

**Supply Chain Transparency and Social
Responsibility: Investigating Consumer and Firm
Perspectives**

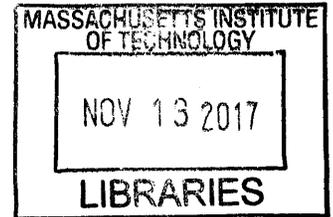
by
Leon Valdes

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Management
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2017

© Massachusetts Institute of Technology 2017. All rights reserved.



ARCHIVES

Signature redacted

Author

Sloan School of Management
August 29, 2017

Signature redacted

Certified by

Yanchong Zheng
Sloan School Career Development Professor
Associate Professor of Operations Management
Thesis Supervisor

Signature redacted

Accepted by

Catherine Tucker
Chair, Sloan School of Management PhD Program



77 Massachusetts Avenue
Cambridge, MA 02139
<http://libraries.mit.edu/ask>

DISCLAIMER NOTICE

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available.

Thank you.

The images contained in this document are of the best quality available.

Supply Chain Transparency and Social Responsibility: Investigating Consumer and Firm Perspectives

by

Leon Valdes

Submitted to the Sloan School of Management
on August 29, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Management

Abstract

Consumers increasingly expect companies to ensure that their products are made in a socially responsible manner. However, most companies do not have good visibility into their supply chains. According to a recent study, 81% of the 1,700 companies surveyed lacked full visibility into the social responsibility practices of their suppliers. Using incentivized laboratory experiments and a game-theoretic model, in this thesis we study how improved transparency about social responsibility practices in the supply chain can positively impact companies' interactions with both consumers and suppliers.

In the first part of this thesis, we design an incentivized laboratory experiment to study two key questions: (i) How does visibility impact consumers' valuations of social responsibility practices in a supply chain? (ii) What roles do indirect reciprocity and prosociality play in affecting consumers' valuations under different levels of visibility? Our results demonstrate that consumers are willing to pay more for greater visibility. Also, high prosocial consumers do not exhibit indirect reciprocity. Conversely, indirect reciprocity increases low prosocial consumers' valuations under high visibility.

In the second part, we study how a manufacturer can improve a supplier's social responsibility practices under incomplete visibility. We consider a game-theoretic model with information asymmetry about the supplier's practices and focus on the manufacturer's investment in the supplier's capabilities. We also consider the potential disclosure of social responsibility information to consumers by the manufacturer or a third party. We find that the manufacturer should invest a high (low) amount of resources in the supplier's capabilities if the information it observes suggests poor (good) practices. Greater visibility helps the manufacturer be more efficient with this investment. The disclosure of social responsibility information by the manufacturer leads to better supplier's practices.

Finally, we conduct an incentivized laboratory experiment to investigate (i) how does visibility affect consumers' trust in companies' communications? (ii) How does visibility impact the effect that trust has on consumers' willingness-to-pay for products? Our results show that the effect of visibility on trust is highly dependent on

consumers' prosociality. In particular, only low prosocial consumers trust companies more when they demonstrate greater visibility – and this translates into a greater willingness-to-pay.

Thesis Supervisor: Yanchong Zheng

Title: *Sloan School Career Development Professor*

Associate Professor of Operations Management

Acknowledgments

First and foremost, I would like to thank my advisors, Karen Zheng and Tim Kraft, for all their support and guidance during my Ph.D. Their dedication as professors, researchers, and advisors was invaluable and it provided me with a great example for my future academic career. I am particularly grateful for their support in conducting research in a topic that I am passionate about, as well as for always challenging me to give my best.

I would also like to thank Retsef Levi and Georgia Perakis for serving in my Thesis Committee and for their insightful questions and comments. They also provided key advice during my first year in the Ph.D. program and during my academic job search. Similarly, I thank all the faculty that I had the opportunity to interact with at MIT, especially Rob Freund – I had the pleasure of being a teaching assistant for his class during my last semester at Sloan. Many thanks also to the faculty and students from the B-Lab; I learned a lot about behavioral research from them.

There are a number of other teachers and professors who have inspired me to always keep learning and who motivated me (directly or indirectly) to want to pursue a career in academia. The list is long, but I would like to especially thank Rene Caldentey, Jose Miguel Cruz, and Rafael Epstein from my undergrad at Universidad de Chile; Libio Baranda and Rafael Cautivo from my high school years; and Sergio Contreras and Hugo Lagos.

My research would not have been possible without the generous help from many people at the UVA and MIT behavioral labs. I want to thank in particular Charles Holt, Jennifer Walker, and Orn Wawsri. Similarly, my years at MIT would not have been the same without the help and support of Hillary Ross, Sarah Massey, Davin Lee, and Tyler Morse.

My years as a Ph.D. student would not have been the same without the unconditional friendship of Francisco, Sarah, and Jorge. I am extremely grateful for having met them during my first weeks in Boston and for sharing these years together – they were my family away from home. I also made great friends in the ORC and OM

programs, whose friendship and support I will always cherish. In particular, many thanks to Yaron and Efrat, Fernanda, and Gonzalo.

I am also grateful to my family, in general, and my parents and siblings, in particular. Thank you for your unconditional love and support, both before and during my Ph.D. For the past six years I have not always been as present as I would have liked to be, but you have always been there for me. My visits to Chile and your visits to Boston always made me extremely happy. I also want to thank Cheryl and Mark for their love and support during these last years.

Finally, and most importantly, I want to thank Naomi. I could not have done this without her and I cannot think of a better partner to share this journey with. She has taken care of me in so many ways during these years, constantly reminding me of what is (and isn't) important, and always pushing me to become a better person. I look forward to our next adventure together and to us continuing to support and take care of each other in the many years to come.

Contents

1	Introduction	15
2	Supply Chain Visibility and Social Responsibility: Investigating Consumers' Behaviors and Motives	19
2.1	Introduction	19
2.2	Experimental Design and Procedure	24
2.2.1	The Consumer Purchase Game: Base Setup	25
2.2.2	The Consumer Purchase Game: Treatment Design	29
2.2.3	Control Tasks and Postexperiment Survey	30
2.2.4	Experimental Procedure	32
2.3	Hypotheses	34
2.4	Experimental Results	38
2.4.1	Hypothesis 2.1: Does Consumers' WTP Increase with the Level of Visibility?	40
2.4.2	Hypothesis 2.2: Does Reciprocity Exist at Each Level of Visibility?	43
2.4.3	The Impact of Consumers' Prosociality	46
2.5	The Product Choice Study	51
2.6	Conclusions and Managerial Insights	53
3	Improving Supplier Social Responsibility under Incomplete Visibility	57
3.1	Introduction	57

3.2	Model Setup	63
3.2.1	No Manufacturer Disclosure	63
3.2.2	Manufacturer Disclosure	68
3.3	Results: No Manufacturer Disclosure	70
3.3.1	The Effect of Visibility on the Manufacturer’s Optimal Investment Strategy and the Supplier’s Expected Final SR Level	74
3.3.2	The Effect of Third-Party Scrutiny on the Manufacturer’s Optimal Investment Strategy and the Supplier’s Final SR Level	77
3.4	Results: Manufacturer Disclosure	78
3.4.1	What Level of SR Should the Manufacturer Disclose?	80
3.4.2	The Effects of Visibility and Third-Party Scrutiny on the Manufacturer’s Disclosure	82
3.5	When Does the Manufacturer Prefer to Voluntarily Disclose SR Information?	85
3.5.1	The Effects of Visibility and Third-Party Scrutiny on the Manufacturer’s Disclosure Choice	89
3.5.2	The Effect of the Manufacturer’s Disclosure Choice on the Final SR Level	90
3.6	Conclusions and Managerial Insights	91
4	The Effect of Visibility on Consumer Trust of Social Responsibility Disclosure	95
4.1	Introduction	95
4.2	Experimental Design and Procedure	100
4.2.1	The Consumer Trust Game: Base Setup	100
4.2.2	The Consumer Trust Game: Experimental Manipulations	104
4.2.3	Control Task and Postexperiment Survey	105
4.2.4	Experimental Procedures	106
4.3	Hypotheses	108
4.4	Experimental Results	112

4.4.1	Hypothesis 4.1: Does Visibility affect Consumers' Trust? . . .	113
4.4.2	Hypothesis 4.2: Is the Effect of Trust on Consumers' WTP Affected by Visibility?	114
4.4.3	Discussion	119
4.5	Conclusions and Managerial Insights	122
5	Conclusions	125
A	Additional Material Chapter 2	129
A.1	Regression Estimates and Marginal Effects of Main Results	129
A.2	Robustness Experiments	129
A.2.1	Implementing the Strategy Method with a Random-Order Pre- sentation	129
A.2.2	Minimizing Efficiency Considerations	131
A.2.3	Varying the Wealth Status of the Disadvantaged Party	131
A.2.4	Analysis on Alternative Classifications of Prosocial Types	135
A.3	Supplementary Analyses	136
A.3.1	The Effects of Demographics on Consumers' WTP	137
A.3.2	Analysis on Order Effects in the CPG	138
A.3.3	Visibility Results for the Random Condition	140
A.4	The Role of Self-Serving Bias in Affecting Consumers' WTP	140
B	Additional Material Chapter 3	143
B.1	Additional Analytical Results	143
B.2	Examining the Expected Final SR Level	145
B.3	The Effect of q on β_L and β_H	147
C	Proofs Chapter 3	149
C.1	The Supplier's Best Response	150
C.2	The Manufacturer's Optimal Strategy	151
C.3	Comparative Statics	175

C.4	When Does the Manufacturer Prefer to Voluntarily Disclose SR Information?	185
C.5	The Effects on Social Responsibility	192

List of Figures

2-1	Summary of Experimental Design	27
2-2	Effect of Visibility on Consumers' WTP	40
2-3	Effect of Reciprocity on Consumers' WTP	44
2-4	Effect of Reciprocity on High Prosocial Consumers' WTP	48
2-5	Effect of Reciprocity on Low Prosocial Consumers' WTP	48
3-1	Supplier's Optimal SR Decision for Different Values of β	72
3-2	Manufacturer's Optimal Investment in the Supplier's SR Capabilities	73
3-3	Comparing the Manufacturer's Optimal Investment in the Supplier's SR Capabilities	79
3-4	Optimal Disclosed SR Level by the Manufacturer	83
3-5	Manufacturer's Optimal Disclosure Strategy	88
4-1	Relationship between the true value of Y and the clue	101
4-2	Summary of Experimental Design	102
4-3	Effect of Experimental Conditions among low prosocial Consumers when $Msg = 0$ or $Clue = 0$	118
A-1	Summary of Results in the Variation with No Provisional Endowment	132
A-2	Consumers' Preferences in the Product Choice Study	138
B-1	Effect of Visibility on the Expected Final Level of SR	146

List of Tables

2.1	Mean, [Median], and (Standard Deviation) of Consumers' WTP	38
2.2	Marginal Effects of V_H , V_M , and Their Differences, Δ_V	41
2.3	Marginal Effects of $D \cdot e_k$	45
3.1	Notation	70
4.1	Percentage of observations where the Consumer's belief is that Clue = Y_{Low}	113
4.2	Summary Statistics of Consumers' WTP in the CTG	115
4.3	Regression Estimates: Equation (4.1)	117
4.4	Summary of Experimental Results	119
A.1	Regression Results: Equation (2.1)	130
A.2	Marginal Effects of V_H , V_M , and Their Differences, Δ_V by Prosocial Type	130
A.3	Marginal Effects of $D \cdot e_k$ by Prosocial Type	130
A.4	Regression Estimates: No Provisional Endowment	132
A.5	Marginal Effects of V_H and $D \cdot e_k$ (by Visibility Condition): No Provi- sional Endowment	133
A.6	Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Con- dition): Excluding Consumers Stating WTP > 0 for $e = 0$	133
A.7	Marginal Effects of V_H : Worker-50 & Worker-80 by Consumer's Self- Serving Bias	133

A.8	Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Equation (2.1) When Low Prosocial Consumers Are Those Who Transfer Exactly 0 in the Dictator Game	133
A.9	Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Equation (2.1) When Prosocial Types Are Classified Based on the Attitudinal Survey	137
A.10	Regression Estimates: Controls for Demographics	139
A.11	Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Controls for Demographics	139
A.12	Regression Estimates: Including Round as an Independent Variable	140
A.13	Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Including Round as an Independent Variable	140
A.14	Marginal Effects of V_H , V_M , and Δ_V for the Random Condition	141
A.15	Regression Estimates and Marginal Effects for Consumers Who Exhibit a Self-Serving Bias and Those Who Do Not	141

Chapter 1

Introduction

Consumers increasingly want and expect companies to ensure that their products are made in a socially responsible manner. According to a 2015 survey of more than 30,000 consumers in 13 countries, 66% of respondents said they were willing to pay more for products from sustainable brands; this represents a significant increase from 50% of respondents in 2013 (Nielsen 2015). Yet most companies cannot fully observe the social responsibility (SR) practices in their supply chains. A recent study by The Sustainability Consortium found that 81% of the 1,700 companies surveyed lacked full visibility into the SR practices of their suppliers (The Sustainability Consortium 2016). This is due to in large part to the fact that improving visibility into a supply chain can be very costly and time-consuming (Doorey 2011). For example, in 2014 companies spent “roughly \$709 million and six million staff hours...to comply with rules to disclose conflict minerals in their supply chains.”¹

In this context, the goal of this thesis is to investigate whether and how companies can benefit from increased supply chain transparency about SR practices. Our insights provide companies with guidance on whether improving transparency is worth the significant investment of time and resources. Employing incentivized laboratory experiments and a game-theoretic model, we study how supply chain transparency can affect companies’ interactions with both consumers and suppliers. In particular, we address the following main research questions: (i) Are consumers willing to pay

¹<https://www.wsj.com/articles/u-s-firms-struggle-to-trace-conflict-minerals-1438636575>

more for a product with more precise information about the SR practices behind it (Chapter 2)? (ii) How can companies improve their suppliers' SR practices and communicate such practices to consumers with only incomplete SR information (Chapter 3)? And (iii) does greater transparency yield a higher level of consumer trust in companies' SR communications (Chapter 4)? In addition, since consumers' valuations of a company's SR practices vary widely (e.g., García-Gallego and Georgantzís 2011), we incorporate consumers' heterogeneity when addressing each of these questions. Throughout this thesis, we focus on the *social* dimension of SR; i.e., working conditions, safe work practices, living wages, etc.

To create a transparent supply chain, a company must (i) gain visibility into its supply chain and (ii) determine what information to disclose to consumers (New and Brown 2011). To clearly delineate these two dimensions of transparency, we refer to them as “visibility” and “disclosure”. In Chapter 2, we focus on visibility and investigate when companies can benefit from it. Specifically, we design an incentivized three-player laboratory experiment to study the following questions: (i) How does supply chain visibility affect consumers' valuations of a company's SR practices in its supply chain? (ii) Are consumers motivated by indirect reciprocity (i.e., rewarding the firm for the responsible treatment of its workers) in their valuations? (iii) How do visibility and indirect reciprocity jointly influence the valuations of consumers based on their prosociality (i.e., their willingness to sacrifice own benefit to improve someone else's well-being)? We choose to focus on visibility – as opposed to disclosure – because it is an equally-important yet understudied dimension of transparency. Furthermore, we incentivize participants' decisions to elicit their *truthful* valuations of SR practices. Our results show that consumers are willing to pay more for greater visibility into the SR practices of a company's supply chain. We also observe a three-way interaction between visibility, indirect reciprocity, and prosociality: high prosocial consumers do not exhibit strong indirect reciprocity, while it significantly increases low prosocial consumers' valuations – but only under a high level of visibility.

In Chapter 3, we turn our attention to the interaction between a manufacturer and a supplier to study how the manufacturer (she) can improve the supplier's (he)

SR practices when these practices cannot be perfectly observed. To capture incomplete visibility, we consider a game-theoretic model where the supplier's initial level of SR is his private information. The manufacturer only observes a signal about this level; how likely the signal captures the supplier's true SR practices depends on the manufacturer's visibility into her supply chain. With this setup, we focus on the manufacturer's investment to improve the supplier's SR capabilities; for example, by training the supplier on topics like health and safety practices. Despite being a highly effective method to improve supplier's SR performance (Porteous et al. 2015), developing a supplier's capabilities remains underutilized in practice (Gillai et al. 2013). To capture how consumer demand can influence the manufacturer's and the supplier's decisions, we also incorporate the potential for SR information to be disclosed to the consumers either directly by the manufacturer or by an external third party (e.g., an NGO). We derive several interesting results from this model. First, the manufacturer should invest a high (low) amount of resources to improve the supplier's capabilities if the information she observes suggests the supplier's current SR practices are poor (good). With greater visibility into her supply chain, the manufacturer can better identify whether the supplier requires significant support. Second, the manufacturer is most likely to prefer not disclosing SR information to consumers when the supplier appears to have only average SR practices. In addition, SR practices in the supply chain are strictly better when the manufacturer discloses than when she does not disclose. Finally, with greater visibility the manufacturer becomes more "truthful" in her disclosure of SR information; i.e., she is less likely to either overstate or understate the supplier's SR performance given her available information.

Finally, in Chapter 4 we study whether supply chain visibility impacts consumers' trust in companies' disclosure of SR information. Consumer trust is key to the success of companies' SR communication strategies: According to a recent survey, brand trust tops the list of factors that influence SR purchases (Nielsen 2015). In this context, transparency has been identified in recent studies as one of the most effective ways for companies to improve trust (e.g., KPMG and Global Reporting Initiative 2014, Marks & Spencer and Globescan 2015). However, these studies focus on the effect of

disclosure on consumer trust, while the effect of visibility remains not well understood. In this chapter we help close this gap by studying: (i) How does visibility affect consumers' trust in companies' SR communications? (ii) How does visibility impact the effect that trust has on consumers' willingness-to-pay for companies' products? (iii) How do the answers to these questions vary based on consumers' prosociality? To address these questions, we design an incentivized three-player laboratory experiment (different from the one in Chapter 2) where a player in the role of the Firm observes a signal about the working conditions in its supply chain. As in Chapter 3, we use the accuracy of this signal to capture the Firm's level of visibility. The Firm then conveys a message – which may or may not be truthful – about this signal to the player in the role of the Consumer. Our results demonstrate that the effect that visibility has on trust highly depends on consumers' prosociality. In particular, low prosocial consumers trust companies' communications more when companies have greater visibility into the SR practices in their supply chain. This increase in trust is also accompanied by an increase in these consumers' willingness-to-pay.

Taken together, the results in this thesis identify several ways in which supply chain transparency can benefit companies. On the supplier side, greater visibility allows a company to better tailor its investment in a supplier's SR capabilities. Similarly, it helps the company to better trust its available information about its supplier's SR practices, which translates into a more truthful disclosure of SR information to consumers. On the consumer side, we find that companies can benefit from greater visibility as consumers are willing to pay more for more precise information about SR practices in the supply chain. We also find that at least some consumers trust companies' SR disclosure more when the company demonstrates a greater level of supply chain visibility. Finally, our results indicate that greater transparency often translates into an improvement of the level of SR in the supply chain.

Chapter 2

Supply Chain Visibility and Social Responsibility: Investigating Consumers' Behaviors and Motives

2.1 Introduction

Creating transparency about the social responsibility (SR) practices in their supply chains is an emerging business challenge for companies. As Patagonia's Director of Environmental Strategy states, "transparency...is really becoming an expectation now. People want to know more about the supply chains making the products they're buying" (Patagonia 2014). However, many companies are at a crossroads in determining how transparent they want to make their supply chains (Marshall et al. 2016). For example, Marks & Spencer reports that it is working with "customers and stakeholders to identify what information they consider to be important about where and how M&S products are produced" and "will respond by improving the information available" by 2020 (Marks & Spencer 2015a, p. 12). To create a transparent supply chain, a company must: (i) gain visibility into its supply chain, and (ii) determine what information to disclose to consumers (New and Brown 2011). In particular, gaining visibility requires a significant investment of time and resources (Doorey 2011). As

such, a critical question arises for companies debating how transparent to make their supply chains – do consumers value greater supply chain visibility?

In this paper, we employ an incentivized human-subject experiment to examine when and how *supply chain visibility* impacts consumers' valuations of a company's SR practices in its upstream supply chain.¹ To clearly delineate the two dimensions of transparency, we refer to "visibility" as the extent to which a company has information about the SR practices in its supply chain and "disclosure" as a company's decision regarding what information to communicate to consumers. We design our experiment to focus on visibility while controlling for disclosure (with information being fully disclosed to consumers). We choose to focus on visibility because having visibility into its supply chain is a prerequisite for a company to be able to disclose SR information to its consumers. In practice, many companies do not have good visibility into their supply chains. For example, a recent study by The Sustainability Consortium found that 81% of the 1,700 companies surveyed lacked full visibility into the SR practices of their supply chains and 54% had no visibility at all (The Sustainability Consortium 2016). Furthermore, the current experimental literature on transparency and SR primarily studies the disclosure of information to consumers (e.g., Rode et al. 2008, Hainmueller et al. 2015), while the dimension of visibility is understudied. We address this gap by identifying when there is a revenue benefit to greater visibility, and thus, providing a market incentive for companies to gain visibility into their supply chains.

To enhance our understanding of consumers' valuations, we design our experiment to investigate the behavioral motives underlying consumers' decisions. Various social preferences can motivate a consumer to value SR, such as altruism (e.g., Levine 1998, Andreoni and Miller 2002), inequality aversion (e.g., Fehr and Schmidt 1999, Bolton and Ockenfels 2000), and reciprocity (e.g., Nowak and Sigmund 1998, Charness and Haruvy 2002). We focus on the preference of *indirect reciprocity*. Following Alexander (1987, p. 5), indirect reciprocity is defined as "the return from a social investment in

¹We follow the European Commission's definition of social responsibility as "[companies integrating] social and environmental concerns in their business operations and in their interaction with their stakeholders on a voluntary basis" (Dahlsrud 2008). We specifically focus on "social" concerns in social responsibility.

another ... from someone other than the recipient of the beneficence.” Thus, indirect reciprocity arises when individuals help those who have helped others. Indirect reciprocity, like direct reciprocity, is an intention-based social preference (Charness and Rabin 2002, Falk and Fischbacher 2006). That is, how much X reciprocates Y for Y’s treatment of Z depends on both the resulting outcome experienced by Z and that Y actively chose the corresponding treatment. In an SR context, consumers motivated by indirect reciprocity care about both the outcomes of a company’s SR practices and how much effort the company exerts to obtain such outcomes. Thus, these consumers would be willing to reward a company for its active engagement in improving practices in its supply chain (e.g., ensuring the proper treatment of workers). Reciprocity has been suggested as an important driver of consumers’ valuations for SR (e.g., Trudel and Cotte 2009, Nielsen 2013). We study when consumers’ behaviors can be attributed to indirect reciprocity, and how this relationship is affected by the level of supply chain visibility.

A third dimension we study in our experiment considers consumers’ heterogeneity in their tendency to care about SR. Examining consumer heterogeneity is important because consumers’ valuations of a company’s SR practices can vary widely (e.g., García-Gallego and Georgantzis 2011). We focus on one particular aspect of heterogeneity – consumers’ *prosociality*. We define an individual’s prosociality as the extent to which the individual is willing to sacrifice his/her own benefit to improve the payoff of another. Prior studies have attributed the consumption of socially responsible products to values such as benevolence, equality, and responsibility (Vermeir and Verbeke 2006). Individuals who adhere to these values tend to be more prosocial. Hence, there is a natural connection between a person’s prosociality and his/her attention to SR.

Based on the above three aspects, we answer the following research questions: (i) How does *supply chain visibility* affect consumers’ valuations of a company’s SR practices in its upstream supply chain? (ii) Are consumers motivated by *indirect reciprocity* in their valuations? (iii) How do visibility and indirect reciprocity jointly influence the valuations of consumers based on their *prosociality*? To address these

questions, we design a three-player game with the following social context: A Worker has helped a Firm to make a product, and the Firm would like to sell the product to a Consumer. In our design, we represent SR by the Firm’s effort to improve the Worker’s payoff and operationalize this effort through monetary transfer from the Firm to the Worker. To capture our focus on SR in an upstream supply chain, we design the players’ initial endowments such that the Worker is a disadvantaged party (in wealth status) compared to the other two players. The Consumer’s valuation of SR is measured by his willingness-to-pay (WTP) for the product given the Firm’s effort to improve the Worker’s payoff. We incentivize Consumers’ decisions to elicit their *truthful* WTP. Employing an incentivized experiment is important because evidence exists that what individuals claim they would do with respect to SR differs from their actual purchase behavior. For example, Devinney et al. (2010, p. 112) observe that “individuals either purposely overstate their social credentials or just want to look good in surveys, making it nearly impossible to believe what they say about their social proclivities.”

We operationalize visibility in our experiment by varying the extent to which the *outcome* of the Firm’s SR effort is precisely known to the Firm, and hence, the Consumer (since information is fully disclosed).² A supply chain with high visibility means that both the Firm and the Consumer observe the outcome of the Firm’s effort to increase the Worker’s payoff. Conversely, in a supply chain with low visibility, even the Firm is uncertain of the outcome of its effort. In practice, although companies have full knowledge about their own efforts, they do not necessarily have the same level of knowledge about their suppliers’ practices (Lee et al. 2012). Greater visibility can help inform companies about such practices, hence reducing uncertainty about SR in their supply chains; e.g., Hewlett-Packard monitors its suppliers on a monthly basis to better assess the SR performance of its supply chain (HP 2016).

