management*, 13(2):146–162, Oct. 2011.

[32] M. Kremer. *Pooling Inventory Risk - The Efficacy of Excess Inventory Markets*. PhD thesis, University of Mannheim, 2007.

[33] J. A. Lavaro and C. J. Corbett. The value of SKU rationalization in practice (the pooling effect under suboptimal inventory policies and nonnormal demand)*. *Production and Operations Management*, 12(1):12–29, 2003.

[34] L. Li. Information Sharing in a Supply Chain with Horizontal Competition. *Management Science*, 48(9):1196–1212, Sept. 2002.

[35] L. Li, I. Duenyas, and S. M. Iravani. Capacity Option Transfer Rights: Do They Benefit Suppliers ? 2011.

[36] L. Li and H. Zhang. Confidentiality and Information Sharing in Supply Chain Coordination. *Management Science*, 54(8):1467–1481, Aug. 2008.

[37] L. X. Lu and M. A. Lariviere. Capacity allocation over a long horizon: the return on turn-and-earn. *Proceedings of the Behavioral and Quantitative Game Theory: Conference on Future Directions*, page 65, 2010.

[38] L. X. Lu and M. A. Lariviere. Capacity Allocation over a Long Horizon: The Return on Turn-and-Earn. *Manufacturing & Service Operations Management*, 14(1):24–41, 2012.

[39] S. Mullainathan and R. H. Thaler. Behavioral economics, 2000.

[40] A. Muller, M. Scarsini, and M. Shaked. The Newsvendor Game Has a Nonempty Core. *Games and Economic Behavior*, 38(1):118–126, Jan. 2002.

[41] S. Netessine and N. Rudi. Notes centralized competitive inventory models with demand substitution. *Operations Research*, 51(2):329–335, 2003.

[42] O. Ozer, Y. Zheng, and K.-Y. Chen. Trust in Forecast Information Sharing. *Management Science*, 57(6):1111–1137, Apr. 2011.

[43] O. Ozer, Y. Zheng, and Y. Ren. Trust, Trustworthiness, and Information Sharing in Supply Chains Bridging China and the U.S. 2012.

[44] Z. J. Ren, M. A. Cohen, T. H. Ho, and C. Terwiesch. Information Sharing in a Long-Term Supply Chain Relationship: The Role of Customer Review Strategy. *Operations Research*, 58(1):81–93, Nov. 2010.

[45] L. W. Robinson. A comment on Gerchak and Guptas On apportioning to customers in centralized continuous review inventory systems. *Journal of Operations Management*, 11:99–102, 1993.

[46] N. Rudi, S. Kapur, and D. F. Pyke. A Two-Location Inventory Model with Transshipment and Local Decision Making. *Management Science*, 47(12):1668–1680, 2001.

[47] M. E. Schweitzer and G. P. Cachon. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):404–420, 2000.

[48] L. S. Shapley. A Value for n-person Games. In *Contributions to the Theory of Games, Vol.2*, pages 307–317. Princeton University Press, Princeton, NJ, 1953.

[49] W. Shih. Optimal inventory policies when stockouts result from defective products. *International Journal of Production Research*, 18(6):677–686, 1980.

[50] E. A. Silver. Establishing the order quantity when the amount received is uncertain. *INFOR*, 14:32–39, 1976.

[51] J. Sobel. Signaling Games. Encyclopedia of Complexity and System Science. Springer, 2008.

[52] Y. Sprumont. The Division Problem with Single-Peaked Preferences: A Characterization of the Uniform Allocation Rule. *Econometrica*, 59(2):509, Mar. 1991.

[53] X. Su. Bounded Rationality in Newsvendor Models. *Manufacturing & Service Operations Management*, 10(4):566–589, 2008.

[54] C. Terwiesch, Z. J. Ren, T. H. Ho, and M. A. Cohen. An Empirical Analysis of Forecast Sharing in the Semiconductor Equipment Supply Chain. *Management Science*, 51(2):208–220, Feb. 2005.

[55] R. H. Thaler. Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior and Organization*, 1:39–60, 1980.

[56] A. Tversky and D. Kahneman. Judgment under Uncertainty: Heuristics and Biases. *Science (New York, N.Y.)*, 185(4157):1124–31, Sept. 1974.

[57] G. Voigt and K. Inderfurth. Supply chain coordination with information sharing in the presence of trust and trustworthiness. *IIE Transactions*, 44(8):637–654, Aug. 2012.

[58] H. Wu and M. Parlar. Games with incomplete information: A simplified exposition with inventory management applications. *International Journal of Production Economics*, 133(2):562–577, Oct. 2011.

[59] C. A. Yano and H. L. Lee. Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334, 1995.

[60] H. Zhang. Vertical information exchange in a supply chain with duopoly retailers. *Production and Operations Management*, 11(4):531–546, Jan. 2002.