Our research contributes to two growing streams of literature. First, regarding

²As an example to delineate outcome versus effort, Starbucks’ 2014 Annual Report states (Starbucks 2014, p. 12): “Starbucks worked with nonprofit organizations ... to contribute more than 520,000 hours of volunteer service around the world ... Altogether, the projects benefitted an estimated 1.4 million people with a value of \$5.2 million for our communities.” The first sentence highlights Starbucks’ efforts, and the second sentence demonstrates the outcomes.

consumers' valuations of SR, an emerging body of experimental studies examine the value of companies being transparent to consumers. These studies have primarily focused on the dimension of disclosure while either implicitly assuming (e.g., Kimeldorf et al. 2006, Hainmueller et al. 2015) or explicitly allowing (e.g., Bartling et al. 2015, Pigors and Rockenbach 2016a) companies to have *full* visibility. For example, Pigors and Rockenbach (2016a) study consumers' purchase decisions in a laboratory experiment where the worker's wage is set by and known to the manager, but may or may not be known to consumers. Hainmueller et al. (2015) conduct a field experiment to study consumers' willingness to pay a premium for coffee with a fair trade label. In this setting, it is implicitly assumed that the company has full visibility into the ethical aspects of the product. We differ from these works by studying an equally-important but understudied dimension of transparency – a company's visibility into its supply chain. As such, we control for the disclosure of information (by assuming full disclosure) while manipulating the level of visibility in our experiment. By doing so, we address when and how consumers' valuations of a company's SR practices depend on the level of supply chain visibility.

Second, within the operations management literature, there is a growing interest in studying how the type and nature of the information presented to individuals affect their perceptions and behavior. For example, Buell and Norton (2011) and Buell et al. (2016) examine how showing information about operational processes may increase customers' perceptions of service value and operational efficiency. Muthulingam et al. (2013) show that the order in which energy-saving recommendations are presented to small and medium firms significantly affects the adoption rates of the recommendations. Roels and Su (2013) show that what type of comparison information a social planner should present to households regarding their energy consumption (e.g., versus neighbors' consumption) depends on whether individuals are ahead-seeking or behind-averse. Kalkan et al. (2016) compare mandatory versus voluntary disclosure of a company's social and environmental impacts and show that voluntary disclosure is perceived more positively by consumers. We differ from these works by studying a new aspect of information – uncertainty about a company's SR practices in its sup-

ply chain. Using a controlled experiment we develop a micro-level understanding of consumers' behaviors and motives in their valuations of SR under different levels of information uncertainty. Furthermore, we underscore the importance of accounting for consumers' heterogeneity when studying their valuations.

We highlight three key results. First, we show that consumers' valuations increase with a higher level of supply chain visibility. This is especially true when consumers use uncertainty as a justification not to pay for SR (i.e., they exhibit a self-serving bias). From a managerial standpoint, this finding highlights a potential revenue benefit to greater supply chain visibility, and thus, identifies a market incentive for companies to create more transparent supply chains. To further strengthen the applicability of this insight, we design and administer a product choice study that considers broader SR contexts and multiple product categories. Our results from this study confirm the positive value of greater visibility.

Second, we show how uncertainty in the outcome of one's action affects indirect reciprocity. Specifically, we observe that indirect reciprocity has a significant positive effect on consumers' valuations only under high visibility. Third, we further examine consumers' behaviors by taking into account their prosociality. We observe that high prosocial consumers are outcome-driven and do not exhibit indirect reciprocity. Conversely, low prosocial consumers' valuations are significantly affected by indirect reciprocity. For example, indirect reciprocity increases low prosocial consumers' valuations when the level of visibility is high. Because consumers motivated by indirect reciprocity are influenced by the knowledge of a company's SR efforts, these findings offer valuable insights into when information regarding these efforts resonates with the company's target consumers.

2.2 Experimental Design and Procedure

Our experimental design consists of three parts: the Consumer Purchase Game (CPG), two control tasks to elicit individuals' prosociality and risk preferences, and a postexperiment survey. The design of each part is grounded in well-established

methods in the experimental and behavioral sciences literatures. In the CPG, we use monetary transfer from the Firm to the Worker to operationalize SR. It is common in the literature for researchers to use simple economic games with monetary transfers to study social preferences and to derive implications for applications beyond monetary transactions (e.g., Ostrom et al. 1992, Kosfeld et al. 2009). Recently, Bartling et al. (2015) and Pigors and Rockenbach (2016a) also use monetary transfers in laboratory experiments to operationalize socially responsible production. Utilizing controlled laboratory experiments to study consumer behavior allows us to design the cleanest possible environment to study our factors of interest (i.e., visibility, indirect reciprocity, and prosociality) through incentivized decisions, while controlling for other confounding factors (Kagel and Roth 1995, Katok 2011). We next discuss our design and the experimental procedure in detail.

2.2.1 The Consumer Purchase Game: Base Setup

We design the CPG as a three-player game with the following roles: a Firm (she), a Consumer (he), and a Worker. At the beginning of the game, the three players are given a background story that the Worker has helped the Firm to make a (hypothetical) product, and the Firm wants to sell the product to the Consumer. The Firm, the Consumer, and the Worker are initially endowed with 160, 120, and 20 tokens. The Firm receives an additional *provisional* payment of 120 tokens and earns this provisional payment only if she manages to sell the product. We design the players' initial endowments to capture wealth differences among a company, consumers, and workers in a supply chain. Because our focus is on SR in the upstream supply chain, we establish the presence of a disadvantaged party, i.e., the Worker player. Such a design is motivated by examples of workers in upstream supply chains (often in developing countries) living in poor economic conditions, e.g., workers in China, Vietnam, and Bangladesh in the electronics and apparel industries. Within the literature, utilizing a zero or lower initial endowment to capture a disadvantaged party is common for ultimatum and dictator games (see Oosterbeek et al. 2004, Engel 2011 for reviews). To minimize pre-experiment biases related to specific practices or products, we call

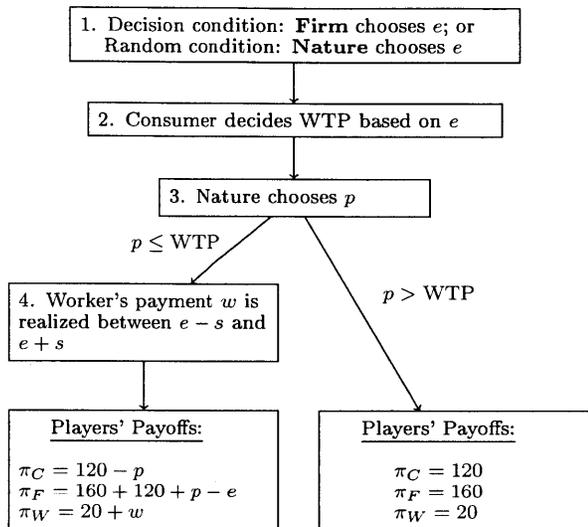
the above three roles Players A, B, and C in the experiment and refer to a generic product with no specific features.

The CPG proceeds as follows. First, the Firm decides how much from the provisional 120 tokens she is willing to use to generate a payment to the Worker. Let e denote the Firm's decision. The Firm can choose one of seven possible values in the set $\{0, 20, 40, 60, 80, 100, 120\}$. This decision models a company's *effort* to improve SR practices in its supply chain. Second, the Consumer states the maximum price that he is willing to pay to the Firm for the product, given the Firm's decision. We call the Consumer's decision his *willingness-to-pay* (WTP). The Worker does not make any decisions. If the product is sold, then the Worker receives a payment w that depends on e .

After the Firm and the Consumer make their decisions, Nature chooses the product price, p , from a uniform distribution between 1 and 120 tokens. If p is lower than or equal to the Consumer's WTP, then the product is sold at price p . In this case, the final payoffs to all players are: (i) the Consumer pays the product price to the Firm and earns $\pi_C = 120 - p$; (ii) the Firm receives the provisional 120 tokens plus the product price minus her decision, earning $\pi_F = 160 + 120 + p - e$; and (iii) the Worker receives the payment and earns $\pi_W = 20 + w$. If instead, the price p is strictly higher than the Consumer's WTP, then the product is not sold and the players earn their initial endowments: $\pi_C = 120$, $\pi_F = 160$, and $\pi_W = 20$. Figure 2-1 (Decision condition) summarizes the above dynamics. The game dynamics and all players' payoff structures are common knowledge.

In the CPG, we use a random price to elicit the Consumer's truthful WTP. This approach is known as the Becker-DeGroot-Marschak mechanism and is commonly used in economic experiments (e.g., Becker et al. 1964, Klos et al. 2005, Halevy 2007). We also apply the strategy method (e.g., Fehr and Fischbacher 2004, Falk et al. 2008) to obtain the Consumer's decision: the Consumer states his WTP for each possible value of e that the Firm may choose *while* the Firm is choosing e (i.e., before knowing the actual value of e). We conducted robustness treatments to verify that using the strategy method does not prime Consumers to state their WTP in a monotonic fashion

Figure 2-1: Summary of Experimental Design



(see Appendix A.2.1). Note that the players' final payoffs are jointly determined by the value of e actually chosen by the Firm and the corresponding WTP stated by the Consumer.

REMARKS: We explain further here the motivation behind a few key design features of our experiment. First, in our design we provide a background story that reflects a supply chain context while using context-neutral terminologies to describe the game dynamics. A long-running debate exists between experimental economists and psychologists regarding whether an experimental design should abstract away all specific contexts or preserve contextual cues (e.g., Goldstein and Weber 1995, Ortman and Gigerenzer 2000). We choose a middle ground to leverage the merits of both approaches. This differentiates our design from other context-neutral economic games and ensures that we study fundamental preferences in the specific context of interest (i.e., SR in an upstream supply chain).

Second, in our design the Consumer's WTP captures the *premium* that he is willing to pay for a product from a supply chain with SR investment versus a product from a supply chain without SR investment. We do not include a surplus from the product for the Consumer so as to isolate the Consumer's motivation to buy purely due to wanting to help the Worker versus wanting to also earn a surplus. This design

allows us to measure the Consumer's WTP specifically for SR in a clean manner.

Third, the payment to the Worker if the product is sold, w , *does not* represent the Worker's wage for his work on the product (the wage is captured by his initial endowment of 20 tokens). Instead, w captures the *improved treatment* of the Worker due to the Firm's SR effort, e . The dynamic between the Firm and the Worker in our design is motivated by examples of companies such as Hewlett-Packard and Starbucks investing to improve SR practices in their upstream supply chains.³ For example, Starbucks has invested more than \$100 million to support coffee communities along its supply chain, including Coffee and Farmer Equity practices, farmer support centers, and other programs aimed at improving farmers' livelihoods. Mapping our design to this example, w corresponds to the improved livelihood of the farmers, and e corresponds to Starbucks' investment of over \$100 million. Rode et al. (2008) and Bartling et al. (2015) are among the recent experimental papers that also operationalize SR with an additional payment being made to a third party *only if* a product is sold. For example, Bartling et al. (2015) model non-socially responsible production creating a (long-term) negative externality on the worker with a negative w . We conversely model socially responsible production having a (long-term) positive externality on the worker.

Finally, the provisional 120 tokens that the Firm receives if the product is sold is included in our design to ensure that (i) the Firm is motivated to sell the product and (ii) the Consumer's WTP is primarily motivated by his intention to help the Worker. The latter combined with the Consumer not receiving a surplus from the product helps to ensure that we are truly measuring the Consumer's WTP for SR. We conducted robustness treatments with a variation of the CPG in which the additional 120 tokens for the Firm were not provisional but instead included in her initial endowment; i.e., the Firm's initial endowment was 280 tokens. We confirm that our results regarding the effects of visibility and indirect reciprocity on Consumers' WTP continue to hold (see Appendix A.2.2).

³See <http://www8.hp.com/us/en/hp-information/global-citizenship/society/supplychain.html> and <https://www.starbucks.com/responsibility/sourcing/coffee>.

2.2.2 The Consumer Purchase Game: Treatment Design

We manipulate the CPG in two dimensions in our experiment. First, to study the effect of *visibility* on the Consumer's WTP, we manipulate the relationship between the Firm's decision e and the payment to the Worker w . Specifically, w is uniformly distributed on $[e - s, e + s]$ if $e > 0$, and $w = 0$ if $e = 0$. The parameter s models the level of visibility. We examine three different values of s , leading to three Visibility conditions: $s = 0$ (*High Visibility*), $s = 10$ (*Medium Visibility*), and $s = 20$ (*Low Visibility*). The condition of $s = 0$ represents high visibility because the payment to the Worker if the product is sold is exactly equal to the Firm's decision e . Conversely, $s = 20$ represents low visibility because given e , there is still significant uncertainty regarding the actual payment to the Worker. Note that by our design, in all three Visibility conditions (i) $w \geq 0$ always holds for all possible values of e , and (ii) the expected payment to the Worker, if the product is sold, is equal to e . We manipulate the level of visibility as an exogenous treatment condition. This approach is in line with recent experiments that study the disclosure of SR information to consumers, where disclosure, or the lack thereof, is often varied as an exogenous treatment condition (e.g., Rode et al. 2008, Bartling et al. 2015).

To focus on visibility while controlling for the effect of disclosure on the Consumer's WTP, we design the CPG such that all information available to the Firm is disclosed to the Consumer; i.e., the Consumer has the same information as the Firm, including knowing the Firm's effort and facing the same level of visibility. This information structure is common knowledge for all three players. We operationalize different levels of visibility by varying the uncertainty faced by the Firm (and thus the Consumer) about the *outcome* (i.e., the payment w to the Worker) of the Firm's effort. Lower visibility results in less precise information being observed and thus, higher uncertainty about SR outcomes in the supply chain. Capturing lower visibility with higher uncertainty in the observed outcomes is a common approach in the literature (e.g., Bloomfield and Wilks 2000, Lang and Maffett 2011).

The second dimension we manipulate in our experiment is the process by which the

decision e is selected. Doing so allows us to study the effect of *indirect reciprocity* on the Consumer’s WTP. We compare two Selection conditions: the *Decision* condition versus the *Random* condition (see Figure 2-1). In the Decision condition, the Firm chooses e as discussed in §2.2.1. In contrast, in the Random condition, Nature chooses e uniformly and randomly from the feasible set $\{0, 20, 40, 60, 80, 100, 120\}$, and the Firm automatically accepts the chosen value. Since the Firm does not actively choose e in the Random condition, a higher or lower value of e in this condition cannot be interpreted by the Consumer as the Firm being more or less responsible for the Worker. That is, the Consumer’s WTP in the Random condition is driven by preferences other than indirect reciprocity. Conversely, the Consumer’s WTP in the Decision condition is driven by both these preferences *and* indirect reciprocity. Thus, differences observed in WTP between the two conditions measure indirect reciprocity. This approach is well established in the literature for studying direct reciprocity (e.g., Blount 1995, Charness and Haruvy 2002). We extend it to study indirect reciprocity. The three Visibility and two Selection conditions combine to yield a 3×2 factorial design (6 treatments). This design allows us to study the individual effects of visibility and indirect reciprocity on WTP, as well as their interaction effect.

2.2.3 Control Tasks and Postexperiment Survey

We include two control tasks in our experiment to measure Consumers’ prosociality and risk preferences. We study the Consumer’s prosociality because it naturally connects with motives that induce a person to care about SR. We examine the Consumer’s risk preference because a person’s risk attitude could affect his/her decision when the decision has uncertain implications on others’ payoffs (e.g., Chakravarty et al. 2011, Andersson et al. 2016) – in our context, the Consumer’s WTP decision influences the Worker’s uncertain payoff.

To measure the Consumer’s prosociality, we employ a variation of the dictator game (Forsythe et al. 1994) as follows. The dictator and the recipient (called Players 1 and 2 in the task) are initially endowed with 120 and 20 tokens. The Consumer from the CPG plays the dictator role and is asked to choose the number of tokens, a , from

the set $\{0, 20, 40, 60, 80, 100, 120\}$ that he is willing to use to generate a payment to the recipient. After the Consumer makes his decision, the recipient receives a random payment t that is uniformly distributed on $[a - s, a + s]$ if $a > 0$, and exactly equal to zero if $a = 0$. The value of s in the dictator game is set equal to the value of s that the Consumer faced in the CPG. Matching these two values ensures that we measure the Consumer's prosociality under the same level of uncertainty regarding outcomes.⁴ The Consumer's decision a as the dictator measures his prosociality. A higher value of a implies that the Consumer is more willing to improve the recipient's payoff at his own cost. Hence, he is more prosocial. We analyze whether visibility and indirect reciprocity have different effects on the WTP decisions made by Consumers with high versus low prosociality.

To measure the Consumer's risk preference, we use the multiple price list method (Holt and Laury 2002). Participants are presented with ten pairs of lotteries. For each pair, they are asked to choose which lottery, A or B, they would prefer to play. For lottery A (B), the two possible outcomes are 60 and 74 (6 and 140) tokens. The expected values of both lotteries increase as one goes down the list, with A (B) having a higher expected value in pairs 1–4 (5–10). By design, a participant is expected to either always prefer B, or prefer A up to a certain pair and prefer B thereafter. Risk-neutral participants (who maximize their expected payoffs) will choose A for pairs 1–4 and B for pairs 5–10. Thus, a participant is considered risk averse if he/she switches from A to B at pair 6 or later; otherwise he/she is considered as non-risk averse. To incentivize their choices, participants are informed that one pair of lotteries will be randomly selected and played by the computer at the end of the session. They earn additional income based on the outcome of the lottery they choose in the selected pair.

After participants finish the CPG and the control tasks, they complete a post-experiment survey. The survey has four parts: (i) questions about participants'

⁴If a participant first experienced low or medium visibility in the CPG and later played the dictator game with no uncertainty, then this difference could affect his/her dictator decision. We show in §2.4.3 that having different levels of uncertainty in the dictator game does not substantially affect the participants' decisions across treatments.

decisions in the CPG, (ii) a product choice study, (iii) an attitudinal survey, and (iv) a demographic survey. Part (i) is used to obtain insights about participants' decisions in the game. In part (ii), we use a simplified conjoint design (Feinberg et al. 2012) to examine how our results in the game may apply to broader SR topics and product categories. We discuss the design and analysis of this part in detail in §2.5. Part (iii) contains a series of binary-choice and 5-point Likert-scale questions to help us better understand participants' general attitudes toward SR and how these attitudes relate to their everyday purchase decisions. The participants' responses provide additional data to further measure their prosociality beyond the dictator game described above. Finally, part (iv) is used to obtain demographic and background information about our participants.

2.2.4 Experimental Procedure

We ran all experimental sessions in the computer laboratories of two large universities. The sessions encompassed a 3 (Visibility: High vs. Medium vs. Low) \times 2 (Selection: Decision vs. Random) factorial design. All 6 treatments were run at both universities. All sessions followed the same procedure with participants completing the following tasks in the same order: (i) two rounds of the CPG, (ii) the dictator game and risk preference tasks, and (iii) the postexperiment survey. We provided participants with the details of the next task only after they completed the previous one, including informing them that there would be a second and final round of the CPG *only after they completed round 1*. We used this procedure to prevent participants from altering their current decisions in anticipation of future decisions. Participants did not observe the outcomes of the tasks until they had completed all tasks. This procedure eliminated the possibility that outcomes from previous tasks affected current and future decisions, including income effects and changes in generosity due to past outcomes. Both methods are commonly used in the experimental literature (e.g., Bolton and Katok 1998, Luhan et al. 2009).

At the beginning of round 1 of the CPG, participants were randomly and anonymously matched into groups of three. Within each group, participants were randomly

assigned to one of three roles: Player A (Firm), Player B (Consumer), or Player C (Worker). Participants were informed of the treatment conditions they were in and played the CPG as described in §2.2.1. After participants completed round 1 of the game, they were told that they would be playing the CPG for a second and final time. The computer performed the following role reassignments for round 2. The Consumer participants in round 1 were assigned the role of the Firm in round 2. The Firm participants in round 1 were assigned the role of the Consumer in round 2. The Worker participants in round 1 remained in that role in round 2. Thus, two thirds of all participants played the role of the Consumer once and the role of the Firm once (in two different rounds); one third of all participants played the role of the Worker twice. All participants were then randomly and anonymously rematched into groups of three. They were informed that they were randomly and anonymously assigned to a new group, but they were not informed of how the roles were reassigned. The Visibility condition used in round 2 was different than the one used in round 1; the Selection condition (i.e., Decision or Random) was the same in both rounds. The order of Visibility conditions used in the two rounds was randomized across sessions. Given a Selection condition, each of the six possible orders of Visibility conditions was implemented in 2 to 4 sessions.

After two rounds of the CPG, participants who played the Consumer role in either round performed the two control tasks: the dictator game and the risk preference task (see §2.2.3). Hereafter, we refer to these participants as the Consumer participants. In the dictator game, the Consumer participants were informed that they had been randomly and anonymously paired with another participant they had not been matched with before, and they were asked to make decisions as the dictator (Player 1). In each pair, the level of uncertainty in the recipient's payoff was designed to match the Visibility condition that the dictator faced when he/she was the Consumer in the CPG. For the risk preference task, the participants made individual decisions. Half of all Consumer participants played the dictator game before performing the risk preference task, whereas the other half followed the reverse order. After both the CPG and the control tasks were completed, all participants were shown the outcomes

of all tasks, including their decisions, the decisions of the participants with whom they were matched, the realized price of the product and the realized payment to the Worker in each round of the CPG, as well as their own payoffs. Finally, all participants completed the postexperiment survey described in §2.2.3. Sample instructions and the postexperiment survey are available upon request.

The experimental tasks were implemented using the z-Tree software (Fischbacher 2007). During a session, participants were not allowed to talk to each other. They made decisions and interacted with each other only through computer terminals. Before each task, participants read detailed instructions on the screen and answered a set of practice questions. They had to answer all practice questions correctly before proceeding to the next task. They were also provided with a reference sheet for the instructions.⁵ A total of 198 participants played the role of the Consumer. All of the participants in the study were students; 80.8% of them were undergraduates and the remaining 19.2% were graduate students. In addition, 61.8% of them were female, and the average age was 21.7 years old (with a standard deviation of 4.8 years). We verify that our experimental results remain unchanged after controlling for participants' demographic factors (see Appendix A.3.1). Participants earned money based on their total payoffs in all tasks. Every 16 tokens were worth 1 U.S. dollar. Participants earned an average of \$29.56, with a minimum of \$20 and a maximum of \$40. Each session lasted on average 90 minutes. The total number of Consumer participants in the Decision (Random) condition was equal to 26, 29, and 31 (38, 34, and 40) in the High, Medium, and Low Visibility conditions.

2.3 Hypotheses

Our experimental design allows us to analyze how Consumers' WTP is affected by three key factors: (i) the level of visibility about the payment to the Worker; (ii)

⁵Participants' responses in the postexperiment survey confirm that they understood the dynamics of the game. For example, they understood that choosing $WTP = 0$ would maximize the Consumer's payoff, that the Worker was more likely to receive the payment w if the WTP was higher, and that a larger e chosen by the Firm would give the Worker a larger share of the gain from selling the product.

Consumers' indirect reciprocity toward the Firm; and (iii) Consumers' prosociality. Note that if Consumers only care about their own payoffs, then they will always state $WTP = 0$ regardless of the value of e . However, an extensive literature has shown that individuals often care about others' well-being in addition to their own payoffs (see Fehr and Schmidt 2006 for a review). For example, inequality aversion may motivate individuals to reduce the payoff differences between themselves and others (Fehr and Schmidt 1999, Bolton and Ockenfels 2000). In the SR literature, a number of studies also show that consumers' WTP for a product and/or perceptions of a company often increase if the company demonstrates improved practices (e.g., Sen and Bhattacharya 2001, De Pelsmacker et al. 2005). Therefore, in our experiment, we expect Consumers' WTP to be increasing in the value of e . Building on this expected behavior, we next discuss our main hypotheses regarding the effects of visibility and indirect reciprocity on Consumers' WTP. Note that our hypothesis on visibility only examines Consumers' WTP in the Decision condition. Recall from §2.2.2 that the Random condition is designed solely to isolate the effect of indirect reciprocity on Consumers' WTP. In addition, to map our results on visibility to a supply chain context, we should measure the effect of visibility on WTP when the Firm is actively making a decision to affect the SR outcome in the upstream supply chain (i.e., the Decision condition).

We first focus on how visibility regarding the payment to the Worker impacts Consumers' WTP. Under low and medium visibility, the exact payment to the Worker if the product is sold is unknown. The current literature on how uncertainty in the outcome of one's actions affects social preferences is scarce and yields mixed results. For example, Brennan et al. (2008) study the effect of people's risk attitudes on social preferences. In their setting, participants are asked to state their WTP or willingness-to-accept (WTA)⁶ for either a sure payoff or a lottery to a recipient, where the expected value of the lottery is equal to the sure payoff. The authors observe that the WTP for the sure payoff to the recipient is similar to that for the

⁶WTA is defined as the minimum price at which a person would be willing to sell something that he/she already owns.

lottery, whereas the WTA for the sure payoff is significantly higher than that for the lottery. Krawczyk and Le Lec (2010) and Brock et al. (2013) both compare the standard dictator game with one where the payoffs to the dictator and/or the recipient are uncertain. When there is uncertainty, the dictator's decision determines the expected payoff(s). Krawczyk and Le Lec (2010) show that if uncertainty exists for both parties' payoffs and these payoffs are independent from each other, then the dictators' decisions are not significantly different from those in the standard dictator game. Conversely, Brock et al. (2013) observe that if uncertainty exists only for the recipient's payoff, then the dictators transfer significantly lower amounts to the recipients. Given these mixed results, we develop the following competing hypotheses.

Hypothesis 2.1A. *In the Decision condition, Consumers' WTP (given e) does not differ across different levels of visibility.*

Hypothesis 2.1B. *In the Decision condition, Consumers' WTP (given e) increases with a higher level of visibility.*

Our second hypothesis examines how indirect reciprocity affects Consumers' WTP. We first distinguish indirect reciprocity from direct reciprocity. Direct reciprocity occurs when an individual treats someone else in a similar manner as that person has treated him/her. Conversely, indirect reciprocity occurs when an individual treats someone else in a similar manner as that person has treated *a third party*. Direct reciprocity has been studied extensively in the literature (e.g., Charness and Rabin 2002, Cox 2004). The most common setting used to study direct reciprocity is the gift exchange game (Akerlof 1982). For example, Charness (2004) shows that the marginal increase in an employee's effort given a unit increase in his wage is larger when the wage increase is actively chosen by the firm than when it is randomly determined. This phenomenon is attributed to direct reciprocity.

Although less studied than direct reciprocity, the existence of indirect reciprocity has also been documented in the literature (e.g., Seinen and Schram 2006, Engelmann and Fischbacher 2009). For example, Stanca (2009) studies the following two-stage game. In stage 1, a sender transfers an amount x to a recipient. In stage 2, a third person (different from the sender and the recipient) observes the amount x sent by the

sender and transfers an amount y to the sender. The author finds that the amount y in stage 2 positively correlates with the amount x in stage 1. He attributes this result to indirect reciprocity. We differ from Stanca (2009) in at least two important aspects: (i) we study indirect reciprocity under different levels of outcome uncertainty; and (ii) we isolate indirect reciprocity from other social preferences by comparing WTP between the Decision and Random conditions.

In the CPG, indirect reciprocity occurs if the Consumer states a higher or lower WTP when he perceives the Firm acting more or less responsibly toward the Worker. We extend Charness's (2004) approach to measure indirect reciprocity: Given a unit increase in the Firm's decision, indirect reciprocity exists if the marginal increase in Consumers' WTP is larger in the Decision condition than in the Random condition (see Figure 2-1). To the best of our knowledge, the current literature has not studied the impact of outcome uncertainty on either direct or indirect reciprocity. Thus, we make the following null hypothesis that indirect reciprocity exists regardless of the level of visibility.

Hypothesis 2.2. *At each level of visibility, the same increase in e yields a larger increase in Consumers' WTP in the Decision condition than in the Random condition.*

Finally, our experiment also allows us to study how the effects of visibility and indirect reciprocity on WTP vary by Consumers' prosociality. Relevant to our study, de Kwaadsteniet et al. (2006) experimentally examine a common resource dilemma where a person's usage of the common resource reduces the availability of the resource to others (i.e., it leads to a negative externality). They compare participants' behavior when the size of the common resource is certain versus uncertain. They show that although all participants increase their usage of the common resource when the size of the resource becomes uncertain, such an increase is much larger for low (versus high) prosocial participants. Nevertheless, researchers have not examined the interactions among uncertainty, indirect reciprocity, and prosociality. Investigating these interaction effects along with Hypotheses 2.1 and 2.2 offers insights into whether consumers value greater supply chain visibility and when information about a company's SR efforts and outcomes resonates with its target consumers.

2.4 Experimental Results

Our data was collected by running sessions at two universities, with two rounds of the CPG being played in each session. We first verify that (i) there are no significant differences in the decisions made by participants from the two universities in any treatment (two-sided Wilcoxon rank-sum test, $p > 0.1$); and (ii) having participants play two rounds of the CPG has no significant impact on our main conclusions (see Appendix A.3.2). Thus, we present our results based on all data collected. Hereafter, we refer to the decision e as the *effort* chosen by the Firm (in the Decision condition) or by Nature (in the Random condition). We also abbreviate indirect reciprocity as reciprocity.

Table 2.1 presents the summary statistics of Consumers' WTP for all treatments. We make two initial observations. First, the average WTP is significantly higher than zero in all treatments and for all effort levels (one-sided Wilcoxon signed-rank test, $p < 0.001$). The average WTP being significantly higher than zero even when $e = 0$ may be due to Consumers' efficiency considerations; i.e., they prefer the product being sold to increase the total payoff for the group. We conducted robustness treatments and analyses to verify that efficiency considerations have a minimal effect on our main conclusions (see Appendix A.2.2). Second, we observe that the average WTP is increasing in effort level. This result is confirmed by the significantly positive correlation between the effort levels and Consumers' WTP, ranging between 0.25 and 0.35 across treatments (one-sided t test, $p < 0.001$).

Table 2.1: Mean, [Median], and (Standard Deviation) of Consumers' WTP

Effort	Decision condition			Random condition		
	Visibility condition			Visibility condition		
	High	Medium	Low	High	Medium	Low
0	17 [10] (27)	9 [0] (25)	12 [2] (24)	13 [1] (19)	13 [0] (22)	11 [0] (19)
20	27 [20] (29)	21 [15] (26)	19 [10] (26)	20 [10] (23)	23 [20] (22)	21 [15] (21)
40	34 [40] (27)	29 [32] (29)	24 [20] (29)	26 [18] (26)	32 [34] (25)	26 [20] (23)
60	41 [40] (33)	37 [25] (35)	29 [20] (33)	31 [20] (31)	42 [51] (28)	33 [30] (29)
80	44 [43] (35)	42 [35] (39)	33 [20] (36)	33 [20] (35)	45 [48] (35)	38 [35] (33)
100	52 [55] (41)	47 [45] (43)	37 [28] (39)	36 [22] (38)	45 [38] (41)	41 [32] (37)
120	55 [45] (48)	51 [50] (48)	41 [30] (43)	37 [20] (40)	48 [30] (47)	43 [32] (40)

Note. The seemingly higher WTP in the Medium Visibility, Random condition than in the High/Low Visibility, Random conditions is not statistically significant. See Appendix A.3.3 for further details.

To formally test our hypotheses, we estimate the following random-effects two-sided Tobit model:

$$\begin{aligned} \text{WTP}_{jk}^* = & \text{Intercept} + \alpha_{V_M} \cdot V_M + \alpha_{V_H} \cdot V_H + \alpha_D \cdot D + \alpha_{DV_M} \cdot D \cdot V_M + \alpha_{DV_H} \cdot D \cdot V_H + \beta \cdot e_k \\ & + \beta_{V_M} \cdot V_M \cdot e_k + \beta_{V_H} \cdot V_H \cdot e_k + \beta_D \cdot D \cdot e_k + \beta_{DV_M} \cdot D \cdot V_M \cdot e_k + \beta_{DV_H} \cdot D \cdot V_H \cdot e_k + \delta_j + \epsilon_{jk}. \end{aligned} \quad (2.1)$$

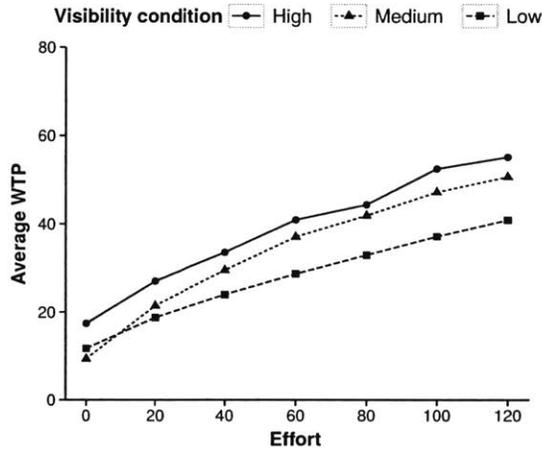
The variable WTP^* is the latent variable for Consumers' WTP in the Tobit model. The subscript j indexes a Consumer participant and k indexes an effort level. The variable e_k with $k = 1, 2, \dots, 7$ represents the seven possible effort levels in $\{0, 20, 40, 60, 80, 100, 120\}$. The dummy variables V_M and V_H indicate medium and high visibility. The dummy variable D indicates the Decision condition. The term δ_j represents the individual-specific error, and ϵ_{jk} is the independent error across WTP decisions. The baseline condition in Equation (2.1) corresponds to the Low Visibility, Random condition at $e_k = 0$.

We use a random-effects two-sided Tobit model for the following reasons. First, Consumers' WTP is bounded between 0 and 120, with 28% of observations taking the value 0. Thus, it is desirable to employ a Tobit model to account for these corner solution outcomes (Wooldridge 2002, pp. 517–542). Second, due to the strategy method, we have seven WTP observations from each Consumer. We accommodate such repeated measures with the random-effects approach; i.e., we include an individual-specific error δ_j . Our choice of regression model is well established in the literature (e.g., List 2006, Bolton et al. 2013). The regression estimates for Equation (2.1) are summarized in Appendix A.1, Table A.1. Since Consumers' WTP is not actually censored, we do not study treatment effects on the latent variable WTP^* . Instead, we focus on the marginal effects of the key independent variables in Equation (2.1). This approach is also well established (e.g., Anderson et al. 2004, Angrist and Pischke 2008).

2.4.1 Hypothesis 2.1: Does Consumers' WTP Increase with the Level of Visibility?

We first analyze the effect of visibility on Consumers' WTP in the Decision condition. Figure 2-2 shows how the observed average WTP changes with effort in each Visibility condition. We observe that the average WTP under high visibility is greater than that under medium or low visibility, and the average WTP under medium visibility is also greater than that under low visibility (except when $e_k = 0$). To formally test Hypothesis 2.1, we analyze the marginal effects of V_H , V_M , and their differences $\Delta_V \equiv V_H - V_M$ at all effort levels for the Decision condition (i.e., $D = 1$) in Equation (2.1). Positive marginal effects of V_M (V_H) indicate that the average WTP is higher under medium (high) visibility than under low visibility. Similarly, positive marginal effects of Δ_V indicate that the average WTP is higher under high visibility than under medium visibility. Table 2.2 summarizes these marginal effects.

Figure 2-2: Effect of Visibility on Consumers' WTP



We first observe that all marginal effects are positive, with the marginal effects of V_H being statistically significant at all effort levels. Thus, we support Hypothesis 2.1B (WTP increases with a higher level of visibility) when comparing high versus low visibility. Second, the marginal effects of V_M are significant only at $e_k = 120$, and the marginal effects of Δ_V are never significant. Therefore, Hypothesis 2.1B is

Table 2.2: Marginal Effects of V_H , V_M , and Their Differences, Δ_V

Effort	Visibility condition		
	V_H	V_M	Δ_V
0	8.98 (6.49)*	2.52 (5.57)	6.46 (6.75)
20	10.44 (7.12)*	3.81 (6.29)	6.63 (7.48)
40	11.93 (7.74)*	5.34 (7.01)	6.60 (8.17)
60	13.39 (8.33)*	7.05 (7.71)	6.34 (8.81)
80	14.78 (8.89)**	8.91 (8.38)	5.87 (9.40)
100	16.04 (9.42)**	10.83 (9.01)	5.21 (9.90)
120	17.13 (9.90)**	12.73 (9.57)*	4.40 (10.31)

Note. Values shown are the marginal effects evaluated at the Decision condition; i.e., at $D = 1$ in Equation (2.1). Values in parentheses are standard errors.

** : $p < 0.05$; * : $p < 0.1$; p values are derived from one-sided t tests.

supported only at $e_k = 120$ for medium versus low visibility; for $e_k < 120$, and when comparing high versus medium visibility, Hypothesis 2.1A (WTP does not differ between Visibility conditions) is instead supported. Finally, the marginal effects of V_H and V_M increase in both magnitude and statistical significance as the Firm's effort increases. Thus, greater visibility has a stronger positive effect on Consumers' WTP at higher effort levels. We summarize our first result as follows.

Result 2.1. *Consumers' WTP is significantly higher under high visibility than under low visibility, with the increase in WTP being larger at higher effort levels. The effect of visibility on WTP is positive but mostly nonsignificant when comparing medium to low or high to medium visibility.*

By Result 2.1, consumers' valuations of SR increase when the uncertainty in the resulting outcomes is reduced; i.e., when visibility improves. This suggests a potential market benefit for greater visibility into a supply chain with a disadvantaged party (e.g., poorly treated labor). Furthermore, companies that exert a high effort to improve SR could benefit the most from greater visibility because the increase in consumers' valuations is most pronounced when a company's SR effort is high. In addition, Result 2.1 suggests that consumers may be fine with a price increase due to greater supply chain visibility. This finding relates to research which has shown that visibility into costs can make increases in price seem fair to consumers (e.g., Kahneman et al. 1986, Bolton and Alba 2006).

One possible behavioral explanation for our finding is that Consumers are risk averse and dislike paying for an uncertain outcome, even if the uncertainty only impacts another participant – the Worker. To investigate this hypothesis, we classify Consumers as risk averse or non-risk averse based on their decisions in the risk preference task (see §2.2.3). We analyze whether the effect of visibility on WTP differs between risk-averse and non-risk-averse Consumers. We observe that only non-risk-averse Consumers state significantly lower WTP under low visibility than under high visibility. Conversely, the WTP decisions of risk-averse Consumers are not statistically different across Visibility conditions. Hence, we find no evidence that the lower WTP observed under low visibility is due to Consumers’ risk aversion. Instead, we postulate that it is a result of Consumers’ self-serving biases.

A self-serving bias is an individual’s tendency to “place greater weight on information that is consistent with [his/her] preferences” (Babcock et al. 1996, p. 3). Under low visibility, a Consumer prone to a self-serving bias may use the presence of uncertainty to justify selecting a lower WTP and hence, keeping a higher payoff for him/herself.⁷ Building on established methods to study self-serving biases (e.g., Babcock and Loewenstein 1997, Pahlia et al. 2013), in the postexperiment survey we ask Consumers to state on a 5-point Likert scale how much they agree or disagree with each of the following three statements: “Not knowing what the exact payment to Player C [the Worker] was, I mainly focused on Player C’s *minimum possible (maximum possible) (average)* payment when making my decision.” The italicized phrase is the only difference across the three statements. A higher score stated for a statement means that the participant agrees with that statement more strongly. Based on their responses, we categorize our Consumer participants as follows. If a participant gives the “minimum possible” statement the highest score, then the participant is classified as exhibiting a self-serving bias. Otherwise, the participant is classified as not exhibiting this bias. Under this categorization, 16 out of 31 (52%) Consumer participants

⁷Similarly, Haisley and Weber (2010) document the presence of a self-serving bias in a dictator game. They find that when the distribution of the recipient’s payoff is unknown (versus when it is known), the dictator overestimates the recipient’s expected payoff and contributes less. By overestimating, the dictator justifies keeping a larger amount for him/herself.

in the Low Visibility, Decision condition exhibit a self-serving bias.⁸

We then analyze if the effect of visibility on WTP depends on whether Consumers exhibit a self-serving bias (see Appendix A.4). We observe that Consumers who exhibit this bias state a significantly lower WTP under low visibility than under high visibility. In sharp contrast, the WTP decisions of Consumers who do not exhibit a self-serving bias are not statistically different between high and low visibility. Therefore, our analysis suggests that the presence of a self-serving bias is a key contributor to the observed effect of visibility on WTP.

Recall from §2.2.1 that to focus on SR in the upstream supply chain, we design the three players' initial endowments in the CPG to capture the presence of a disadvantaged party. Nevertheless, one may wonder how the effect of visibility on Consumers' WTP might be influenced by the Workers' wealth status. In Appendix A.2.3, we discuss additional treatments and analyses of the Decision condition with high and low visibility, where the Worker's initial endowment is increased to 50 or 80 tokens. Results in these additional treatments again show that greater visibility has a positive effect on WTP for Consumers who exhibit a self-serving bias. In particular, these Consumers' WTP is significantly lower under low visibility than under high visibility. Conversely, we do not observe such a difference in the WTP decisions of Consumers who do not exhibit a self-serving bias. These results further highlight the role of self-serving biases in affecting consumers' valuations of SR under different levels of visibility.

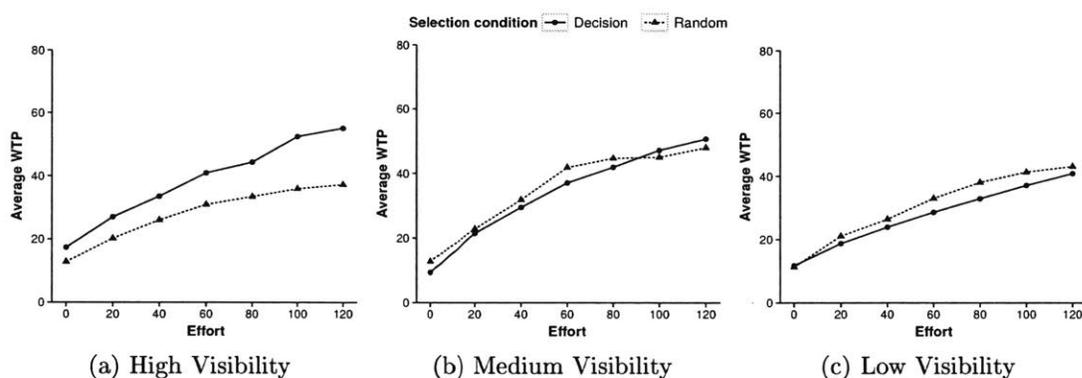
2.4.2 Hypothesis 2.2: Does Reciprocity Exist at Each Level of Visibility?

Next, we address whether Consumers exhibit reciprocity in their WTP at each level of visibility. If reciprocity exists, then we also examine whether it has a positive or negative impact on Consumers' WTP; i.e., whether it generates a higher or lower

⁸Consumers under medium visibility were not included in this analysis as the survey data was corrupted during one of the sessions. Consumers under high visibility were not categorized because they did not evaluate these statements.

WTP in the Decision condition than in the Random condition. Figure 2-3 shows the average WTP observed under the Decision and Random conditions for all three levels of visibility. Based on Equation (2.1), to test Hypothesis 2.2 we examine the marginal effects of $D \cdot e_k$ at each effort level and each Visibility condition. Table 2.3 summarizes these marginal effects. A positive value indicates that the slope of WTP as a function of effort is larger in the Decision condition than in the Random condition and hence, reciprocity exists.

Figure 2-3: Effect of Reciprocity on Consumers' WTP



We first examine the High Visibility condition. In Figure 2-3a, we observe that the average WTP increases in the effort level at a faster rate in the Decision condition than in the Random condition; i.e., the solid line has a steeper slope than the dotted line. This observation is confirmed by the significantly positive marginal effects under high visibility shown in Table 2.3. Thus, reciprocity exists and Hypothesis 2.2 is supported when visibility is high. In addition, the observed WTP is significantly higher in the Decision condition than in the Random condition for $e \geq 100$ (one-sided Wilcoxon rank-sum test, $p < 0.1$). Hence, the presence of reciprocity generates a positive effect on Consumers' WTP at high effort levels when visibility is high.

Next, consider the Medium and Low Visibility conditions. Our results show that the effect of reciprocity diminishes with lower levels of visibility. First, note for the Medium Visibility condition in Table 2.3 that although the marginal effects of $D \cdot e_k$ are positive at all effort levels, they are significant only when $e_k \geq 100$. Also, the presence of reciprocity at $e_k \geq 100$ does not result in a significantly higher WTP in

Table 2.3: Marginal Effects of $D \cdot e_k$

Effort	Visibility condition		
	High	Medium	Low
0	0.09 (0.04)***	0.02 (0.04)	-0.02 (0.03)
20	0.10 (0.04)***	0.03 (0.04)	-0.01 (0.04)
40	0.11 (0.04)***	0.04 (0.05)	-0.00 (0.04)
60	0.12 (0.04)***	0.05 (0.05)	0.00 (0.04)
80	0.13 (0.04)***	0.06 (0.05)	0.01 (0.04)
100	0.13 (0.04)***	0.07 (0.05)*	0.02 (0.04)
120	0.13 (0.04)***	0.07 (0.04)*	0.02 (0.04)

Note. Values in parentheses are standard errors.
 ***: $p < 0.01$; *: $p < 0.1$; p values are derived from one-sided t tests.

the Decision condition than in the Random condition (one-sided Wilcoxon rank-sum test, $p > 0.1$; see also Figure 2-3b). Second, observe for the Low Visibility condition in Table 2.3 and Figure 2-3c that the slopes of WTP as a function of effort are very similar between the Decision and Random conditions. Thus, we do not observe reciprocity under low visibility.

To summarize, we find partial support for Hypothesis 2.2 and obtain the following result.

Result 2.2. *There exists an interaction between visibility and indirect reciprocity.*

- (i) *Under high visibility, indirect reciprocity exists at all effort levels and has a positive effect on Consumers' WTP at high effort levels.*
- (ii) *Under medium visibility, indirect reciprocity exists only at high effort levels and does not have an effect on Consumers' WTP.*
- (iii) *Under low visibility, we do not find evidence of indirect reciprocity.*

Result 2.2 suggests that consumers value a company's SR efforts (i.e., indirect reciprocity exists) but only if the outcomes of these efforts are highly visible. As a result, companies that can provide a high level of visibility into their supply chains could benefit from emphasizing their active engagement (e.g., time and capital investment) when communicating their SR practices to consumers.

Our findings add to the limited literature on how uncertainty may impact reciprocity. Prior studies in this area mainly focus on situations where a person's ini-

tial endowment is not perfectly observable by his/her partner, and show that self-interested actions are more likely to occur (e.g., Güth et al. 1996, Kanagaretnam et al. 2010). Recently, Rubin and Sheremeta (2016) use a gift exchange game with potential reward/punishment to examine whether the presence of a random shock to the agent’s effort can affect the principal’s wage offer and how much the principal rewards/punishes the agent. They find that when a random shock exists and is perfectly observable by the principal, the principal offers a lower wage and rewards the agent’s effort less, compared to when there is no random shock. However, none of these decisions change significantly when the principal cannot observe the value of the shock (i.e., when there is uncertainty in the agent’s effort). Furthermore, in their design, the principal always perfectly observe the outcome of the agent’s effort. Our design and results differ from this work as we study outcome uncertainty and show that its presence *diminishes* indirect reciprocity. These varying results suggest that our understanding of how uncertainty impacts direct or indirect reciprocity is far from conclusive. Thus, further research with different decision contexts would be valuable.

2.4.3 The Impact of Consumers’ Prosociality

Next, we investigate how Consumers’ prosociality (elicited via the dictator game control task) impacts their WTP decisions. We first note that the distributions of the dictators’ decisions are very similar across all six treatments (Kruskal-Wallis test, $p = 0.35$). Therefore, we use the median dictator decision among all Consumers (which is equal to 20) to classify two types of Consumers. Those Consumers who transferred 0 or 20 tokens are classified as *low prosocial* and account for 64% of the entire sample (127 participants). Those Consumers who transferred 40 or more tokens are classified as *high prosocial* and account for the remaining 36% of the sample (71 participants). To examine our hypotheses for high and low prosocial Consumers, we first reestimate Equation (2.1) with the data from these two groups separately. We then compute for each prosocial type the same marginal effects as before. The regression results are summarized in Appendix A.1, Table A.1. To

confirm the robustness of our results, we repeat our analysis with two alternative classifications of the Consumers' prosocial types and obtain similar results as discussed below. See Appendix A.2.4 for more details.

We first confirm that the effect of visibility on Consumers' WTP does not vary significantly between prosocial types (Appendix A.1, Table A.2). Both types of Consumers state significantly higher WTP under high visibility than under low visibility, whereas their WTP does not differ significantly between high and medium visibility (Result 2.1). When comparing medium to low visibility, high prosocial Consumers state significantly higher WTP under medium visibility at all positive effort levels, whereas low prosocial Consumers do so when $e_k \geq 100$. Compared to the pooled results, the increases in Consumers' WTP from low to medium visibility are more pronounced when examined by prosocial type.

The key difference in the observed behavior between high and low prosocial Consumers lies in the role reciprocity plays in affecting their WTP decisions (Appendix A.1, Table A.3). Figures 2-4 and 2-5 illustrate the high and low prosocial Consumers' average WTP in each treatment. We first observe in Figure 2-4 that the slope of the solid line is never significantly steeper than that of the dotted line (see also Table A.3, columns 2-4). Therefore, *we do not observe reciprocity among high prosocial Consumers in any of the Visibility conditions*. Hence, Hypothesis 2.2 is not supported for high prosocial Consumers. That is, these Consumers are not significantly influenced by knowing whether the outcome is due to the Firm's active effort or determined by Nature. Instead, their WTP decisions are mainly driven by the expected payment to the Worker (i.e., the value of e) and the level of visibility.

In sharp contrast, *low prosocial Consumers' WTP is significantly affected by reciprocity*. First, Figure 2-5a is very similar to Figure 2-3a. The solid line being steeper than the dotted line shows that reciprocity exists for low prosocial Consumers under high visibility (see also Table A.3, column 5). In addition, similar to Figure 2-3a, reciprocity leads to a significantly higher WTP at $e = 120$ in the Decision condition than in the Random condition ($p < 0.1$).⁹ Second, we continue to observe the presence

⁹We compare low prosocial Consumers' WTP decisions between the Decision and Random con-

Figure 2-4: Effect of Reciprocity on High Prosocial Consumers' WTP

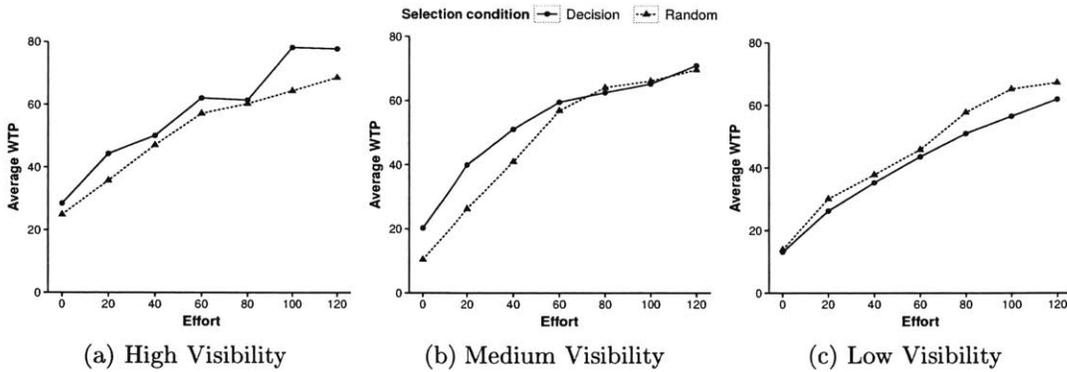
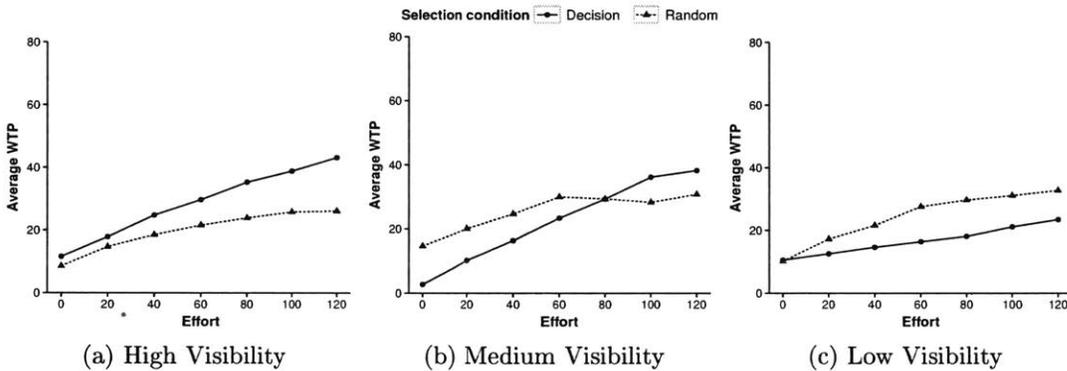


Figure 2-5: Effect of Reciprocity on Low Prosocial Consumers' WTP



of reciprocity for low prosocial Consumers under medium visibility (see Figure 2-5b and Table A.3, column 6). However, unlike the case of high visibility, reciprocity now has a *negative* effect on WTP that results in a lower WTP at low effort levels in the Decision condition than in the Random condition ($p < 0.1$ for $e \in \{20, 40\}$). Third, we find no evidence of reciprocity under low visibility for low prosocial Consumers, as the slope of the solid line is not significantly steeper than that of the dotted line (Figure 2-5c and Table A.3, column 7). Hence, we partially support Hypothesis 2.2 among low prosocial Consumers. We summarize our findings regarding Consumers' prosociality.

Result 2.3. *Both high and low prosocial Consumers' WTP is significantly higher*

conditions for a given effort in two aspects: (i) the proportion of zero values (one-sided χ^2 test) and (ii) the distribution of non-zero values (one-sided Wilcoxon rank-sum test). See Appendix A.3 for further details.

under high visibility than under low visibility. With respect to indirect reciprocity:

- (i) For high prosocial Consumers, we do not observe indirect reciprocity at any level of visibility.
- (ii) For low prosocial Consumers: (a) There is significant indirect reciprocity under both high and medium visibility. Moreover, indirect reciprocity has a positive (negative) effect on WTP under high (medium) visibility at high (low) effort levels. (b) We do not find evidence of indirect reciprocity under low visibility.

Result 2.3 suggests that whether consumers reward a company's effort to improve SR practices in its supply chain highly depends on the interaction between visibility and consumer heterogeneity. Consumers who naturally care about others' well-being (i.e., high prosocial consumers) are less interested in knowing the amount of effort that a company exerts and more interested in observing precise information about the outcomes of such effort. Conversely, consumers who are driven by self-interests (i.e., low prosocial consumers) are influenced by information about a company's SR effort, but this information may positively or negatively affect their valuations depending on the level of visibility.

We also observe that low prosocial Consumers demonstrate a unique behavior under low visibility. Their WTP is significantly *lower* in the Decision condition than in the Random condition at all positive effort levels except for $e = 120$ ($p < 0.05$ for $e \in \{20, 40, 60, 80\}$, $p < 0.1$ for $e = 100$; see Figure 2-5c). This result is driven by a significantly larger fraction of low prosocial Consumers stating $WTP = 0$ for all effort levels in the Decision versus the Random condition under low visibility (53% versus 25%; one-sided χ^2 test, $p < 0.1$). We postulate that low prosocial Consumers' behavior under low visibility is affected by the different levels of responsibility they perceive for themselves in the Decision condition versus the Random condition. In the Random condition, the Consumer is the *only* player who makes an active decision to influence the final payoffs of all three players. Conversely, in the Decision condition, both the Firm's *and* the Consumer's decisions determine the final payoffs. Thus, the Consumer is likely to feel less responsible for the Worker's well-being in the

Decision condition. With a decreased sense of responsibility for the final outcomes, the Consumer's low prosociality motivates him to state a lower WTP in the Decision condition than in the Random condition when visibility is low. The phenomenon of responsibility alleviation has been documented in other games examining social preferences (e.g., Charness 2000, Dana et al. 2007).

We indeed find evidence in the postexperiment survey that our participants perceived different levels of responsibility between the Decision and Random conditions. For example, a participant who played the Consumer role in the Random condition under low visibility stated that "being Player B [the Consumer] made me feel pressured as the decision maker...that affects other people." In contrast, another participant who played the Consumer role in the Decision condition under low visibility stated that "It is Player A [the Firm] which decides whether to be fair or not with Player C [the Worker]." Our observations are in line with Gneezy et al.'s (2010) results from a field experiment. The authors study consumers' purchase decisions for souvenir photos in a setting where half of the sales revenue goes to a charity. They compare the case when consumers name their own prices for the photos (i.e., a "pay-what-you-want" strategy) with the case when listed prices are used. The authors find that consumers make larger donations to the charity under pay-what-you-want pricing. They attribute this outcome to consumers' perceptions of "shared social responsibility" under this strategy.

Finally, we note an interesting connection between potential responsibility alleviation and self-serving biases. In particular, the result that low prosocial Consumers state lower WTP in the Decision (versus Random) condition under low visibility occurs only for those who exhibit a self-serving bias. Moreover, a significantly larger fraction of these Consumers state $WTP = 0$ for all effort levels in the Decision (versus Random) condition (80% versus 33%; one-sided χ^2 test, $p < 0.05$). We conjecture that the following behavior is occurring. First, low prosocial Consumers who exhibit a self-serving bias use the presence of uncertainty to justify not contributing or contributing very little to increase the Workers' payoffs. Second, when they become the sole decision maker in the Random condition, the increased sense of responsibility

counteracts this justification and motivates Consumers to state a higher WTP. These exploratory observations suggest that low prosocial individuals' behavior is more malleable and easily influenced by environmental factors (e.g., uncertainty and whether they are the only decision makers who can affect others' payoffs). This malleability has been shown in studies of social dilemmas where one's actions have negative externalities on others' welfare (e.g., Roch and Samuelson 1997, Au and Kwong 2004). Further examining how low prosocial individuals' behavior is influenced by environmental factors in an SR context could be a valuable future research direction.

2.5 The Product Choice Study

We next examine how well our result regarding the effect of visibility on Consumers' WTP in the CPG generalizes to broader SR contexts and multiple product categories. In our postexperiment survey, we include a product choice study in which participants are shown multiple pairs of products with different attributes and are asked to state which product in each pair they prefer to buy. We test three product categories – coffee, t-shirts, and laptop computers – and five SR topics – treatment of the company's (suppliers') employees, community development in regions where the company (suppliers) operates, and charitable donations. Each Consumer evaluates 12 pairs of products divided into two 6-pair blocks, with each block consisting of a distinct product category and SR topic. We manipulate two attributes within a block: the product prices and the message attached to each product about the company's SR practices for one of the topics. Within a pair, one product is labeled with a vague message for a specific SR topic, while the other product is labeled with a precise message for the same topic. A precise (vague) message contains information showing a company's high (low) visibility into a given SR practice in its supply chain. The product with precise information is priced equal to or greater than the product with vague information. We examine six different price premiums for precise information within a block (0% to 20%, 18.75%, and 12% for coffee, t-shirts, and computers). The product choice study is based on a simplified conjoint design commonly used in the

marketing literature (e.g., Arora and Henderson 2007, Rooderkerk et al. 2011). Like most conjoint studies, our product choice study is not incentivized; i.e., participants do not actually purchase the products they prefer.

Our findings in the product choice study confirm and complement our experimental result that Consumers value greater visibility (i.e., support Hypothesis 2.1B). We examine the effect of visibility on Consumers' preferences by analyzing their choices between vague and precise information for each combination of product and SR topic (Appendix A.3, Figure A-2). First, we observe that for all three product categories and all five SR topics, a vast majority of Consumers prefer products with more precise information when they are not offered at a premium. Second, we analyze the average premium that Consumers are willing to pay for products with more precise information. We calculate this premium as follows. Let p_v denote the price of a product with vague information (which is fixed for a given product category). If a Consumer prefers the product with precise information over one with vague information at all prices up to and including \bar{p} , then the premium that the Consumer is willing to pay is $(\bar{p} - p_v)/p_v \times 100\%$. We find that our Consumer participants are on average willing to pay a premium ranging from 1.6% to 11% when more precise information is provided (one-sided Wilcoxon signed-rank test against the null hypothesis of zero premium, $p < 0.05$).

To compare our results in the product choice study with those in the CPG, we perform the following additional analysis. We define two groups of Consumers based on the premium they are willing to pay for products with precise information. The "premium" group contains Consumers who are willing to pay a *positive* premium for both products and SR topics they evaluated. The "no-premium" group contains Consumers who are not willing to pay a premium for at least one product and SR topic. We then analyze whether the effect of visibility on WTP in the CPG is stronger for the premium group than for the no-premium group. We find that Consumers in the premium group state a significantly higher WTP under high visibility than under medium or low visibility, for almost all effort levels. For Consumers in the no-premium group, the marginal effects of the visibility variables in the regression are positive but

mostly not significant. Thus, Consumers' stated preferences in the product choice study align with the observed WTP differences across Visibility conditions in the CPG.

2.6 Conclusions and Managerial Insights

As the role of social responsibility in business continues to evolve, consumer demands are forcing companies to adapt their operations to meet the needs of a changing marketplace. Companies must address not only how to establish socially responsible practices throughout their supply chains but also how to make their supply chains more transparent to themselves and the public. In this paper, we employ an incentivized controlled laboratory experiment with human subjects to investigate how supply chain visibility, indirect reciprocity, and individual prosociality jointly affect consumers' valuations of a company's SR practices. In what follows, we conclude the paper by discussing two specific insights on (i) whether a company can benefit from greater visibility, and (ii) what SR information resonates with a company's target consumers.

Insight 2.1. *Consumers value greater visibility regarding a company's SR practices in its upstream supply chain where there exist disadvantaged parties (e.g., poorly treated labor). This is especially true if visibility is currently low and if consumers exhibit a self-serving bias.*

Recent events such as the Foxconn suicides in Shenzhen in 2010 and the Rana Plaza collapse in Bangladesh in 2013, have highlighted the need for companies to improve visibility into their supply chains. However, gaining such visibility can be costly and time-consuming (Doorey 2011). Our results demonstrate a potential market benefit for greater visibility into a supply chain that operates in developing countries where visibility is lacking and the labor is often subject to poor economic conditions. This result applies to a broad range of industries and contexts, including the apparel and consumer electronics industries where the use of inexpensive, overseas labor is a common practice.

We find that visibility is particularly valuable when a company's consumers tend to assume poor outcomes under low visibility; i.e., exhibit a self-serving bias. In this case, greater visibility can help to minimize these consumers' tendency to use information uncertainty as an excuse not to pay for SR. Although we are unaware of any existing market surveys that directly examine whether consumers exhibit a self-serving bias, this bias has been shown to be prevalent in many practical situations (see Gino et al. 2016 for a recent review). Therefore, it is crucial for companies to recognize this bias and understand when its presence could affect consumers' decision making.

Insight 2.2. *Information regarding SR efforts can induce higher valuations from consumers when the outcomes of the efforts are highly visible. This is particularly true for low prosocial consumers.*

Our results highlight when information regarding a company's SR *efforts* resonates with its target consumers. Specifically, we observe that if consumers naturally care about others' well-being (i.e., if they are high prosocial consumers), then they are less interested in learning about the amount of effort that a company exerts and more interested in observing precise information regarding the outcomes of such effort. This is because high prosocial consumers do not exhibit indirect reciprocity but do value greater visibility. If instead, consumers are more driven by self-interests (i.e., if they are low prosocial consumers), then information about the company's efforts could impact these consumers' valuations depending on the level of visibility. Under high visibility, low prosocial consumers demonstrate strong indirect reciprocity, which increases their valuations of a company's SR efforts. However, under lower levels of visibility, the presence of indirect reciprocity may decrease low prosocial consumers' valuations at low effort levels. In addition, these consumers may justify a low willingness-to-pay by shifting responsibility for the workers' well-being onto the company.

In practice, companies can develop an understanding of the heterogeneity of their consumers (e.g., their different prosociality) through psychographic segmentation, a method commonly used in marketing to segment consumers based on psychological

dimensions “including activities, interests, opinions, needs, values, attitudes, and personality traits” (Wells 1975, p. 197). For example, leading market research firms conduct consumer surveys to segment consumers based on their attitudes toward ethics and the environment to help companies develop better strategies for targeting specific segments (e.g., Cone Communications/Echo 2013, Mintel Group 2015). In Appendix A.2.4, we discuss how we utilize our attitudinal survey and follow a similar approach to segment consumers. We find that classifying participants’ prosocial types based on their responses in the attitudinal survey yields similar results as classifying them based on their decisions in the dictator game. This result further strengthens the relevance of our experimental findings on the behaviors of consumers with different prosociality.

To conclude, our research yields valuable insights into the impact of supply chain visibility on consumers’ valuations of a company’s SR practices, as well as the behavioral motives behind consumers’ valuations. Our results also underscore the importance of accounting for consumers’ heterogeneity when studying their valuations of SR. We hope that these insights will motivate others to further study behavioral factors that can drive consumers’ and companies’ decisions in various SR contexts. Related to our work, studying the impact of both visibility and disclosure on consumers’ valuations would be a valuable future research direction.

Chapter 3

Improving Supplier Social Responsibility under Incomplete Visibility

3.1 Introduction

Establishing socially responsible practices in a supply chain is a difficult task. This is especially true when a company lacks visibility into its suppliers' practices. In this paper, we examine how a manufacturer (she) can improve a supplier's (he) social responsibility (SR) practices when these practices cannot be perfectly observed by the manufacturer. Specifically, we focus on the manufacturer's investment to improve the supplier's SR capabilities; for example, a manufacturer developing training programs at a supplier to educate employees on human rights topics, or a manufacturer investing to help improve the safety of a supplier's facility.¹ To capture how consumer demand can influence the manufacturer's and the supplier's decisions, we incorporate the potential for SR information to be disclosed to the consumers either directly by the manufacturer or by an external third party (e.g., an NGO). Our results address (i)

¹We follow the European Commission's definition of social responsibility as "[companies integrat- ing] social and environmental concerns in their business operations and in their interaction with their stakeholders on a voluntary basis," (Dahlsrud 2008). We specifically focus on social concerns in social responsibility.

how the manufacturer should invest to improve the supplier's SR capabilities and what SR information to disclose given the current level of visibility the manufacturer has into her supply chain, and (ii) when the potential disclosure of SR information to the consumers (either voluntarily by the manufacturer or involuntarily by a third party) can positively or negatively impact the supplier's SR performance.

A recent study by The Sustainability Consortium found that 81% of the 1,700 companies surveyed lacked full visibility into the SR practices of their suppliers (The Sustainability Consortium 2016). At the same time, improving visibility into a supply chain can be a very costly and time-consuming task for a company (Doorey 2011). As a result, complete visibility into a supply chain is rarely achieved. Instead, companies often make decisions on how to improve upstream suppliers' SR practices and what information to disclose to consumers under incomplete visibility. We focus on a manufacturer's decisions for a given level of visibility rather than examining how the manufacturer can improve her visibility (e.g., through more frequent audits). We make this choice in part because increasing supply chain visibility is a much longer-term endeavor than a manufacturer's investment in a supplier's SR capabilities and her disclosure of SR information (the two manufacturer decisions we study). For example, retailer Marks & Spencer discusses in its 2015 Plan A report that it strives to improve the information available about its supply chain by 2020 (Marks & Spencer 2015b).

To motivate a supplier to improve his SR practices, a manufacturer can either offer the supplier incentives (e.g., preferred supplier status or investment in the supplier's capabilities) or threaten the supplier with penalties (e.g., contract termination). We focus on a manufacturer's investment in a supplier's SR capabilities because it has been shown to be highly effective in improving supplier performance (Porteous et al. 2015); however, it remains an underutilized and understudied method (Gillai et al. 2013). Our choice is motivated by recent examples of companies investing in their suppliers' SR capabilities. Starbucks has invested more than \$100 million in Coffee and Farmer Equity Practices, farmer loans, and other programs aimed at improving farmers' livelihoods along its supply chain (Starbucks 2017). In working with its suppliers,

Hewlett-Packard “invests in programs that empower and protect workers” by training them on topics like anti-discrimination and labor rights (HP 2016). Similarly, Nike organizes training sessions for their contract manufacturers to address topics that impact the ethical treatment of workers such as human resource management and health and safety practices (Porteous and Rammohan 2013).

In this paper, we study a supply chain with one supplier and one manufacturer. The manufacturer sells a product in a market where at least some consumers care whether the product is made in a socially responsible manner. To capture incomplete visibility, we model the supplier’s initial level of SR as his *private information*. The manufacturer has a prior belief about the supplier’s current level of SR and observes a signal about this level. How likely the signal captures the supplier’s true initial level of SR depends on the manufacturer’s visibility into her supply chain. In practice, an audit report is a good example of a signal since it only provides a snapshot of a supplier’s current practices, and therefore, is subject to noise (EY and UN Global Compact 2016).

In our model, the supplier is the only party that can directly improve the SR practices of the supply chain; the manufacturer can only help to reduce the cost of SR by investing in the supplier’s capabilities. The manufacturer is motivated to invest in the supplier in part because there is a possibility that the supplier’s level of SR will be disclosed to the consumers by an external third party. Similar to Chen and Lee (2016), we treat the level of SR as a “soft” quality attribute, which can only be observed and verified at the supplier’s site (e.g., working conditions). This feature differentiates our work from most of the quality management literature (e.g., Zhu et al. 2007, Babich and Tang 2012) because the final level of SR cannot be inferred via product inspection by either the manufacturer or the consumers.

We analyze and compare two settings: (i) The manufacturer does not disclose SR information to the consumers (*No Manufacturer Disclosure*), and (ii) the manufacturer voluntarily discloses SR information to the consumers based on the information available to her (*Manufacturer Disclosure*). In the latter case, the manufacturer can increase her demand by disclosing, but she may also incur a penalty if she overstates

the supplier's SR performance. We compare the two settings in order to study when the manufacturer prefers to voluntarily disclose SR information. However, there are regulations in practice that require manufacturers to disclose SR information; e.g., the California Transparency in Supply Chains Act and the EU Sustainability Reporting Directive. Our results in the Manufacturer Disclosure setting can inform manufacturers' decisions under such regulations. Within the operations management literature there is a growing interest in the topic of supply chain transparency, both with respect to companies' visibility into their supply chains and the disclosure of SR information to consumers (e.g., Kalkanici and Plambeck 2015, Chen and Lee 2016, Kraft et al. 2017). We contribute to this stream of work by examining how a manufacturer's incomplete visibility into a supplier's SR practices impacts her investment to improve the supplier's capabilities and her decision to disclose SR information to the consumers.

We derive a number of key results from our model. First, the manufacturer should invest a high (low) amount of resources to improve the supplier's capabilities if the information she observes suggests the supplier's current SR practices are poor (good). With greater visibility into her supply chain, the manufacturer can better tailor her level of investment as she can better identify whether the supplier needs significant support. Second, if the manufacturer plans to disclose SR information to the consumers, then she should always be more aggressive with her investment in the supplier's capabilities (as compared to if she plans not to disclose). This more aggressive strategy ensures a better SR performance from the supplier when the manufacturer discloses. Third, when the manufacturer voluntarily discloses SR information to the consumers, she is likely to overstate (understate) the supplier's SR performance if the observed information suggests very poor (very good) current practices. However, with greater visibility, the manufacturer becomes more "truthful" in her disclosure; i.e., she is less likely to either overstate or understate the supplier's SR performance given her available information. Finally, when choosing between disclosing and not disclosing SR information to the consumers, the manufacturer is most likely to prefer not disclosing when the supplier appears to have only average

SR practices. Furthermore, an increase in visibility or the probability of third-party disclosure may cause the manufacturer to prefer not disclosing information to the consumers, thus hurting the resulting SR performance of the supplier.

Literature Review: Our work is closely related to two streams of research: socially responsible supply chains and supplier development. Within the socially responsible supply chain literature, analytical studies that examine suppliers' SR practices typically address one of two topics: (i) identifying potential SR risks in a supply chain through tools such as audits and inspections (e.g., Huang et al. 2016, Plambeck and Taylor 2016, Wang et al. 2016); or (ii) motivating better SR performance of upstream suppliers. Recent works that address topic (i) have begun to examine the effect of a firm's voluntary disclosure decision on supply chain SR performance (e.g., Chen et al. 2015, Kalkanci and Plambeck 2015). The papers that examine topic (ii) span a number of different areas including sourcing strategies (e.g., Agrawal and Lee 2016, Guo et al. 2016), supply chain design (e.g., Letizia and Hendrikse 2016, Orsdemir et al. 2016, Zhang et al. 2017), and supplier development (e.g., Mendoza and Clemen 2013).

The works most closely related to ours consider both topics (i) and (ii). Cho et al. (2016) examine a company's choice of inspection policy and wholesale price to combat a supplier's use of child labor. They study two separate scenarios: one in which the company's inspection policy is only known to the company, and the other in which the policy is also known to the supplier and third parties. Chen and Lee (2016) analyze how a company can design incentive schemes in sourcing contracts (e.g., contingency payments) and invest in screening mechanisms (e.g., supplier certifications and process audits) to prevent unethical actions by a supplier. Lewis et al. (2016) investigate a mechanism design problem where a company can invest to develop a supplier's capabilities to achieve sustainable quality. In their setting, both the company's demand and the supplier's production cost are private information.

We contribute to this growing literature in two aspects. First, we model and examine supply chain transparency in a more holistic manner by capturing *both* a manufacturer's visibility into a supplier's practices *and* the manufacturer's disclosure

of information to consumers. In particular, visibility refers to the extent to which the manufacturer can observe the SR practices of the supplier. To better represent practice, we model a continuous level of visibility as opposed to a binary state with either full visibility or no visibility at all. In addition, we capture incomplete visibility by modeling a game with asymmetric information between the supplier and the manufacturer, where the supplier's private information about his current SR practices follows a continuous distribution (rather than a two-type distribution). To the best of our knowledge, Lewis et al. (2016) is the only paper in the SR literature with a similar information asymmetry setup as ours; however, they do not examine a company's potential voluntary disclosure of information to consumers. Our model thus allows us to determine how supply chain visibility and the manufacturer's endogenous disclosure decision jointly impact the SR performance of the supplier. Second, our paper addresses topic (ii) by analyzing an understudied approach – a manufacturer's investment in a supplier's SR capabilities. This analysis is timely, as companies are increasingly realizing the importance of going beyond monitoring to actually develop better capabilities at their suppliers when addressing SR challenges (Locke et al. 2007, EY and UN Global Compact 2016). Our results offer valuable guidance on how companies should better leverage this underutilized approach to improve SR, particularly under the constraint of incomplete visibility.

Regarding supplier development, a number of papers analyze how a buyer can use a contracting approach (e.g., offer a price premium) to improve a supplier's quality or process (e.g., Corbett and DeCroix 2001, Zhu et al. 2007, Li and Debo 2009, Kim and Netessine 2013). We do not consider supply chain contracts because (i) a supplier's SR performance is difficult to measure and verify, and thus, often noncontractible (Norman and MacDonald 2004); and (ii) in practice, the use of price premiums is rare and often ineffective for improving a supplier's SR performance (Porteous et al. 2015). A group of works outside the contracting literature have applied analytical models to examine buyers' supply development decisions (e.g., Kim 2000, Babich 2010, Liu et al. 2010, Talluri et al. 2010, Wang et al. 2010, Karaer et al. 2017). Of these papers, the studies that are most closely related to ours examine a company's

investment in supplier reliability, where the investment addresses the uncertainty in the outcome (e.g., production yield) of a known supplier type (Babich 2010, Liu et al. 2010, Wang et al. 2010). We instead consider a setting with information asymmetry, where the manufacturer invests in the supplier’s SR capabilities when the supplier’s type is his private information.

3.2 Model Setup

In this section, we review our model formulation and assumptions. We first discuss the setting where the manufacturer does not disclose SR information. We then build on this setting to introduce the case where the manufacturer discloses.

3.2.1 No Manufacturer Disclosure

We consider a supply chain with one manufacturer (she) and one supplier (he). The manufacturer sells a product in a market where at least some consumers care whether the product is made in a socially responsible manner. The supplier’s initial level of SR, s_0 , is his *private* information. We also refer to s_0 as the supplier’s type. The manufacturer does not know the supplier’s type but has a prior belief that s_0 is distributed on $[m, M]$ with cumulative distribution function (CDF) $\Phi(\cdot)$ and probability density function (PDF) $\phi(\cdot)$. We assume $\phi(x) > 0$ for all $x \in [m, M]$. Before any decision is made, the manufacturer observes a signal, \tilde{s} , that contains some information about s_0 . The accuracy of the signal (i.e., how likely the signal is equal to s_0) depends on the level of supply chain visibility, v , that the manufacturer has. Specifically, with probability v , $\tilde{s} = s_0$; otherwise, with probability $1 - v$, \tilde{s} is equal to a random value drawn from the manufacturer’s prior belief. Therefore, with higher supply chain visibility (i.e., a higher value of v), the manufacturer is more certain that the observed signal corresponds to the true initial SR level of the supplier.

In our setting, the supplier is the only party that can directly impact the final level of SR. Specifically, he can choose to increase from s_0 to any $s \geq s_0$. We assume that the supplier will not reduce his level of SR to $s < s_0$; i.e., he does not engage in

practices that deteriorate his current SR practices. The manufacturer cannot directly impact SR, but she can indirectly influence the supplier's choice of s by investing $\beta \in [0, 1]$ to improve the supplier's capabilities and reduce his cost of SR. Note that investing in a supplier's capabilities differs from direct cost sharing because the former does not require the manufacturer to know s_0 with certainty (and hence the supplier's cost of improving SR). As a result, such an investment can be made with incomplete visibility, whereas cost sharing would be difficult if not impossible. We model visibility v as exogenously fixed and independent of the manufacturer's investment β . This modeling choice is reasonable because improving visibility (e.g., through audits) is typically a lengthy endeavor for a firm (Doorey 2011). We are interested in studying the manufacturer's investment (e.g., training suppliers on human resource practices) that would likely not drastically impact v in the short term.

The market consists of two types of consumers: an $\alpha \in (0, 1)$ fraction of the consumers are socially conscious (SC); the remaining $(1 - \alpha)$ fraction are socially neutral (SN). Demand from both types of consumers depends on the retail price r , which is fixed and exogenous. Demand from the SC consumers also depends on the supplier's final level of SR if this information is disclosed. Specifically, the final level of SR may be observed by a third party (e.g., an NGO) with a given, exogenous probability q . If the supplier is found to have a final level of SR below the consumers' minimum acceptable SR standard, then the third party discloses this information and the demand from SC consumers decreases.

The sequence of events is as follows: (i) After observing the signal \tilde{s} , the manufacturer chooses her investment β to improve the supplier's SR capabilities. (ii) The supplier selects $s \geq s_0$ to be the final level of SR. (iii) With probability q , the third party observes s . If s is below the SC consumers' minimum acceptable SR standard, then the third party discloses s to the consumers. (iv) Finally, demand is realized, and the supplier and the manufacturer earn their profits. Note that the manufacturer is the Stackelberg leader in our setup. Thus, by backward induction, she can fully anticipate the supplier's best response and take this into account when choosing β in step (i). Next, we discuss each step in more detail.

(i) *Manufacturer's investment in the supplier's SR capabilities:* The manufacturer makes her investment decision to maximize her expected profit, Π_M . Given signal \tilde{s} , the manufacturer's expected profit can be written as

$$\mathbb{E}_{s_0}[\Pi_M(\beta, s^*(s_0, \beta)) | \tilde{s}] = \left(D_{SN} + \mathbb{E}_{s_0}[D_{SC}(s^*(s_0, \beta)) | \tilde{s}] \right) (r - w) - \delta(\beta), \quad (3.1)$$

where D_{SN} is the demand from SN consumers and D_{SC} is the expected demand from SC consumers. The notation $s^*(s_0, \beta)$ represents the supplier's optimal decision on the final level of SR given his initial level s_0 and the manufacturer's investment β . Demand D_{SN} is independent from $s^*(s_0, \beta)$, whereas demand D_{SC} depends on $s^*(s_0, \beta)$. Since s_0 is the supplier's private information, when calculating her expected profit, the manufacturer has to take the expectation over all possible values of s_0 based on her posterior belief given \tilde{s} . The term $(r - w)$ represents the manufacturer's profit margin, where w is the wholesale price paid by the manufacturer to the supplier. We model the wholesale price w as exogenous and independent of the supplier's SR level. This approach is reasonable because in a survey of 334 companies, Porteous et al. (2015) find that price premiums are (i) not commonly used to improve suppliers' SR practices and (ii) less effective than investing in suppliers' SR capabilities. Taking w as an exogenous parameter is also common in the socially responsible operations literature (see, e.g., Guo et al. 2016, Plambeck and Taylor 2016). The function $\delta(\beta)$ captures the manufacturer's cost to invest in the supplier's capabilities. We assume $\delta(\beta)$ to be strictly increasing, strictly convex, and twice-continuously differentiable in β with $\delta'(0) = 0$.

(ii) *Supplier's SR decision:* The supplier selects $s \geq s_0$ to maximize his expected profit

$$\Pi_S(\beta, s) = (D_{SN} + D_{SC}(s)) (w - c) - (1 - \beta) (\rho(s) - \rho(s_0)). \quad (3.2)$$

The function $\rho(s)$ captures the supplier's fixed cost for his final level of SR. We assume $\rho(s)$ to be strictly increasing, convex, and twice-continuously differentiable in s . We model the supplier's cost of improving SR as $\rho(s) - \rho(s_0)$ to capture that the same increase in SR (i.e., from s_0 to s) is more costly at a higher value of s_0 . In this regard,

we also assume that the supplier's current cost of SR at s_0 is a sunk cost.

The parameter c in Equation (3.2) represents the per-unit cost of production. This cost is exogenous, known to all parties, and independent of s and s_0 . We make this modeling choice because our motivating examples involve non-production related SR investments at a supplier, such as human resource training and facility safety improvements. In addition, Rangan et al. (2015) find that 32% (35%) of the 142 managers they surveyed reported decreased (increased) production costs due to SR investments. Given this variation and our motivating examples, we model the supplier's unit cost as being independent of the manufacturer's SR investment. Finally, we capture fixed rather than variable (i.e., per unit sold) costs of SR. The literature on supplier development has examined companies working to decrease suppliers' unit costs (e.g., Kim 2000, Kim and Netessine 2013) or fixed investment costs (e.g., Babich 2010, Mendoza and Clemen 2013). Our focus on reducing the supplier's fixed SR costs is in line with the common challenge that SR development typically requires high fixed costs (but low variable costs; Borzaga and Becchetti 2010).

(iii) Third-party disclosure: After the manufacturer and the supplier make their decisions, the third party observes s with probability q . To delineate between what consumers feel are good and bad practices, we consider a minimum SR standard, $\hat{s} \in (m, M)$, such that the disclosure of $s = \hat{s}$ would have no impact on the demand from SC consumers. We focus on the third party only disclosing bad practices with respect to SC consumers' minimum standard. That is, the third party discloses s to the consumers only if $s < \hat{s}$. If instead $s \geq \hat{s}$, then no information is disclosed.

(iv) Demand and profits are realized: Following a common approach in the literature (e.g., Moorthy 1988, Bagnoli and Watts 2003), we model a continuum of consumers characterized by their private valuations of the product, θ , with θ being uniformly distributed on $[0, 1]$. Without loss of generality we normalize the total market size to 1. A consumer's type (SC or SN) is independent of his/her private valuation, and each consumer buys at most one unit of the product. We consider a given retail price $r < 1$ which applies to both types of consumers if they purchase the product.

SC consumers incur an additional utility equal to $\gamma(\min\{s, \hat{s}\})$ if they purchase the product and the third party discloses s (which occurs only if $s < \hat{s}$). The function $\gamma(\cdot)$ is assumed to be strictly increasing, strictly concave, and twice-continuously differentiable. As previously stated, disclosing the minimum SR standard \hat{s} has no impact on the demand from SC consumers; i.e., $\gamma(\hat{s}) = 0$. Since the third party only discloses $s < \hat{s}$, we have $\gamma(\min\{s, \hat{s}\}) < 0$ for any third-party disclosure.

Given this setup, a consumer with private valuation θ will buy the product if and only if: (i) $\theta - r \geq 0$, if the consumer is SN, or if the consumer is SC and no SR information is disclosed; or (ii) $\theta - r + \gamma(\min\{s, \hat{s}\}) \geq 0$, if the consumer is SC and the third party discloses s . If no SR information is disclosed, then SC consumers behave like SN consumers and do not take SR into account when making their purchase decisions. Recall that an α fraction of the consumers are SC. Thus, the expected demand from SN and SC consumers can be written as

$$\begin{aligned}
 D_{SN} &= (1 - \alpha) \int_0^1 \mathbb{1}_{\theta \geq r} d\theta = (1 - \alpha)(1 - r), \\
 D_{SC}(s) &= \alpha(1 - q) \int_0^1 \mathbb{1}_{\theta \geq r} d\theta + \alpha q \int_0^1 \mathbb{1}_{\theta \geq r - \gamma(\min\{s, \hat{s}\})} d\theta \\
 &= \alpha(1 - q)(1 - r) + \alpha q \max \left\{ 0, \min \left\{ 1, 1 - r + \gamma(\min\{s, \hat{s}\}) \right\} \right\}, \quad (3.3)
 \end{aligned}$$

where $\mathbb{1}$ denotes the indicator function. We make the following assumption regarding demand.

- Assumption 1.** (i) *A consumer with the highest valuation $\theta = 1$ will always buy the product for any SR level above the minimum level of m ; i.e., $\gamma(m) > r - 1$.*
- (ii) *A consumer with the lowest valuation $\theta = 0$ will never purchase the product, regardless of the SR level; i.e., for any s , $\gamma(s) < r$.*

Assumption 1 ensures that the manufacturer captures at least some but not all of the market. Under this assumption, the term $\max \left\{ 0, \min \left\{ 1, 1 - r + \gamma(\min\{s, \hat{s}\}) \right\} \right\}$ in Equation (3.3) simplifies to $1 - r + \gamma(\min\{s, \hat{s}\})$. The resulting total expected de-

mand in the market (including both SN and SC consumers) is therefore equal to $1 - r + \alpha q \gamma (\min\{s, \hat{s}\})$.

3.2.2 Manufacturer Disclosure

The key difference between the Manufacturer Disclosure setting and the No Manufacturer Disclosure setting is the following: After step (ii) (the supplier's SR decision) and before step (iii) (third-party disclosure), the manufacturer chooses a (final) level of SR, s_D , to disclose to the consumers. The value s_D disclosed by the manufacturer does not necessarily match the supplier's final level of SR; instead, s_D represents the level of SR that the manufacturer prefers to communicate to the consumers. The manufacturer discloses s_D before the third party may observe s . If the third party observes s , then it discloses s to the consumers only if $s < s_D$. In this case, the manufacturer has overstated the supplier's SR level. As a result, the manufacturer incurs a penalty proportional to $(s_D - s)$ and the demand from SC consumers depends on s instead of s_D . Conversely, if the third party observes $s \geq s_D$ (i.e., the manufacturer has understated the supplier's SR level) or if the third party does not observe s , then no information is disclosed by the third party, and the demand from SC consumers depends on s_D . Disclosing SR information to the consumers while still lacking full visibility into their supply chains is not an uncommon practice for companies. For example, in 2014 Unilever published a progress report about its Sustainable Palm Oil Policy and disclosed that 58% of the palm oil in their supply chain was traceable to known mills. Despite lacking visibility into the remaining 42%, they still shared information about the progress made in the year and some of the social responsibility initiatives occurring throughout their palm oil supply chain (Unilever 2014).

In this setting, the manufacturer's expected profit for a given signal \tilde{s} can be written as

$$\begin{aligned} \mathbb{E}_{s_0} [\Pi_M(\beta, s_D, s^*(s_0, \beta)) | \tilde{s}] = & \left(D_{SN} + \mathbb{E}_{s_0} [D_{SC}(s_D, s^*(s_0, \beta)) | \tilde{s}] \right) (r - w) \\ & - \delta(\beta) - pq \mathbb{E}_{s_0} [\max\{s_D - s^*(s_0, \beta), 0\} | \tilde{s}]. \end{aligned} \quad (3.4)$$

The impact of disclosure on the expected value of D_{SC} depends on how $\min\{s^*, s_D\}$ compares to \hat{s} if the third party discloses, and how s_D compares to \hat{s} if the third party does not disclose. In addition to the potential impact on demand, the manufacturer may incur a penalty if she overstates the supplier's SR level in her disclosure. The last term in Equation (3.4) corresponds to the expected penalty incurred by the manufacturer if $s^* < s_D$. This penalty captures the loss of goodwill the manufacturer suffers (e.g., brand damage) due to the third party disclosing that she has overstated the SR level of her supplier (see, e.g., Chen and Lee 2016, Cho et al. 2016, Plambeck and Taylor 2016 for similar goodwill costs). Therefore, when selecting the optimal value of s_D to disclose, the manufacturer must balance the tradeoff between possibly increasing demand and the risk of incurring a penalty from overstating.

The supplier's expected profit remains the same as in the No Manufacturer Disclosure setting except that the expected value of D_{SC} now depends on both s and s_D . We do not include a penalty for the supplier because (i) we are not investigating an audit setting where the manufacturer imposes a penalty on the supplier for the discovery of poor performance (as in Plambeck and Taylor 2016); and (ii) levying penalties on suppliers in developing countries can be very difficult (as discussed in Chen and Lee 2016). Note however that like the manufacturer, the supplier does potentially suffer a loss of demand if the third party discloses $s < \hat{s}$.

With respect to demand, an SC consumer will buy the product if and only if $\theta - r + \gamma(y) \geq 0$, where $y = s_D$ when the third party does not observe or disclose s , and $y = \min\{s, s_D\}$ when the third party observes and discloses s (which occurs only if $s < s_D$). Given Assumption 1 stated earlier, the expected demand from SC consumers can be written as

$$\begin{aligned} D_{SC}(s_D, s) &= \alpha(1-q) \int_0^1 \mathbb{1}_{\theta \geq r - \gamma(s_D)} d\theta + \alpha q \int_0^1 \mathbb{1}_{\theta \geq r - \gamma(\min\{s, s_D\})} d\theta \\ &= \alpha(1-q) (1 - r + \gamma(s_D)) + \alpha q (1 - r + \gamma(\min\{s, s_D\})). \end{aligned}$$

We next analyze the manufacturer's and the supplier's decisions in the No Manufacturer Disclosure (§3.3) and Manufacturer Disclosure (§3.4) settings. In addition,

we study how these decisions depend on the level of visibility and the probability of third-party scrutiny. In §3.5 we compare the two settings and identify when the manufacturer prefers to disclose SR information versus not disclose, and how the manufacturer’s disclosure choice impacts the supplier’s level of SR. Table 3.1 summarizes our notation.

Table 3.1: Notation

<i>Decision Variables</i>	
β	Manufacturer’s investment to improve the supplier’s SR capabilities; $\beta \in [0, 1]$
s	Supplier’s final level of SR; $s \geq s_0$
s_D	Supplier’s final level of SR <i>disclosed</i> by the manufacturer (Manufacturer Disclosure setting only)
<i>SR-related Variables</i>	
s_0	Supplier’s initial level of SR; $s_0 \in [m, M]$ and is the supplier’s private information
\tilde{s}	Signal the manufacturer observes regarding s_0
\hat{s}	Minimum SR standard of SC consumers; $\hat{s} \in (m, M)$
q	Probability that the third party observes s ; $q \in (0, 1)$
v	Level of supply chain visibility the manufacturer has; $v \in (0, 1)$
<i>Cost and Demand Parameters</i>	
r	Unit retail price
w	Unit wholesale price paid by the manufacturer to the supplier
c	Supplier’s per-unit cost of production
p	Penalty factor the manufacturer incurs if the third party discloses $s < s_D$ (Manufacturer Disclosure setting only); $p > 0$
θ	Consumers’ private valuation for the product; $\theta \sim U[0, 1]$
α	Fraction of consumers who are socially conscious; $\alpha \in (0, 1)$
$\delta(\beta)$	Manufacturer’s cost to invest β in the supplier’s SR capabilities; $\delta'(\beta) > 0$ for $\beta > 0$, $\delta'(0) = 0$, and $\delta''(\beta) > 0$
$\rho(s)$	Supplier’s fixed cost of SR at level s ; $\rho'(s) > 0$ and $\rho''(s) > 0$
$\gamma(\cdot)$	SC consumer’s additional utility from the disclosure of SR information; $\gamma'(\cdot) > 0$ and $\gamma''(\cdot) < 0$

3.3 Results: No Manufacturer Disclosure

We first analyze the setting where the manufacturer does not voluntarily disclose any SR information to the consumers. We address the following questions: (i) What is the manufacturer’s optimal investment to improve the supplier’s SR capabilities? (ii) Given this investment, what is the supplier’s optimal choice of SR? (iii) How are these decisions affected by the signal about the supplier’s type, the manufacturer’s visibility, and the probability of third-party scrutiny? Throughout our analysis we define the

ratio $R_f(x)$ as $\frac{f''(x)}{f'(x)}$ for any function $f(x)$ and make the following assumptions.

Assumption 2. For any positive investment $\beta > 0$, the lowest supplier type (i.e., a supplier with initial level of SR $s_0 = m$) always increases his SR level; i.e., the optimal final level of SR $s^*(m, \beta) > m$ for any $\beta > 0$.

Assumption 3. It is never optimal for the manufacturer to choose a value of β so high that the highest supplier type (i.e., a supplier with initial level of SR $s_0 = M$) would increase his SR level.

Assumption 4. $R_\rho(s)$ is non-decreasing in s , and $R_\gamma(s)$ is non-increasing in s .

Assumption 5. For any $\beta \in (0, 1]$, $R_\delta(\beta)(1 - \beta) \geq 1 + \frac{\phi(s_m^*)}{\Phi(s_m^*)} \frac{1}{R_\rho(s_m^*) - R_\gamma(s_m^*)}$, where $s_m^* \equiv s^*(m, \beta)$.

Assumptions 2 and 3 ensure that in the optimal solution the lowest supplier type always improves SR, while the highest supplier type never does. Assumption 4 means that the supplier's cost of SR, $\rho(s)$, increases sharply as s increases when s is already high, and SC consumers' additional utility from SR, $\gamma(s)$, decreases sharply as s decreases when s is already low. That is, the cost of SR increases rapidly if the supplier attempts to achieve very good SR practices. Conversely, SC consumers' demand for the product drops significantly when the supplier's SR practices approach a very poor level. Note that Assumption 4 is satisfied, for example, if $\rho(\cdot)$ and $\gamma(\cdot)$ are exponential functions. Finally, Assumption 5 can be interpreted as follows. First, $R_\delta(\beta)(1 - \beta) \geq 1$ implies that, for any investment β , the rate of increase in the marginal cost $\delta'(\beta)$ (captured by $R_\delta(\beta)$) is faster than the rate of reduction in the supplier's SR cost (captured by $1 - \beta$). For example, as β approaches 1, Assumption 5 guarantees that R_δ goes to $+\infty$. In other words, this assumption captures the existence of inefficiencies in the manufacturer's investment in the supplier's capabilities. The second term in the right hand side captures the additional inefficiency that results from the manufacturer not observing the supplier's type with certainty.

We first analyze the supplier's optimal choice of SR for a given value of β . Fol-

lowing §3.2.1, the supplier's problem can be specified as

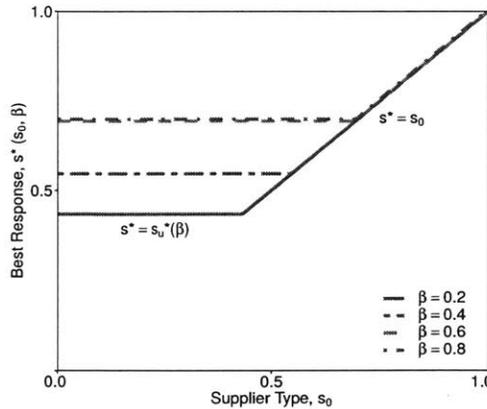
$$\max_{s \geq s_0} \left(1 - r + \alpha q \gamma (\min\{s, \hat{s}\}) \right) (w - c) - (1 - \beta) (\rho(s) - \rho(s_0)). \quad (3.5)$$

We define the *unconstrained supplier's problem* as the maximization problem in Equation (3.5) without the constraint $s \geq s_0$. Let $s_u^*(\beta)$ denote the unique optimal solution to this unconstrained problem. The following result characterizes the supplier's optimal decision.

Theorem 1. *For a given β , the supplier's optimal SR decision is given by $s^*(s_0, \beta) = \max\{s_0, s_u^*(\beta)\}$. Furthermore, $\frac{\partial s^*(s_0, \beta)}{\partial \beta} \geq 0$, $\frac{\partial s^*(s_0, \beta)}{\partial q} \geq 0$, and $s_u^*(\beta) \leq \hat{s}$ for all $\beta \in [0, 1]$.*

Figure 3-1 illustrates Theorem 1. For any s_0 , the supplier's best response to the manufacturer's investment is as follows: if s_0 is strictly lower than $s_u^*(\beta)$, then the supplier increases his SR level to $s_u^*(\beta)$; otherwise, the supplier does not improve SR; i.e., $s^*(s_0, \beta) = s_0$. Furthermore, since the third party never communicates values of s greater than \hat{s} to consumers, the supplier has no incentive to improve SR beyond \hat{s} . Thus, $s_u^*(\beta) \leq \hat{s}$ for all β . This explains why the supplier's best response is identical for $\beta = 0.6$ and $\beta = 0.8$ in Figure 3-1.

Figure 3-1: Supplier's Optimal SR Decision for Different Values of β



Based on the supplier's best response, we next characterize the manufacturer's optimal investment to improve the supplier's SR capabilities. The manufacturer's

optimal investment is determined by the solution to the following problem:

$$\max_{\beta \in [0,1]} \left(1 - r + \alpha q \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta), \hat{s}\}) \mid \tilde{s} \right] \right) (r - w) - \delta(\beta), \quad (3.6)$$

where $\mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta), \hat{s}\}) \mid \tilde{s} \right]$ captures the expected change in demand from SC consumers if the third party observes and discloses the supplier's final level of SR. We define $\hat{\beta}$ such that $s_u^*(\hat{\beta}) = \hat{s}$; i.e., $\hat{\beta}$ is the minimum investment needed from the manufacturer to ensure that the supplier's final SR level is at least \hat{s} . Our next result summarizes the manufacturer's optimal investment decision.

Theorem 2. *For a given level of visibility v , there exist $\beta_L(v)$ and $\beta_H(v)$ such that $0 \leq \beta_L(v) \leq \beta_H(v) < 1$. The manufacturer's optimal investment, $\beta^*(v, \tilde{s})$, is defined as follows:*

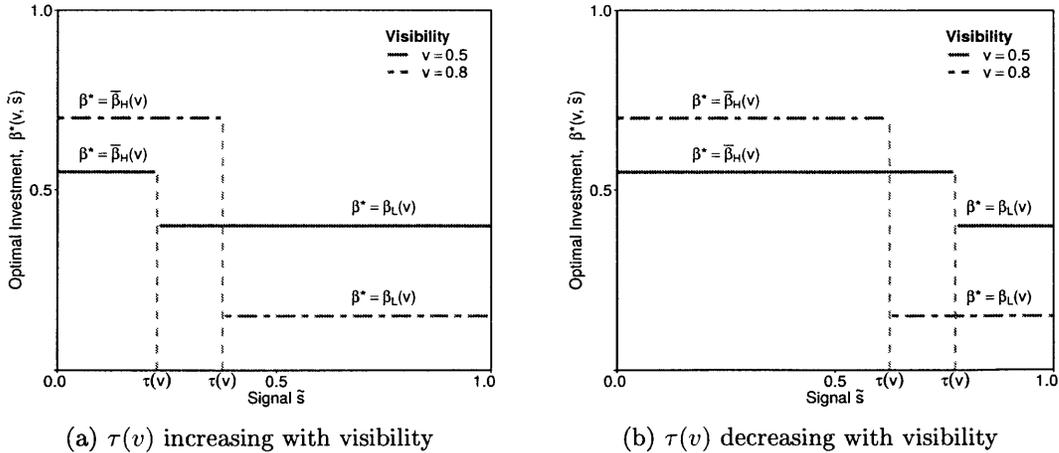
(a) *If $\beta_L(v) \geq \hat{\beta}$, then $\beta^*(v, \tilde{s}) = \hat{\beta}$ for all \tilde{s} .*

(b) *If $\beta_L(v) < \hat{\beta}$, then there exists a threshold $\tau(v) \in (m, \hat{s})$ such that:*

(i) *If $\tilde{s} \leq \tau(v)$, then $\beta^*(v, \tilde{s}) = \bar{\beta}_H(v) \equiv \min\{\beta_H(v), \hat{\beta}\}$;*

(ii) *If $\tilde{s} > \tau(v)$, then $\beta^*(v, \tilde{s}) = \beta_L(v)$.*

Figure 3-2: Manufacturer's Optimal Investment in the Supplier's SR Capabilities



Theorem 2 characterizes the structure of the manufacturer's optimal investment to improve the supplier's SR capabilities for a given level of visibility, v . In case (a),

motivating the supplier to increase his SR level to at least \hat{s} requires a low investment from the manufacturer. Hence, the manufacturer invests $\hat{\beta}$ for any signal. Case (b) is illustrated in Figure 3-2 for two levels of visibility. In this case, the manufacturer only needs to consider two possible values of β : a low investment $\beta_L(v)$ or a high investment $\bar{\beta}_H(v)$. If the signal \tilde{s} is below the threshold $\tau(v)$, then it is optimal for the manufacturer to make a high investment, $\bar{\beta}_H(v)$. This is because the signal indicates that the supplier may initially have poor SR practices. Conversely, if \tilde{s} is above the threshold $\tau(v)$, then it is optimal for the manufacturer to make a low investment, $\beta_L(v)$, because the signal indicates that the supplier may already have good SR practices. The manufacturer makes this low investment only to help improve the supplier's SR level in case the signal is incorrect (i.e., in case s_0 is in fact low). If the signal is correct and $s_0 = \tilde{s} > \tau(v)$, then the supplier's best response to $\beta_L(v)$ is simply to stay at s_0 . Two extreme cases for the value of v help to better understand the structure of the manufacturer's choice of β . First, if the manufacturer had full visibility (i.e., $v = 1$ and hence $\tilde{s} = s_0$ with certainty), then the manufacturer would not invest at all for high supplier types (i.e., $\beta_L(1) = 0$). Conversely, if $v = 0$, then the signal would be completely uninformative and the manufacturer would select the same value of β regardless of \tilde{s} (i.e., $\beta_L(0) = \bar{\beta}_H(0)$). In our subsequent analysis, we will focus on case (b) in Theorem 2, when the manufacturer has incomplete (but some) visibility and it is too costly for her to motivate all supplier types to achieve a final SR level that is at least \hat{s} .

3.3.1 The Effect of Visibility on the Manufacturer's Optimal Investment Strategy and the Supplier's Expected Final SR Level

Our next proposition demonstrates how improved supply chain visibility helps the manufacturer to better tailor her investment in the supplier's SR capabilities. To simplify notation, we will drop the argument (v) and write β_H , $\bar{\beta}_H$, β_L , and τ .

Proposition 1. *(i) β_L is strictly decreasing in v and β_H is strictly increasing in v .*

(ii) The threshold τ is increasing in v if and only if

$$\delta(\bar{\beta}_H) - \delta(\beta_L) > \alpha q (r - w) \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \bar{\beta}_H), \hat{s}\}) - \gamma(\min\{s^*(s_0, \beta_L), \hat{s}\}) \mid \tilde{s} \neq s_0 \right]. \quad (3.7)$$

Proposition 1(i) demonstrates how the two potential investment values in the manufacturer's optimal strategy change with visibility. As visibility increases, the manufacturer becomes more confident that the signal she observes captures the supplier's true initial SR level. As a result, the manufacturer offers potentially more help to the supplier (i.e., β_H increases) if the supplier likely has poor SR practices as suggested by the signal (i.e., when $\tilde{s} \leq \tau$). Conversely, the manufacturer decreases her investment (i.e., β_L decreases) if the signal suggests that the supplier already has good SR practices (i.e., when $\tilde{s} > \tau$). This differentiation exists as long as the signal contains some information about the supplier's type (i.e., $v \neq 0$). Therefore, better visibility into the supplier's practices allows the manufacturer to utilize her resources more efficiently and to ensure that she helps those suppliers who need it most.

Proposition 1(ii) specifies the condition under which the threshold τ , and hence the range of supplier types for which the manufacturer should make a high investment, increases or decreases with v . The left hand side of Equation (3.7) captures the manufacturer's cost difference between a high and a low investment. The right hand side of Equation (3.7) captures the manufacturer's expected revenue gain from making a high versus a low investment *if* the manufacturer observes an incorrect signal about the supplier's initial SR level (i.e., $\tilde{s} \neq s_0$). The intuition behind Proposition 1(ii) is as follows. If the cost difference between a high and a low investment is large relative to the expected revenue gain, then the manufacturer is reluctant to choose a high investment. In this case, greater visibility – i.e., increased confidence that the signal captures the true supplier type – can convince the manufacturer that investing $\bar{\beta}_H$ is worthwhile for a wider range of \tilde{s} values. Thus, the threshold τ increases as v increases (Figure 3-2a). Conversely, if the cost difference is small relative to the expected revenue gain, then the manufacturer is willing to make a high investment in general. With worse visibility, the manufacturer is less confident that the signal is

correct. Hence, she opts for investing $\bar{\beta}_H$ for a wider range of \tilde{s} values to ensure that the supplier's final level of SR is acceptable, just in case the signal is not correct. As a result, the threshold τ increases as v decreases (Figure 3-2b).

We next discuss how the manufacturer's level of visibility impacts the supplier's final level of SR. Note that visibility impacts SR in two ways: (i) it affects the accuracy of the signal \tilde{s} , and (ii) it affects the manufacturer's investment given a signal. Therefore, to understand the combined effect, we analyze how visibility impacts the *expected* final level of SR for a given supplier type s_0 , with the expectation taken over all possible signals that the manufacturer may observe. Specifically, the expected final level of SR given s_0 is defined as

$$ES(v, s_0) \equiv vs^*(s_0, \beta^*(v, s_0)) + (1 - v) \int_{\tilde{s} \in [m, M]} s^*(s_0, \beta^*(v, \tilde{s})) \phi(\tilde{s}) d\tilde{s}.$$

We numerically observe that while greater visibility generally increases the expected final SR level, it can lead to a lower expected final SR level if the supplier's type takes an intermediate value. In this case, the manufacturer would often make a high investment to ensure good practices when visibility is not high (i.e., when the signal is often incorrect). Since the supplier's type takes an intermediate value, he would benefit from such a high investment by improving his SR practices (i.e., by selecting $s > s_0$). Greater visibility, on the other hand, increases the chance that the manufacturer observes s_0 and realizes that the supplier's current practices are acceptable. As a result, the manufacturer reduces her investment when observing $\tilde{s} = s_0$, leading to a lower expected final SR level. Proposition 11 shows that this case always exists for some intermediate s_0 values. See Appendix B.2 for further details.

3.3.2 The Effect of Third-Party Scrutiny on the Manufacturer's Optimal Investment Strategy and the Supplier's Final SR Level

Regarding the effect of third-party scrutiny on the manufacturer's investment strategy, we find that β_L and β_H are not monotone in q (Appendix B.3). Nevertheless, the following result always holds.

Proposition 2. *The threshold τ is strictly increasing in q .*

As the probability of third-party scrutiny increases, the manufacturer prefers the high investment $\bar{\beta}_H$ over the low investment β_L for a wider range of signals. To understand the reason behind this result, consider the signal $\tilde{s} = \tau$ for which the manufacturer is indifferent between making a high or a low investment. As q increases, the final SR level $s^*(\tau, \beta)$ is more likely to be disclosed to the consumers. Hence, it is in the manufacturer's best interest to make a high investment to ensure acceptable practices at the supplier. Therefore, she prefers $\bar{\beta}_H$ over β_L for a wider range of signals when q increases.

Although an increase in q can lead to a decrease in the manufacturer's investment, we show that it never decreases the supplier's final level of SR.

Proposition 3. $\frac{ds^*(s_0, \beta^*(\tilde{s}, q), q)}{dq} \geq 0$ for all $s_0, \tilde{s} \in [m, M]$.

In our model, third-party scrutiny impacts the supplier's SR decision in two ways. First, it directly impacts the supplier's decision by changing the demand from SC consumers due to possible disclosure. As disclosure becomes more likely (with a higher q), the supplier is more likely to improve SR. Second, third-party scrutiny indirectly impacts the supplier's SR level through its (nonmonotone) effect on the manufacturer's investment. Proposition 3 demonstrates that the direct impact of q on demand dominates the indirect impact of q on the manufacturer's investment. Thus, the supplier is always motivated to improve his SR practices when there is a higher chance that these practices will be observed by the third party and disclosed to the consumers.

3.4 Results: Manufacturer Disclosure

We next examine the setting where the manufacturer voluntarily discloses SR information to the consumers. Similar to the No Manufacturer Disclosure setting, third-party scrutiny can still affect the demand from SC consumers. In addition, the manufacturer may suffer a penalty if she overstates the supplier's final SR level and the third party later discovers and discloses this information. We address the following questions in this setting: (i) What is the manufacturer's optimal investment to improve the supplier's SR capabilities? (ii) What is the supplier's optimal choice of SR? (iii) What level of SR should the manufacturer disclose? (iv) How are these decisions affected by the signal about the supplier's type, the manufacturer's visibility, and the probability of third-party scrutiny? We will use subscript D on the relevant variables to indicate the Manufacturer Disclosure setting.

First, we analyze the supplier's optimal choice of SR given the manufacturer's investment β_D and the final SR level disclosed, s_D .² The supplier's problem can be written as

$$\max_{s \geq s_0} \left(1 - r + \alpha(1 - q)\gamma(s_D) + \alpha q\gamma(\min\{s, s_D\}) \right) (w - c) - (1 - \beta_D)(\rho(s) - \rho(s_0)). \quad (3.8)$$

Comparing to Equation (3.5), the only difference in the Manufacturer Disclosure setting is that the expected demand from SC consumers is now affected by the manufacturer's disclosure, s_D . Specifically, if the third party does not observe the supplier's final level of SR (which happens with probability $1 - q$), then the demand from SC consumers (α fraction of the market) increases by $\gamma(s_D)$. If instead the third party observes the final level of SR (which happens with probability q), then it discloses this information only if the manufacturer has overstated the supplier's level of SR. Hence, the impact of disclosure on the demand from SC consumers becomes $\gamma(\min\{s, s_D\})$. Nevertheless, the structure of the supplier's optimal SR decision, $s^*(s_0, \beta_D)$, remains

²In our model, the manufacturer does not observe the true supplier type s_0 due to incomplete visibility. As a result, the manufacturer cannot condition her disclosure decision on s_0 or the corresponding optimal SR decision by the supplier. Therefore, from a mathematical standpoint, the supplier's SR decision and the manufacturer's disclosure decision can be considered as being made simultaneously.

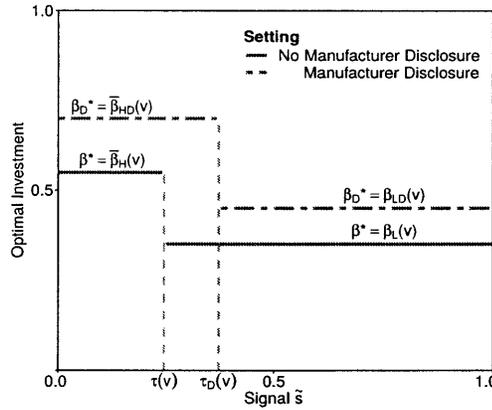
the same as in Theorem 1, except that now the unconstrained optimal solution satisfies $s_{uD}^*(\beta_D) \leq s_D$ as opposed to $s_{uD}^*(\beta_D) \leq \hat{s}$ (Theorem 5).³

We next analyze the manufacturer's investment in the supplier's SR capabilities and her choice of what final SR level to disclose. The manufacturer solves the following problem:

$$\max_{\beta_D \in [0,1]} \left\{ \max_{s_D \in [m,M]} \left\{ \left(1 - r + \alpha(1-q)\gamma(s_D) + \alpha q \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta_D), s_D\}) \mid \tilde{s} \right] \right) (r - w) - \delta(\beta_D) - pq \mathbb{E}_{s_0} \left[\max\{s_D - s^*(s_0, \beta_D), 0\} \mid \tilde{s} \right] \right\} \right\}. \quad (3.9)$$

The inner maximization solves for the manufacturer's optimal disclosure decision given the investment she has put into the supplier. The outer maximization solves for the manufacturer's optimal investment, taking into account its effect on the supplier's SR decision and the manufacturer's own disclosure decision. The manufacturer anticipates the supplier's optimal SR decision $s^*(s_0, \beta_D)$ for each possible s_0 when making her decisions. The last term in Equation (3.9) captures the expected penalty incurred by the manufacturer if she overstates the supplier's SR performance and the third party later discloses the supplier's actual final SR level.

Figure 3-3: Comparing the Manufacturer's Optimal Investment in the Supplier's SR Capabilities



³We can simplify the notation by removing the dependence of s^* on s_D , because the latter only constitutes an upper bound on the supplier's optimal level of SR (similar to \hat{s} in the No Manufacturer Disclosure setting). However, in equilibrium this upper bound is already captured by the maximum value of β_D , $\hat{\beta}_D$, that the manufacturer may invest.

The structure of the manufacturer's optimal investment in this setting remains the same as in the No Manufacturer Disclosure setting (Theorem 6). As illustrated in Figure 3-3, the manufacturer makes a high investment $\bar{\beta}_{HD}(v) \equiv \min\{\beta_{HD}(v), \hat{\beta}_D\}$ if the observed signal \tilde{s} is less than or equal to a threshold $\tau_D(v)$. Otherwise, if $\tilde{s} > \tau_D(v)$, then the manufacturer makes a low investment $\beta_{LD}(v)$.⁴ The effects of visibility and third-party scrutiny on the manufacturer's optimal investment decision are also qualitatively the same as in the No Manufacturer Disclosure setting (Propositions 8–9). To simplify notation, hereafter we will drop the argument (v) and write $\bar{\beta}_{HD}$, β_{HD} , β_{LD} , and τ_D .

Our next result shows that the manufacturer's optimal investment in the Manufacturer Disclosure setting is strictly higher than that in the No Manufacturer Disclosure setting.

Proposition 4. *For any level of visibility v and signal \tilde{s} , $\beta_{LD} > \beta_L$, $\bar{\beta}_{HD} > \bar{\beta}_H$, $\tau_D > \tau$, and therefore, $\beta_D^*(v, \tilde{s}) > \beta^*(v, \tilde{s})$.*

Proposition 4 follows from the fact that in the Manufacturer Disclosure setting, there is always a positive chance for the manufacturer to incur a penalty for overstating the supplier's final level of SR. This result implies that when the manufacturer discloses SR information to the consumers, she invests a *strictly larger* amount of resources to improve the supplier's SR capabilities as compared to when she does not disclose SR information (Figure 3-3).

3.4.1 What Level of SR Should the Manufacturer Disclose?

We next discuss the additional decision that the manufacturer needs to make in the Manufacturer Disclosure setting: her choice of final SR level to disclose to the consumers, s_D . When determining s_D , the manufacturer faces a tradeoff between capturing additional demand from disclosing a high s_D versus risking a penalty if s_D

⁴As in Theorem 2(a), if it is inexpensive to motivate the supplier to improve his SR practices, then regardless of the signal, it is optimal for the manufacturer to do so with a high (and uniform) investment, $\hat{\beta}_D$. We focus here on the more interesting case when this strategy is too costly for the manufacturer.

turns out to be higher than the supplier's actual SR performance. The manufacturer's optimal disclosed level $s_D^*(v, \tilde{s})$ is characterized in the following theorem.

Theorem 3. (a) *If $\beta_{LD} \geq \hat{\beta}_D$, then for any v and \tilde{s} , $s_D^*(v, \tilde{s}) = s^*(m, \hat{\beta}_D)$.*

(b) *If $\beta_{LD} < \hat{\beta}_D$ and $\alpha(1-q)(r-w)\gamma'(M) \geq pq$, then for any level of visibility v and signal \tilde{s} , $s_D^*(v, \tilde{s}) = M$.*

(c) *If $\beta_{LD} < \hat{\beta}_D$ and $\alpha(1-q)(r-w)\gamma'(M) < pq$, then for a given level of visibility v , there exist two thresholds $\tau_L(v)$ and $\tau_H(v)$ satisfying $m < \tau_D \leq \tau_L(v) < \tau_H(v) \leq M$, such that $s_D^*(v, \tilde{s})$ is as follows:*

Region (i) If $\tilde{s} < \tau_L(v)$, then $s_D^(v, \tilde{s}) \in [s^*(\tilde{s}, \beta_D^*(v, \tilde{s})), M)$. In addition, $s_D^*(v, \tilde{s})$ is either constant or piece-wise constant in \tilde{s} , with at most one discontinuous drop at $\tilde{s} = \tau_D$.*

Region (ii) If $\tau_L(v) \leq \tilde{s} \leq \tau_H(v)$, then $s_D^(v, \tilde{s}) = s^*(\tilde{s}, \beta_D^*(v, \tilde{s})) = \tilde{s}$.*

Region (iii) If $\tilde{s} > \tau_H(v)$, then $s_D^(v, \tilde{s}) = \tau_H(v) < s^*(\tilde{s}, \beta_D^*(v, \tilde{s})) = \tilde{s}$.*

Furthermore, $\tau_H(v) < M$ if and only if $\alpha(1-q)(1-v), (r-w)\gamma'(M) < pq(1-v)$.

Theorem 3(a) captures a scenario when an overstated disclosure is very costly to the manufacturer (e.g., under a high penalty factor p or a high possibility of third-party disclosure q). In this case, the manufacturer is overly cautious. She invests the maximum amount $\hat{\beta}_D$ regardless of the supplier type but discloses the minimum possible final level of SR given her investment; i.e., $s^*(m, \hat{\beta}_D)$. In contrast, Theorem 3(b) captures the opposite scenario when the potential penalty associated with an overstated disclosure is very small (e.g., under low p and q). As a result, the manufacturer discloses the maximum possible final level of SR, M , regardless of the observed signal \tilde{s} . Note that, by Assumption 3, the supplier never increases his SR level to $s > M$ and thus the manufacturer cannot credibly disclose $s_D > M$.

In our subsequent analysis, we will focus on case (c) of Theorem 3, which is illustrated in Figure 3-4. First note that in our setup, given $v > 0$ and signal \tilde{s} , the most likely value of s_0 from the manufacturer's perspective is $s_0 = \tilde{s}$. Thus, the

value $s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$ represents the final SR level of the most likely supplier type from the manufacturer's standpoint. We refer to this value as the manufacturer's *best estimate* of the supplier's final SR level. When the manufacturer observes a low signal (i.e., Region (i) of Theorem 3(c) and Figure 3-4), the manufacturer chooses to potentially overstate the final level of SR with respect to her best estimate ($s_D^*(v, \tilde{s}) \geq s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$). This is in part because with a low signal, the actual supplier type is likely to be higher than the signal (i.e., $s_0 > \tilde{s}$), since the latter can be wrong. Thus, the risk of overstating and the associated potential penalty are small relative to the benefit of attracting higher demand from SC consumers with a higher s_D .

The manufacturer's decision when she observes a high signal (i.e., Region (iii) of Theorem 3(c) and Figure 3-4) is exactly the opposite. The manufacturer now faces a high risk from overstating if the signal turns out to be wrong (and hence the actual supplier type is lower). Therefore, she chooses to be conservative and understates the final level of SR compared to her best estimate ($s_D^*(v, \tilde{s}) < s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$). Finally, when the manufacturer observes a signal that lies in the intermediate range (i.e., Region (ii) of Theorem 3(c) and Figure 3-4), she addresses the tradeoff between increasing demand and potentially incurring a penalty by disclosing her best estimate, $s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$, to the consumers. In this region, the optimal disclosed level is exactly equal to \tilde{s} . This is because the manufacturer invests $\beta_D^*(v, \tilde{s}) = \beta_{LD}$ for $\tilde{s} > \tau_D$. The best response of a supplier with type $s_0 = \tilde{s} > \tau_D$ is not to improve SR, but to simply stay at \tilde{s} (Proposition 12).

3.4.2 The Effects of Visibility and Third-Party Scrutiny on the Manufacturer's Disclosure

Our next result summarizes the effect of visibility on the optimal SR level disclosed by the manufacturer.

Proposition 5. (a) $\tau_L(v)$ is strictly decreasing in v ; $\tau_H(v)$ is non-decreasing in v .

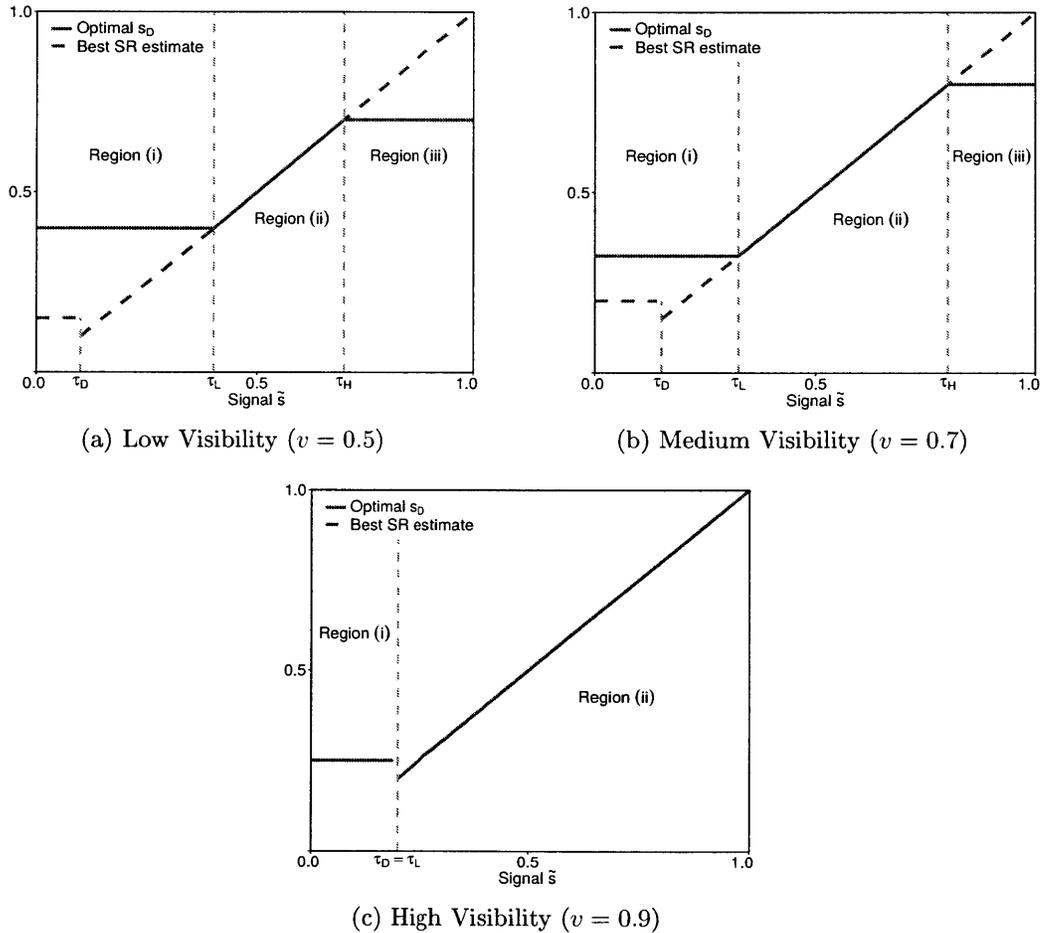
(b) The manufacturer's optimal disclosed level $s_D^*(v, \tilde{s})$ is

- (i) non-increasing in v if $\tilde{s} < \tau_L(v)$;
- (ii) independent of v if $\tilde{s} \in [\tau_L(v), \tau_H(v)]$; and
- (iii) increasing in v if $\tilde{s} > \tau_H(v)$.

Proposition 5(a) implies that the range of \tilde{s} values for which it is optimal for the manufacturer to disclose her best estimate of the supplier’s final SR level widens as visibility increases (see Region (ii) in Figures 3-4a to 3-4c). That is, as the manufacturer becomes more confident that the signal she observes accurately captures the supplier’s initial level of SR, she is more likely to be “truthful” in her disclosure (i.e., neither overstate nor understate relative to her best estimate).

For the range of \tilde{s} values where the manufacturer finds it optimal to overstate

Figure 3-4: Optimal Disclosed SR Level by the Manufacturer



relative to her best estimate (i.e., Region (i) in Figure 3-4), the optimal disclosed level decreases as visibility increases (Proposition 5(b)-(i)). Since the manufacturer's best estimate mostly increases with v in this region, the decreased disclosed level implies that the extent to which she overstates decreases as visibility increases.⁵ This is because as the manufacturer becomes more certain what she observes is the true supplier type, the benefit of overstating decreases and she becomes more cautious about the risk of a potential penalty. Finally, for the range of \tilde{s} values where the manufacturer finds it optimal to understate relative to her best estimate (i.e., Region (iii) in Figure 3-4), the optimal disclosed level increases as visibility increases (Proposition 5(b)-(iii)). Hence, the extent to which the manufacturer understates decreases with visibility. Again, the manufacturer's increased certainty that the signal is correct makes her less concerned about the potential penalty risk. Instead, she finds it beneficial to disclose a level closer to her best estimate. Observe that when visibility is sufficiently high, Region (iii) no longer exists and thus, the manufacturer does not understate the supplier's SR performance for any \tilde{s} value (see the last condition in Theorem 3 and Figure 3-4c).

To summarize, greater visibility motivates the manufacturer to be more "truthful" in her disclosure by (i) reducing overstatement in general when she observes a low signal; (ii) expanding the range of signals for which disclosing her best estimate is optimal; and (iii) reducing or eliminating understatement when she observes a high signal.

Next, we examine how the probability of third-party scrutiny, q , affects the manufacturer's optimal disclosed SR level. Define $\Omega(v, \tilde{s}) \equiv s_D^*(v, \tilde{s}) - s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$ as the difference between the manufacturer's optimal disclosed level and her best estimate of the supplier's final SR level. We have the following result.

Proposition 6. *For any given v and \tilde{s} , $\Omega(v, \tilde{s})$ is non-increasing in q . Furthermore, τ_H is non-increasing in q , and strictly decreasing in q if $\tau_H < M$.*

⁵One possible exception to this result occurs for \tilde{s} values close to the threshold τ_D , when τ_D is decreasing in visibility. In this case, the manufacturer's best estimate decreases with v for some $\tilde{s} < \tau_D$ because the optimal investment drops from $\bar{\beta}_{HD}$ to β_{LD} . Thus, the gap between the optimal disclosed level and the best estimate can increase with visibility.

As the probability of third-party scrutiny increases, the manufacturer becomes more cautious in her disclosure. For low signals (i.e., $\tilde{s} < \tau_L$; Region (i) in Figure 3-4), the manufacturer discloses a level that is closer to her best estimate. Conversely, for high signals (i.e., $\tilde{s} > \tau_H$; Region (iii) in Figure 3-4 if it exists), the manufacturer understates with respect to her best estimate even more as q increases. Finally, τ_H being non-increasing in q (and strictly decreasing when Region (iii) exists) implies that as q increases, the manufacturer understates the supplier's SR performance for a wider range of high signals.

With respect to the effects of v and q on the supplier's final level of SR, our findings in the Manufacturer Disclosure setting are qualitatively the same as in the No Manufacturer Disclosure setting, and hence, will not be repeated.

3.5 When Does the Manufacturer Prefer to Voluntarily Disclose SR Information?

We next investigate under what conditions the manufacturer prefers to disclose versus not to disclose SR information to the consumers. Define $\Pi_M^*(v, \tilde{s})$ and $\Pi_{MD}^*(v, \tilde{s})$ as the manufacturer's optimal expected profits in the No Manufacturer Disclosure and the Manufacturer Disclosure settings. For a given level of visibility v and signal \tilde{s} , the manufacturer will choose to disclose SR information to the consumers if and only if $\Pi_{MD}^*(v, \tilde{s}) > \Pi_M^*(v, \tilde{s})$. Theorem 4 summarizes how the difference between these two profits, $\Delta_\Pi(v, \tilde{s}) \equiv \Pi_{MD}^*(v, \tilde{s}) - \Pi_M^*(v, \tilde{s})$, depends on \tilde{s} .⁶

Theorem 4. *For a given level of visibility v , $\Delta_\Pi(v, \tilde{s})$ is continuous and*

- (i) *constant in \tilde{s} for $\tilde{s} \in [m, \tau]$;*
- (ii) *non-increasing in \tilde{s} for $\tilde{s} \in [\tau, \tau_D]$; and*
- (iii) *non-decreasing in \tilde{s} for $\tilde{s} \in [\tau_D, M]$.*

⁶Theorem 4 and our discussion hereafter consider the case when $\tau_H \geq \hat{s}$, with τ_H defined in Theorem 3. If instead $\tau_H < \hat{s}$, then the manufacturer generally prefers not to disclose, regardless of the \tilde{s} value. See Theorem 7 and the related discussion in Appendix B.1 for further details.

Therefore, $\tau_D \in \arg \min_{\tilde{s}} \{\Delta_{\Pi}(v, \tilde{s})\}$.

Theorem 4 implies the following observation regarding the manufacturer's choice of whether or not to disclose SR information.

Corollary 1. *The manufacturer is least likely to disclose SR information when the observed signal is close to τ_D .*

Theorem 4 can be interpreted as follows. When the observed signal is low (i.e., $\tilde{s} \in [m, \tau)$), the manufacturer infers that the supplier's current practices are poor and thus makes a high investment to improve the supplier's SR capabilities, regardless of whether or not she plans to disclose SR information. Note that for $\tilde{s} \in [m, \tau)$, $s^*(\tilde{s}, \bar{\beta}_H)$ is equal to $s_u^*(\bar{\beta}_H)$ in the No Manufacturer Disclosure setting and to $s_{uD}^*(\bar{\beta}_{HD})$ in the Manufacturer Disclosure setting (Proposition 12). Both of these values are independent of \tilde{s} . Therefore, neither $\Pi_M^*(v, \tilde{s})$ nor $\Pi_{MD}^*(v, \tilde{s})$ depends on \tilde{s} , and $\Delta_{\Pi}(v, \tilde{s})$ is constant in \tilde{s} (part (i)).

When the observed signal is in the intermediate range (i.e., $\tilde{s} \in [\tau, \tau_D)$), the manufacturer continues to make a high investment ($\bar{\beta}_{HD}$) if she plans to disclose SR information (and hence, $\Pi_{MD}^*(v, \tilde{s})$ remains constant in \tilde{s}). In doing so she reduces the potential penalty in case of overstating. If instead the manufacturer plans not to disclose SR information, then she prefers a low investment (β_L). Note that $s^*(\tilde{s}, \beta_L) = \tilde{s}$ for $\tilde{s} > \tau$ (Proposition 12). Thus, $\Pi_M^*(v, \tilde{s})$ is strictly increasing in \tilde{s} for $\tilde{s} \in [\tau, \hat{s})$ and constant thereafter (since the third party does not disclose any SR level above \hat{s}). Consequently, $\Delta_{\Pi}(v, \tilde{s})$ is non-increasing for $\tilde{s} \in [\tau, \tau_D)$ (part (ii)). Finally, when the observed signal is high (i.e., $\tilde{s} \in [\tau_D, M]$), in both settings the manufacturer chooses a low investment and her expected profit is first strictly increasing, then constant in \tilde{s} . Since the manufacturer can potentially increase demand by disclosing SR information, $\Pi_{MD}^*(v, \tilde{s})$ increases faster or at the same rate as $\Pi_M^*(v, \tilde{s})$. As a result, $\Delta_{\Pi}(v, \tilde{s})$ is non-decreasing in \tilde{s} for $\tilde{s} \in [\tau_D, M]$ (part (iii)).

Theorem 4 leads to five possible disclosure strategies for the manufacturer.

Corollary 2. *There exist five possible disclosure strategies for the manufacturer based on the signal \tilde{s} :*

(i) Do not disclose for all $\tilde{s} \in [m, M]$;

(ii) Disclose for all $\tilde{s} \in [m, M]$;

(iii) Do not disclose if $\tilde{s} \leq \kappa$ and disclose if $\tilde{s} > \kappa$, where $\kappa > \tau_D$;

(iv) Disclose if $\tilde{s} \leq \kappa'$ and do not disclose if $\tilde{s} > \kappa'$, where $\kappa' < \tau_D$;

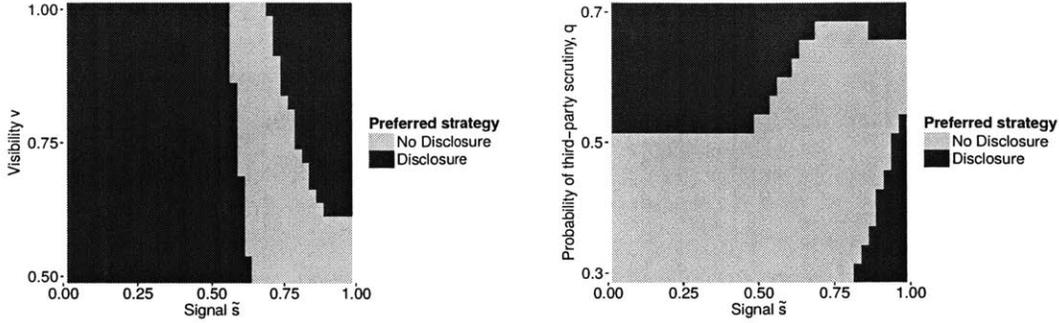
(v) Do not disclose if $\tilde{s} \in (\kappa_L, \kappa_H)$, and disclose if $\tilde{s} \leq \kappa_L$ or $\tilde{s} \geq \kappa_H$, where $\kappa_L < \tau_D < \kappa_H$.

In addition, strategies (i) and (ii) are more likely to be optimal under a low level of visibility; strategies (iii)–(v) are more likely to be optimal under a high level of visibility.

Strategies (i)–(iii) are relatively intuitive. For example, if visibility is low and both the possibility of third-party disclosure (q) and the penalty factor for overstating the SR level (p) are high, then the manufacturer may prefer to be cautious and not to disclose any information to the consumers, regardless of the signal (strategy (i)). Conversely, if q and/or p are very low, then the manufacturer may prefer to disclose SR information regardless of the signal (strategy (ii)) because the benefit from increased demand outweighs the risk of a potential penalty. Strategy (iii) can be optimal if visibility is high and q and p are not too low. In this case, it is preferable for the manufacturer to disclose only when \tilde{s} indicates that the supplier has good SR practices.

Strategies (iv) and (v) are less intuitive. Figure 3-5a presents a numerical illustration of when these strategies can be optimal. First observe that strategy (iv) is optimal when visibility is low (i.e., when $v < 0.6$) whereas strategy (v) is optimal when visibility is higher. Recall from Theorem 2 and §3.4 that the manufacturer makes a high (low) investment when the observed signal is low (high). The high investment (and the resulting high final SR level) at low signals motivates the manufacturer to disclose to increase demand (the left half of Figure 3-5a). In contrast, she is more cautious when observing a high signal, especially when visibility is low.

Figure 3-5: Manufacturer's Optimal Disclosure Strategy



(a) Optimal disclosure strategy as a function of the signal and visibility ($q = 0.6$)

(b) Optimal disclosure strategy as a function of the signal and third-party disclosure ($v = 0.6$)

Parameters: $\alpha = 0.4$, $r = 0.5$, $w = 0.15$, $c = 0$, $p = 17.5$. Functions: $\rho(s) = e^{1.75s}$, $\gamma(s) = 0.5 - e^s$, $\delta(\beta) = \frac{0.5\beta^2}{1-\beta}$. The supplier's type s_0 is distributed uniformly between $m = 0$ and $M = 1$.

In this case, the manufacturer is not very confident about the signal being correct and hence, faces a considerable risk of penalty from disclosing and overstating her supplier's SR performance. Hence, she refrains from disclosing (the lower right corner of Figure 3-5a) and saves on her investment cost (as her optimal investment is lower if she does not disclose; see Proposition 4).

With higher visibility, the manufacturer becomes more certain that a high signal indicates good performance at the supplier and therefore she chooses to disclose (the upper right corner of Figure 3-5a). However, for intermediate signals, the manufacturer still prefers not to disclose even under high visibility (the middle top region of Figure 3-5a). This is most likely to occur for $\tilde{s} \in (\tau, \tau_D)$, where τ and τ_D are the threshold signals at which the manufacturer decreases her investment under No Manufacturer Disclosure and Manufacturer Disclosure. Within this range, the manufacturer invests a high (low) amount of resources if she plans to disclose (not to disclose). The substantially higher investment cost under disclosure cannot be justified by the moderate increase in demand from disclosing. Thus, the manufacturer opts not to disclose and to instead save on her investment cost in this region.

3.5.1 The Effects of Visibility and Third-Party Scrutiny on the Manufacturer's Disclosure Choice

The reasoning behind strategy (v) above also explains why increasing visibility (i.e., going from the bottom to the top of Figure 3-5a) may induce the manufacturer to switch from disclosing to not disclosing for intermediate signals (e.g., the region when \tilde{s} is close to 0.6 in Figure 3-5a). This situation occurs when τ and τ_D change with visibility such that \tilde{s} begins to fall inside (τ, τ_D) . In particular, for some intermediate signals, as visibility increases, the manufacturer decreases her investment from high to low under No Manufacturer Disclosure because she is more certain that the supplier's initial SR level is acceptable. However, her investment remains high under Manufacturer Disclosure so that she can better support her claim in the case of third-party scrutiny and avoid the potential penalty. Thus, the substantial investment savings cause the manufacturer to switch from disclosing to not disclosing SR information as visibility increases.

Finally, Figure 3-5b shows how an increase in the probability of third-party scrutiny (q) can affect the manufacturer's choice between disclosing and not disclosing. We observe that for a large range of \tilde{s} values that are not too high (e.g., when $\tilde{s} \leq 0.75$ in Figure 3-5b), increased scrutiny motivates the manufacturer to switch from not disclosing to disclosing SR information to the consumers. This change is driven by the fact that increased scrutiny (i.e., a higher value of q) incentivizes the supplier to improve his SR performance (see Propositions 3 and 10). Therefore, the manufacturer finds it beneficial to disclose and increase demand. However, for higher values of \tilde{s} (e.g., $\tilde{s} > 0.8$ in Figure 3-5b), we observe that as q increases, the manufacturer switches from disclosing to not disclosing, and then back to disclosing. The first switch is because a higher q implies a larger potential penalty for disclosing and overstating the supplier's SR level. Hence, the manufacturer becomes more cautious and chooses not to disclose. As q further increases, the direct effect that q has on increasing the supplier's SR decision outweighs the potential penalty. As a result, the manufacturer again finds it beneficial to disclose SR information.

3.5.2 The Effect of the Manufacturer's Disclosure Choice on the Final SR Level

The manufacturer's decision of whether or not to disclose can have a significant impact on the supplier's final SR level. Specifically, the manufacturer's disclosure of SR information leads to an equal or better final SR level from the supplier, as compared to when she does not disclose.

Proposition 7. *For any supplier type s_0 , the final level of SR in the Manufacturer Disclosure setting is greater than or equal to that in the No Manufacturer Disclosure setting. Furthermore, there exists $\tau_{SR} < \tau_D$ such that it is strictly greater for $s_0 \in [m, \tau_{SR}]$.*

Proposition 7 follows directly from the result that the manufacturer's optimal investment is strictly higher under Manufacturer Disclosure than under No Manufacturer Disclosure (Proposition 4). Moreover, it applies both for any given (and fixed) signal \tilde{s} and in expectation over all possible signal values.

As discussed in §3.5.1, an increase in visibility or the probability of third-party scrutiny may lead the manufacturer to switch from disclosing to not disclosing SR information (see Figure 3-5). Therefore, an increase in either v or q can result in a *decrease* in the final SR performance of the supplier, as summarized below.

Corollary 3. *An increase in supply chain visibility, v , or the probability of third-party scrutiny, q , can lead to a decrease in the supplier's final SR level if such an increase motivates the manufacturer to prefer not disclosing SR information over disclosing.*

This result is noteworthy given our findings that under the individual cases of either No Manufacturer Disclosure or Manufacturer Disclosure, increasing third-party scrutiny *never* decreases the supplier's final level of SR (Propositions 3 and 10). In other words, when disclosure is not mandatory but instead an option for the manufacturer, intensifying third-party scrutiny can result in undesirable consequences and actually hurt the level of SR in the supply chain.

3.6 Conclusions and Managerial Insights

Companies are increasingly facing pressure from consumers and external stakeholders to guarantee good SR practices in their supply chains. However, most companies do not have good visibility into their suppliers' practices. Our research provides guidance to manufacturers on how to improve a supplier's SR performance under incomplete supply chain visibility. Specifically, we examine two decisions for a manufacturer – her investment to improve a supplier's SR capabilities and her disclosure (or not) of SR information to consumers. We show how these decisions are impacted by the manufacturer's available information about the supplier's current practices, the level of supply chain visibility she has, and the probability of third-party scrutiny. To capture the social impact of the manufacturer's actions, we investigate the effect of her decisions on the resulting SR performance of the supplier.

We conclude the paper by discussing three specific insights on (i) a manufacturer's strategy for investing in a supplier's SR capabilities, (ii) how the manufacturer's investment strategy should change when she discloses SR information to consumers, and (iii) the impact of the manufacturer's investment and disclosure decisions on the supplier's SR performance.

Insight 3.1. *A manufacturer should invest a high (low) amount of resources to improve a supplier's SR capabilities if the information she observes suggests poor (good) practices. Greater visibility into the supply chain helps the manufacturer to be more efficient with her investment.*

Our analysis demonstrates when a manufacturer should provide a supplier with significant support to improve his SR practices. Specifically, if the supplier appears to currently have poor practices, then the manufacturer should invest a large amount of resources to improve the supplier's capabilities. If instead the supplier appears to already have good practices, then the manufacturer does not need to make a significant investment. The higher the level of visibility the manufacturer has into her supplier's practices, the more she can trust her available information, and thus, better recognize whether the supplier truly needs support. This can then help the

manufacturer to tailor the extent of her investment and increase (decrease) it as she becomes more certain that the supplier's practices are poor (good).

Our insight aligns with investment strategies observed in practice. For example, Starbucks, who has extensive visibility into its supply chain, focuses most of its efforts toward improving the practices of disadvantaged coffee bean farmers in developing countries (Starbucks Corporation 2015). More generally, a recent study by the Organization for Economic Cooperation and Development (OECD) and the World Trade Organization (WTO) found that over 65% of the 219 companies surveyed engaged in development activities with suppliers in developing countries. Of these activities, more than 40% were driven by the companies' SR agendas, with 31% of the companies citing better working conditions as one of the main results achieved (OECD/WTO 2013pp. 111–113).

Insight 3.2. *If a manufacturer plans to disclose SR information to consumers, then her investment in the supplier's SR capabilities should be more aggressive as compared to when she plans not to disclose. This more aggressive strategy when the manufacturer discloses leads to better SR practices by the supplier.*

If the manufacturer plans to disclose, then she should follow the same approach outlined in Insight 1 but with a more aggressive investment (as compared to if she does not disclose). This includes (i) always increasing the amount of resources she offers to a supplier, and (ii) investing significant resources in a supplier whose current practices appear to be “good enough,” and hence, in whom she would not invest significantly if she were not to disclose. With a more aggressive strategy, the manufacturer can reduce the possibility and the potential impact of a third party uncovering her overstating a supplier's SR performance. The higher level of investment also ensures that the resulting SR performance of the supplier is always greater when the manufacturer discloses information to consumers.

Insight 2 suggests that any initiative which requires manufacturers to disclose SR information to consumers (e.g., government regulation such as the California Transparency in Supply Chains Act) will have a positive impact on suppliers' SR practices. However, manufacturers typically have the freedom to choose whether

or not to disclose SR information to consumers. Our last insight highlights when a manufacturer having this choice may negatively impact the SR practices of a supplier.

Insight 3.3. *Voluntary disclosure by a manufacturer is least effective in improving the SR practices of a supplier when the supplier appears to currently have only average practices.*

We find that a manufacturer is most likely to prefer not disclosing SR information when a supplier appears to have only average SR practices (hereafter referred to as an average supplier). This is in large part due to Insight 2 and the fact that if a manufacturer discloses, then she would need to offer an average supplier significant support. Conversely, if the manufacturer does not disclose, then she would not need to offer this same supplier as much support (see Figure 3-3). Due to this difference in investment strategy, the manufacturer may find it more economical to invest fewer resources and to not disclose SR information. Note that this is also why increased supply chain visibility may cause a manufacturer to prefer not to disclose when working with an average supplier. The improved visibility helps the manufacturer to confirm that there may be an opportunity to save costs by not disclosing SR information to the consumers and only making a low investment in the supplier's SR capabilities.

As previously discussed, the resulting SR performance of the supplier is always higher when the manufacturer discloses. Therefore, anything that causes the manufacturer to prefer not to disclose (rather than to disclose) has the potential to decrease the supplier's level of SR. For example, we show that an increase in third-party scrutiny improves SR when the manufacturer is required to disclose. However, increased scrutiny may discourage the manufacturer from disclosing (due to the threat of penalty from disclosing incorrect information) and hence hurt the supplier's level of SR if the manufacturer has a choice regarding disclosure.

We indeed observe examples of companies not focusing on their average suppliers' SR practices. For example, Inditex (the parent company of Zara) is considered a leader in terms of supply chain transparency (Fashion Revolution 2016). However, a 2016 report from the European workers' rights campaign *Labour Behind the Label* found that footwear workers at Zara's suppliers in Albania "were earning as little as

49p [pound pennies] an hour including overtime” (The Guardian 2016). Interestingly, Inditex’s 2015 annual report primarily discusses its work with suppliers operating in developing countries such as Bangladesh, Cambodia, and India, and does not discuss much about Inditex’s Eastern European suppliers (Inditex 2015). This lack of disclosure suggests that Inditex may not have invested enough in its average suppliers (i.e., those located in Eastern Europe) to improve their SR practices.

Supply chain transparency is an emerging topic both in practice and the academic literature. Given the lack of visibility into the SR practices of their suppliers, many companies are now facing the challenge of (i) how to invest resources to improve suppliers’ SR performance and (ii) deciding whether and what SR information to disclose to consumers. We hope that our work will motivate other researchers to continue investigating how to improve suppliers’ SR practices under incomplete supply chain visibility.

Chapter 4

The Effect of Visibility on Consumer Trust of Social Responsibility Disclosure

4.1 Introduction

Consumer trust is key to the success of companies' social responsibility (SR) communication strategies. According to a 2015 Nielsen survey of more than 30,000 consumers, brand trust tops the list of factors that influence socially responsible purchases (Nielsen 2015). Similarly, in a recent meta-analysis of empirical studies, Joshi and Rahman (2015) identify trust in a company as a key variable that positively impacts the purchase of SR products. Conversely, they identify lack of trust (or distrust) in a company's claims about SR as a significant barrier towards such purchases.¹

In this context, supply chain transparency has been identified as one of the most effective ways for companies to improve trust among consumers. For example, in a recent poll, 75% of respondents considered transparency to be effective in building trust between businesses and consumers (KPMG and Global Reporting Initiative

¹We follow the European Commission's definition of social responsibility as "[companies integrat- ing] social and environmental concerns in their business operations and in their interaction with their stakeholders on a voluntary basis" (Dahlsrud 2008). We specifically focus on social concerns in social responsibility.

2014). Similarly, Marks & Spencer identified “building trust” as the main benefit of transparency for companies (Marks & Spencer and Globescan 2015). However, these studies focus on the effect that disclosing information to consumers has on trust, and assume that companies already have good information about SR practices in their supply chains. In practice, to create a transparent supply chain, a company must not only determine what information to disclose to consumers, but also gain visibility into its supply chain (New and Brown 2011). In this paper, we employ an incentivized human-subject experiment to study whether and how *supply chain visibility* impacts consumers’ trust in companies’ communications of SR.

We choose to focus on visibility because in practice, many companies do not have good visibility into their supply chains. A recent study by The Sustainability Consortium found that 81% of the 1,700 companies surveyed lacked full visibility into the SR practices of their supply chains and 54% had no visibility at all (The Sustainability Consortium 2016). Furthermore, gaining visibility requires a significant investment of time and resources (Doorey 2011). The current SR literature on trust and supply chain transparency has focused on the effect that disclosure has on consumer trust (MacLean and Rebernak 2007, Kang and Hustvedt 2014e.g.), while the effect of visibility on trust is not well understood. We address this gap by studying when a company that has greater visibility into its suppliers’ SR practices can enjoy a higher level of consumer trust.

To enhance our understanding of consumer trust, we also investigate how our results vary with consumers’ heterogeneity. Examining consumer heterogeneity is important because consumers’ valuations of a company’s SR practices can vary widely (e.g., García-Gallego and Georgantzís 2011). We focus on one particular aspect of heterogeneity – consumers’ *prosociality*. We define an individual’s prosociality as the extent to which the individual is willing to sacrifice his/her own benefit to improve the payoff of another. Prior studies have attributed the consumption of socially responsible products to values such as benevolence, equality, and responsibility (Vermeir and Verbeke 2006). Individuals who adhere to these values tend to be more prosocial. Hence, there is a natural connection between a person’s prosociality and his/her

attention to SR.

Based on the above aspects, we address the following research questions: (i) How does *supply chain visibility* affect consumers' trust in companies' SR communications? (ii) How does visibility impact the effect that trust has on consumers' *willingness-to-pay* for companies' products? (iii) How do the answers to these questions vary based on consumers' *prosociality*? We focus on the effect of supply chain visibility when a company communicates good (as opposed to poor) SR practices in its supply chain, as this is when a lack of trust among consumers is most likely to arise.

To address our research questions, we design a three-player game with the following context: A Worker helped a Firm (she) to make a product, and the Firm would like to sell the product to a Consumer (he). If the product is sold, then the Consumer and the Firm earn additional payments in addition to their endowments while the Worker may incur a decrease in his endowment (i.e., a negative externality). The sequence of events in our experiment is as follows. The Firm first receives a clue about whether the Worker is likely to incur a negative externality. She then conveys a message about the clue to the Consumer which may or may not be truthful. A greater level of supply chain visibility is represented by a higher probability that the clue the Firm receives is correct. Finally, based on the message received, the Consumer decides the maximum price that he is willing to pay (WTP) for the product. A higher WTP ensures that the product is more likely to be sold. We incentivize participants to elicit Consumers' and Firms' *truthful* WTP and message decisions, respectively. Employing an incentivized experiment is important because evidence exists that what individuals claim they would do with respect to SR differs from their actual purchase behavior. For example, Devinney et al. (2010p. 112) observe that "individuals either purposely overstate their social credentials or just want to look good in surveys, making it nearly impossible to believe what they say about their social proclivities."

We follow the definition of trust in Rousseau et al. (1998) as the "intention to accept vulnerability based upon positive expectations of the intentions or behavior of another." In our experiment, a Consumer who cares about the Worker's well-being

is vulnerable because he may cause a reduction in the Worker's payoff (via the negative externality associated with the product), particularly if he states a high WTP. Similarly, the Consumer may have a positive expectation that the Firm truthfully communicates her available information (i.e., the clue) to the Consumer.

Based on the above definition, we use two measures to study trust and its effect on the Consumer's WTP. First, to capture the Consumer's expectation about the Firm's behavior, we measure the Consumer's stated belief in the truthfulness of the message sent by the Firm with an incentivized question after he states his WTP. For simplicity, we refer to this measure as the Consumer's trust in the Firm's message. The use of beliefs questions to measure trust is a common technique in the experimental economics literature (see, e.g., Charness and Dufwenberg 2006, Rode 2010). Second, to measure the effect of trust (or lack thereof) on Consumers' WTP, we manipulate our experimental design and introduce a control condition where both the Firm and the Consumer directly observe the clue about the negative externality that the Worker may incur if the product is sold. There is no communication between the two players in the control condition. Thus, distrust of the Firm does not play a role in explaining the Consumer's WTP. By comparing Consumers' WTP between our main condition and this control condition, we can obtain a measure of the effect of distrust on Consumers' WTP (for a similar approach, see Özer et al. 2014).

Our research contributes to two growing streams of literature. First, regarding consumers' valuations of SR, an emerging body of empirical studies examine the impact of SR communications on consumer decision-making. For example, Pigors and Rockenbach (2016a) study consumers' purchase decisions in a laboratory experiment where the worker's wage is set by and known to the manager, but may or may not be known to consumers. Hainmueller et al. (2015) conduct a field experiment to study consumers' willingness to pay a premium for coffee with a fair trade label. Egels-Zandén and Hansson (2016) study the transition towards greater transparency from Nudie Jeans Co. and show that it increases consumers' willingness-to-buy the company's products. However, none of these studies investigate how disclosure affects trust. Kang and Hustvedt (2014) show that consumers' perceptions of efforts to dis-

close SR information significantly improve consumers' trust in companies. Similarly, MacLean and Rebernak (2007) identifies transparency in sustainability reporting (i.e., disclosure) as the best way to build trust among stakeholders. We contribute to this literature by studying how the equally-important but understudied dimension of supply chain visibility affects consumer trust. The paper that is closest to ours and that also studies supply chain visibility is Kraft et al. (2017). However, in their design the Consumer and the Firm always have the same information about the Worker's payoff, as they do not study consumers' trust in companies' communications.

Second, our work also adds to a growing experimental literature that examines the value of information in supply chain management. For example, Kalkan et al. (2016) compare mandatory versus voluntary disclosure of a company's social and environmental impacts and show that voluntary disclosure is perceived more positively by consumers. Croson and Donohue (2003, 2006) study information sharing among players in the beer game and show that it can partly alleviate the bullwhip effect. Regarding the combined effects of information and trust, Özer et al. (2011) use a cheap-talk game to study demand information sharing between a manufacturer and a supplier. They show that, because of suppliers' trust and manufacturers' trustworthiness, channel efficiency is greater than what is predicted by standard game theory models. Similarly, Inderfurth et al. (2013) show through a lab experiment that sharing information about holding costs between a supplier and a manufacturer helps to reduce supply chain inefficiencies if there exists trust among participants. We contribute to this literature by studying how the extent to which a company can be certain about the SR practices occurring in its upstream supply chain can help to increase consumer trust in the company's SR communications.

Our insights demonstrate how a company can improve consumer trust through supply chain visibility and why it is important to account for consumers' heterogeneity when developing SR communications. For example, we find that high prosocial consumers' trust is not affected by visibility and their willingness-to-pay is not affected by their trust or distrust in companies' claims about SR practices. Instead, since there is always a positive probability of poor SR practices, high prosocial consumers are

generally reluctant to pay a high price for the product. Conversely, low prosocial consumers trust companies' communications more when companies have greater visibility into the SR practices in its supply chain. Consequently, these consumers are also willing to pay more for products when there is higher supply chain visibility. Our results thus help companies understand when they can expect to improve consumer trust and derive benefits from increased visibility.

4.2 Experimental Design and Procedure

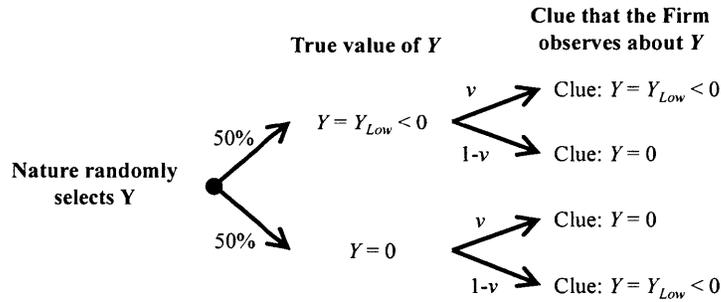
Our experimental design is composed of three parts: the Consumer Trust Game (CTG); a control task where we elicit participants' prosociality; and a postexperiment survey. In the CTG, we operationalize SR through monetary payments to Worker participants. The use of monetary transfers to study social preferences and their implications beyond monetary transactions is common in the experimental economics literature (e.g., Ostrom et al. 1992, Kosfeld et al. 2009) and in the recent experimental literature on socially responsible production (e.g., Bartling et al. 2015, Pigors and Rockenbach 2016a). We design our experiment to examine two measures – consumers' trust of companies' SR communications and the impact of trust on consumers' willingness-to-pay. We next discuss our design and experimental procedures in detail.

4.2.1 The Consumer Trust Game: Base Setup

We design the CTG as a three-player game with the following roles: a Firm (she), a Consumer (he), and a Worker. To minimize any pre-experiment biases related to specific SR practices, in the experiment we call these roles Players A, B, and C, respectively. At the beginning of the game, all participants are endowed with 150 experimental tokens and they are given the following background story: "Player C has helped Player A to make a product, and Player A wants to sell the product to Player B." No further information is provided about the context of our study and we only refer to a generic product with no specific attributes. In addition, all participants

are informed that if the product is sold, then the Worker will receive a total payoff equal to $150 + Y$. The value of Y is randomly selected by Nature and it can be either $Y = 0$ or $Y = Y_{Low} < 0$, with both values equally likely. The parameter Y_{Low} is equal to one of seven possible values in the set $\{-140, -120, -100, \dots, -20\}$ and it is known to all participants. However, only the Worker observes the true value of Y . With this setup, the sequence of events is as follows:

Figure 4-1: Relationship between the true value of Y and the clue

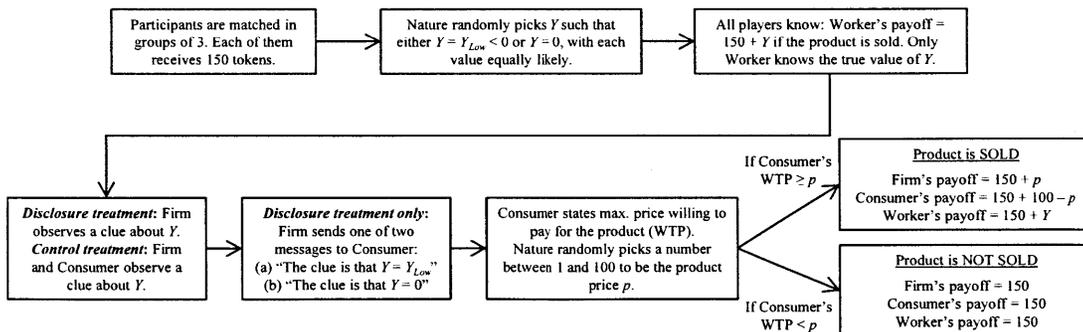


- (1) The Firm observes a *clue* about Y that is correct with probability v . For example, if the Firm observes “The clue you observe is that $Y = 0$ ” and $v = 0.6$, then the Firm knows that there is a 60% chance that the clue is correct and $Y = 0$; and there is a 40% chance that the clue is incorrect and instead $Y = Y_{Low}$. Figure 4-1 summarizes how the value of Y and the clue are generated. The value of v , which is also known to the Consumer and the Worker, represents the Firm’s visibility into the social responsibility practices in her supply chain. For simplicity, we hereafter refer to the clue as a binary variable $\text{Clue} \in \{Y_{Low}, 0\}$.

- (2) The Firm sends a *message* to the Consumer about the clue she observed. This message can be truthful or not. For example, if the Firm receives $\text{Clue} = Y_{Low}$, then she can send to the Consumer the truthful message “The clue I observe is that $Y = Y_{Low}$ ”, or she can instead send him the message “The clue I observe is that $Y = 0$.” We refer to the message as a binary variable $\text{Msg} \in \{Y_{Low}, 0\}$.

- (3) The consumer selects the maximum price that he is willing to pay for the product. We call this decision his *willingness-to-pay* (WTP). The Worker does not make any decisions in our game.
- (4) After the Firm and the Consumer make their decisions, Nature selects the product price, p , from a uniform distribution between 1 and 100 tokens. Two cases are then possible. First, if p strictly greater than the Consumer's WTP, then the product is not sold and all players earn their initial endowments of 150 tokens. Otherwise, if p is less than or equal to the Consumer's WTP, then the product is sold at price p and the players' final payoffs are: (i) the Consumer receives a consumption utility equal to 100 tokens and pays the product price to the Firm, earning $\pi_C = 150 + 100 - p$; (ii) the Firm receives the product price, earning $\pi_F = 150 + p$; and (iii) the Worker earns $\pi_W = 150 + Y$. The true value of Y and the Worker's payoff are never revealed to the Firm or the Consumer. With the exception of the value of Y , which is private knowledge, all the game dynamics and players' payoff structure outlined above are common knowledge. Figure 4-2 summarizes our experimental design and the players' payoffs.

Figure 4-2: Summary of Experimental Design



Our use of a random price to determine whether the product is sold is known as the Becker-DeGroot-Marschak or BDM mechanism (Becker et al. 1964). It is commonly used in economic experiments (e.g., Klos et al. 2005, Halevy 2007) and it ensures the truthful elicitation of Consumers' WTP. We set the maximum possible price of the

product at 100 tokens because of evidence from the literature. Specifically, Bohm et al. (1997) show that setting an upper bound that exceeds what “any real buyer would be willing to pay” increases participants’ valuations versus a market setting where buyer-participants confront prices set by seller-participants. Conversely, the prices they obtain when the upper bound does not exceed such a realistic maximum value are not significantly different between the two elicitation mechanisms. In our case, the most we should expect a real buyer to be willing to pay for the product is 100 tokens, the consumption utility that the Consumer derives from buying the product.

We follow the approach in Bartling et al. (2015) and consider $Y = Y_{Low} < 0$ to capture a long-term negative externality on workers that results from non-socially responsible production. That is, Y *does not* represent the Worker’s wage for his work on the product. Instead, lower values of $Y = Y_{Low}$ capture poorer treatment of workers in the supply chain. Conversely, if $Y = 0$, then socially responsible practices ensure that no negative externality is associated with the product. Similar to Bartling et al. (2015) and to Rode et al. (2008), we also consider that the potential decrease in the Worker’s payoff occurs *only if* the product is sold.

Measuring Consumer Trust: To measure the Consumer’s *trust* in the message sent by the Firm, the Consumer answers the following belief question: “What do you think is the actual clue observed by Player A [the Firm]?” He is presented with this question immediately after selecting his WTP in the CTG and he selects from two possible answers, Y_{Low} or 0. To motivate participants to truthfully state their beliefs, the Consumer receives 10 additional tokens if his answer is correct; i.e., if it matches the clue actually observed by the Firm (see Gächter and Renner 2010).² The Consumer trusts the Firm’s message if he believes that the clue is equal to the message (i.e., if he believes that $Msg = Clue$); and the Consumer distrusts the Firm otherwise (i.e., if he believes that $Msg \neq Clue$). This approach of measuring trust through the belief

²Because participants play the CTG for five rounds and are not paid until the end of the experiment (see §4.2.4), the Consumer cannot use this payment to infer what clue the Firm actually observed in a given round. Furthermore, the Firm does not know that the Consumer answers this question and hence the Firm’s decisions is not affected by it.

that someone else truthfully shares her private information is well established in the experimental literature (see, e.g., Charness and Dufwenberg 2006, Rode 2010). To measure the *effect of trust on WTP*, we use the experimental manipulations described in the next subsection.

4.2.2 The Consumer Trust Game: Experimental Manipulations

We manipulate the CTG in two dimensions. First, to study the impact of trust (or lack thereof) on the Consumer's WTP, we manipulate the information that the Consumer observes about the value of Y . In the *Disclosure condition*, described in §4.2.1, the Firm observes a clue about the value of Y and then sends a message about this clue to the Consumer. The Consumer only observes the message sent by the Firm and therefore he does not know whether this message is equal to the clue that the Firm actually observed. As a result, distrust in the Firm's communication can affect the Consumer's WTP. In contrast, in the *Control condition* both the Firm and the Consumer directly observe the clue about Y . Thus, the Firm does not make a decision in this condition and she does not send a message to the Consumer (see Figure 4-2). Since in the Control condition both players have the same information about Y , the Consumer's WTP cannot be influenced by distrust in the Firm's communication. Thus, differences in WTP between the Disclosure and Control conditions measure the effect of distrust on Consumers' WTP. The use of information asymmetry to quantify the effects of trust and distrust is common in the experimental economics and psychology literatures (see, e.g., Gneezy 2005, Cohen et al. 2009, Özer et al. 2014).

The second dimension that we manipulate in our experiment is the level of visibility v that the Firm has about the Worker's payment, Y . Specifically, we consider two experimental conditions: *Low Visibility*, corresponding to $v = 0.6$, and *High Visibility*, corresponding to $v = 0.9$. The parameter v captures the Firm's level of visibility because the greater its value, the more confident the Firm can be that the true value

of Y is equal to the clue (see Figure 4-1). The level of accuracy of a clue or signal has been previously used in both the modeling and experimental literature to capture visibility (e.g., Anctil et al. 2004, Morris and Shin 2004).

We study the two visibility conditions in both the Disclosure and Control conditions, which yields a 2×2 factorial design. Since the Consumer also observes the value of v , this experimental design allows us to address our main research questions, namely: (i) what is the *effect of visibility on Consumer trust*, by studying how visibility affects the Consumer's belief that $\text{Clue} = \text{Msg}$ in the Disclosure condition; and (ii) *how is the impact of trust on Consumer WTP affected by visibility*, by comparing the differences in WTP between the Disclosure and Control conditions under low visibility versus under high visibility.

4.2.3 Control Task and Postexperiment Survey

In addition to the CTG, our experiment also includes a control task to measure Consumers' and Firms' prosociality. We study participants' prosociality because of its natural connection with a person's reasons to care about SR practices (Vermeir and Verbeke 2006). To measure prosociality, we use the dictator game (Forsythe et al. 1994). In the game, Consumers and Firms play the role of the dictator (he). Each of them is matched with another player, the recipient. The dictator receives an initial endowment of 100 tokens, the recipient does not receive any tokens upfront, and the dictator is asked to decide how many tokens, t , he is willing to transfer to the recipient. The dictator's decision is thus a measure of his prosociality: a higher value of t means that the dictator is willing to spend a greater amount to improve the recipient's payoff. Therefore, he is more prosocial (see, e.g., Henrich et al. 2006, Gummerum et al. 2010).

After participants finish the CTG and the control task, they complete a postexperiment survey. The survey has four parts: (i) questions about players' decisions in the CTG; (ii) a trust and trustworthiness questionnaire; (iii) questions about participants' attitudes towards SR; and (iv) a demographic survey. We use the first part of this survey to gain insights into how participants make their decisions in the CTG.

In part (ii), we use Likert-scale questions that are commonly used in the literature to study individuals' attitudes towards trust and trustworthiness (Yamagishi and Yamagishi 1994, Yamagishi et al. 2015, Smith et al. 2011, Glaeser et al. 2000, Evans and Revelle 2008, Mayer and Davis 1999). Part (iii) consists of a series of Likert-scale and binary-choice questions that help us assess participants' general attitudes towards SR. Finally, we use part (iv) to obtain demographic and background information about our participants.

4.2.4 Experimental Procedures

All experimental sessions were run in the computer laboratory of a large university in the United States. In each session, we ran one of the four different treatments from our 2 (Disclosure vs. Control condition) \times 2 (High vs. Low Visibility) factorial design. The tasks were always performed in the following order: (i) five rounds of the CTG; (ii) the dictator game; and (iii) the postexperiment survey. Participants were informed at the beginning of each session that they would perform multiple tasks and a survey, but they were not provided with the details of a task until they had completed the previous one. We followed this procedure to prevent participants from altering their current decisions in anticipation of future decisions. Similarly, participants only observed the outcomes of the tasks after they had been completed. This procedure eliminated the possibility that outcomes from previous tasks or rounds affected current and future decisions (e.g., due to income effects). These methods are commonly used in the experimental literature (e.g., Bolton and Katok 1998, Luhan et al. 2009).

At the beginning of each round of the CTG, participants were provided with their initial endowments of 150 tokens each and they were randomly and anonymously matched into groups of three. Within each group, participants were randomly assigned to the roles of Player A (Firm), Player B (Consumer), or Player C (Worker). They were informed of the treatment conditions they were in, that they would play the CTG for a total of five rounds, and that they would stay in the same role for all five rounds. They were presented with detailed instructions, including the fact

that the value of Y_{Low} could vary across rounds³ and that they would be randomly and anonymously rematched into groups of three at the beginning of each round. To ensure that participants correctly understood these instructions, they were provided with examples and a reference sheet. They also had to correctly answer a set of practice questions before making their decisions.

After the five rounds of the CTG, participants performed the dictator game described in §4.2.3. In the dictator game, participants who played either the Consumer or the Firm role in the CTG were informed that they had been randomly and anonymously paired with another participant and they were asked to make decisions as the dictator (referred to as Player 1 in the instructions). The role of the recipient (referred to as Player 2) was performed by a participant who played the role of the Worker in the CTG. The relationship between participants' roles in the CTG and the dictator game was not known to participants.

Once the CTG and the dictator game were completed, participants were shown the outcomes of all tasks, including their own payoffs, their decisions, and the decisions made by the participants that they were matched with. Importantly, the true value of Y was never shown to Consumers and Firms, so the Consumer could never infer whether the Firm may have lied in a round of the CTG. Finally, all participants individually answered the postexperiment survey. Sample instructions and questions used in the postexperiment survey are available from the authors upon request.

All tasks were implemented using a web-based interface. During each session, participants were not allowed to communicate with each other and they only interacted via computer terminals. We had a total of 294 participants, all of which were students. 66.2% of them were undergraduates and the remaining 33.8% were graduate students. In addition, 62.5% of them identified as female, while 36.9% of them identified as male. Their average age was 21.2 years old, with a standard deviation of 2.6 years. In all experimental sessions, 50 tokens were worth 1 U.S. dollar. Participants earned an average of approximately \$22.85, with a minimum of \$15 and a maximum

³Consumers observed a different randomly selected value of Y_{Low} in each round; the other two players could observe the same value of Y_{Low} in multiple rounds.

of \$31. Each session typically lasted between 60 and 75 minutes. The total number of participants in the Disclosure (Control) condition was equal to 78 and 78 (72 and 66) in the High and Low Visibility conditions, respectively.

4.3 Hypotheses

Our experimental design allows us to study our main research questions, namely: (i) how Consumer trust in the Firm's message is affected by visibility; (ii) how the effect of trust on Consumers' WTP varies with visibility; and (iii) how the answers to these questions vary with Consumers' prosociality. In this section, we review some of the existing literature on these topics and present our main hypotheses.

We first note that all of our hypotheses focus on the case where $\text{Msg} = 0$ in the Disclosure condition (or $\text{Clue} = 0$ in the Control condition). This is because the Firm has a monetary incentive to sell the product and stating $\text{Msg} = 0$ is more likely to increase the Consumer's WTP, since $Y = 0$ would guarantee that the Worker's pay does not decrease. As a result, when the Consumer receives $\text{Msg} = 0$ in the Disclosure condition, he may believe this message not to be truthful; i.e., he may believe that $\text{Clue} = Y_{\text{Low}} < \text{Msg}$. In addition, our study of the case where $\text{Msg} = 0$ allows us to examine whether the Consumer trusts the Firm when she communicates good SR practices and how the Consumer's trust is affected by the Firm's level of visibility. In other words, this allows us to focus on the effect of supply chain visibility while keeping the disclosure of SR information constant.

The existing literature on individuals' trust of other parties' communications suggests that at least some individuals will distrust the other parties when these have monetary incentives to lie (see, e.g., the cheap-talk game in Gneezy 2005). Therefore, based on the above discussion, we hypothesize that in the Disclosure condition a fraction of Consumers significantly greater than zero distrusts the Firm's message when $\text{Msg} = 0$ and believe instead that $\text{Clue} = Y_{\text{Low}}$.

Our first hypothesis also addresses the effect of visibility on Consumers' trust. Studies in both the SR (MacLean and Rebernak 2007, Kang and Hustvedt 2014e.g.)

and the experimental economics (e.g., Pigors and Rockenbach 2016b, Kanagaretnam et al. 2010) literatures have shown that a higher level of transparency can yield a higher level of trust in other parties' communications. However, these studies focus on transparency as it relates to the *disclosure* of information. We instead focus on the level of *visibility* that the Firm has about the impact on the Worker if the product is sold. To address this, we first note that as discussed in the above paragraph, lack of trust is driven by the Consumer anticipating that the Firm may not be trustworthy and may send a message $Msg = 0$ when $Clue = Y_{Low}$. Thus, the effect of visibility on trust is directly influenced by the effect that Consumers expect visibility to have on Firms' trustworthiness. A higher level of visibility makes it more difficult for the Firm to hide her action under the pretense of ignorance about the value of Y (see, e.g., Dana et al. 2007). Similarly, lying about the clue may affect the Firm's self-image more under high visibility than under low visibility (for the importance of self-image, see Bénabou and Tirole 2006). Finally, the more negative the impact from a lie, the more likely individuals are to be trustworthy (e.g., Gneezy 2005, Lundquist et al. 2009). A higher level of visibility when $Clue = Y_{Low}$ means that there is a greater chance that Y is strictly negative and hence the expected negative impact from lying about this clue is greater. Based on these results, Consumers can expect Firms to be more trustworthy in the High Visibility condition than in the Low Visibility condition and they can particularly trust $Msg = 0$ more under higher visibility. Therefore, we make the following hypothesis about Consumers' distrust in Firms' communications.

Hypothesis 4.1A. *At each level of visibility, there is a significant fraction of Consumers who distrust the message $Msg = 0$ in the Disclosure condition.*

Hypothesis 4.1B. *The fraction of Consumers who distrust the message $Msg = 0$ is smaller in the High Visibility condition than in the Low Visibility condition.*

Second, regarding the effect of trust on WTP, we note that if Consumers only cared about their own payoffs, then they would always state $WTP = 99$ or 100 in the CTG, as either of these values maximize their expected earnings. By stating $WTP = 99$ or 100 , Consumers would ensure they earn a strictly positive amount of

tokens for any $p < 100$. However, there is a vast literature on social preferences that shows that individuals often care about others' payoffs in addition to their own (see Fehr and Schmidt 2006 for a review). In the context of SR, numerous studies have also shown that consumers' WTP for a product, or their positive perceptions of a company, tend to increase (decrease) if the company demonstrates good (poor) SR practices (e.g., Sen and Bhattacharya 2001, De Pelsmacker et al. 2005). Therefore, in our experiment we expect Consumers' WTP to be (on average) lower than 100 tokens, as there is always a positive chance that $Y = Y_{Low}$ and thus that the Worker's payoff will be reduced if the product is sold. Furthermore, the more a Consumer cares about the Worker's payoff and the more he believes that $Y = Y_{Low}$, the lower his WTP should be. Building on this expected Consumer behavior, we next discuss our main hypothesis about the effect of trust on Consumers' WTP and how it may be affected by visibility.

To measure the effect of distrust on the Consumer's WTP, we note that in the Control condition, both the Consumer and the Firm observe the clue about the value of Y . Thus, distrust in the Firm does not affect the Consumer's WTP in this condition. Conversely, as discussed above, we expect distrust in the Firm's message to exist in the Disclosure condition when $Msg = 0$. This, in turn, can reduce the Consumer's WTP if the Consumer receives $Msg = 0$ but believes instead that $Clue = Y_{Low}$. To formally measure the impact of trust on Consumers' WTP, we define the *trust gap* $\Delta(v)$ as the difference between Consumers' WTP when they observe the clue $Clue = 0$ in the Control condition versus when they receive the message $Msg = 0$ in the Disclosure condition. That is, $\Delta(v) = WTP(\text{Control}, v, Clue = 0) - WTP(\text{Disclosure}, v, Msg = 0)$ for any fixed value of v . A positive value of $\Delta(v)$ suggests an impact of distrust on Consumers' WTP.

We note that the above approach is similar to the one used in recent papers to measure the effect of trust or distrust on individuals' decisions. For example, Özer et al. (2014) experimentally investigate whether and how trust and trustworthiness vary between participants in China and the United States, in the context of forecast information sharing in a supply chain. To measure trust, they compare suppliers' (the

trustors) decisions when they only observe manufacturers' (the trustees) reports of their private forecast information to suppliers' decisions when the forecast information is public. Similarly, in the deception game (Gneezy 2005), Player As have two possible actions but do not observe the payoffs associated with each of them. Player Bs observes the relationship between actions and payoffs and communicate (truthfully or not) with Player As to inform them of the action that will yield them the highest payoff. If Player As had complete information and wanted to maximize their own payoffs, then 100% of them would choose the profit-maximizing action. Thus, in both the economics and psychology literatures the difference between the percentage of Player As who follow Player Bs suggested actions in the deception game (i.e., under incomplete information) and 100% (i.e., the optimal decision under complete information) is used to capture distrust (see, e.g., Cohen et al. 2009, Dreber and Johannesson 2008).

Finally, based on the above discussion and the fact that we expect visibility to reduce distrust in the Firm's message (see Hypothesis 4.1B), we conjecture that distrust reduces Consumers' WTP and that this effect is diminished under greater visibility. Formally, our second hypothesis is as follows.

Hypothesis 4.2A. *At each level of visibility, distrust in the message $Msg = 0$ significantly impacts Consumers' WTP. That is, $\Delta(v = 0.6) > 0$ and $\Delta(v = 0.9) > 0$.*

Hypothesis 4.2B. *The impact of distrust on Consumers' WTP when $Msg = 0$ is smaller in the High Visibility condition than in the Low Visibility condition. That is, $\Delta(v = 0.6) > \Delta(v = 0.9)$.*

Our experiment also allows us to study how the effect of visibility on trust (and its effect on WTP) varies by Consumers' prosociality. Relevant to our study, de Kwaadsteniet et al. (2006) experimentally examine a common resource dilemma where a person's usage of the common resource reduces the availability of the resource to others (i.e., it leads to a negative externality). They compare participants' behavior when the size of the common resource is certain versus uncertain. They show that although all participants increase their usage of the common resource when the size

of the resource becomes uncertain, such an increase is much larger for low (versus high) prosocial participants. Similarly, Kraft et al. (2017) study a market situation where consumers state their WTP for a product that has a positive externality for a third party. They show that the motives for buying the product are significantly affected by uncertainty in the value of the externality only among consumers with low prosociality. Kanagaretnam et al. (2009), on the other hand, shows that individuals with higher prosociality are more trusting than individuals with lower prosociality. Nevertheless, researchers have not examined the interactions among visibility, trust, and prosociality. Investigating these interaction effects offers insights into (i) whether consumers trust companies more when they demonstrate a high level of supply chain visibility and (ii) how to account for prosociality when evaluating the potential benefits of supply chain visibility.

4.4 Experimental Results

In this section, we present the main results from the CTG. We first discuss how we incorporate Consumers' prosociality into our analyses. Then we analyze each of our hypotheses from §4.3. We conclude this section with a brief discussion of our results.

As will be discussed throughout this section, our results vary significantly by Consumers' prosociality. This not only affects the level of Consumers' WTP, but also how Consumers respond to our experimental manipulations. To formally study how prosociality affects our results, we first note that the distributions of the dictators' decisions in the dictator game are very similar across all four experimental treatments (Kruskal-Wallis test, $p = 0.81$). Therefore, we use the median dictator decision among all Consumers (equal to 20 tokens) to classify two types of Consumers. Those who transferred 20 or more tokens are classified as *high prosocial* and account for 66.8% of the sample (63 participants). Those who transferred strictly less than 20 tokens are classified as *low prosocial* and account for the remaining 43.2% of the sample (48 participants). In the remainder of this section, we present results for both the entire sample of Consumer participants and by prosocial type.

4.4.1 Hypothesis 4.1: Does Visibility affect Consumers' Trust?

We first analyze whether Consumers trust the message $\text{Msg} = 0$ in the Disclosure condition by studying their responses to the belief question “What do you think is the actual clue observed by Player A [the Firm]?” introduced in §4.2.1. In each visibility condition, we study the proportion of responses where the belief is that $\text{Clue} = Y_{\text{Low}}$ when $\text{Msg} = 0$. If this proportion is significantly greater than zero, then there exists a non-negligible level of distrust in the Firm’s message.

Table 4.1: Percentage of observations where the Consumer’s belief is that $\text{Clue} = Y_{\text{Low}}$

	All Consumers	High Prosocial	Low Prosocial
Low Visibility	31.6% (76)	35.0% (40)	27.8% (36)
High Visibility	26.7% (90)	36.0% (50)	15.0% (40)

Note. Total number of observations in parentheses. $\text{Msg} = 0$ cases only.

Table 4.1 summarizes our results about the level of distrust in the Firm’s message. First, we note that in both visibility conditions and for all Consumer types (all Consumers, only high prosocial, and only low prosocial), the fraction of Consumers who believe that $\text{Clue} = Y_{\text{Low}}$ when $\text{Msg} = 0$ is significantly greater than zero (one-sided χ^2 test, $p < 0.001$). Thus, we support Hypothesis 4.1A under both visibility levels and for all prosocial types.

We also note from Table 4.1 that the percentage of Consumers who distrust the message is noticeably lower in the High Visibility condition than in the Low Visibility condition if we consider low prosocial Consumers. To formally test whether higher visibility reduces the level of distrust in the Firm’s message, we conduct a one-sided χ^2 test where the null hypothesis is that the level of distrust is the same across visibility conditions and the alternative hypothesis is that it is lower in the High Visibility condition. We note that we use one-sided statistical tests in all of our analyses due to the directionality of our hypotheses. Our results confirm that low prosocial Consumers distrust the message $\text{Msg} = 0$ significantly less under high visibility than under low visibility ($p < 0.1$). In addition, the magnitude of the effect is considerable, as it drops by almost half (27.8% in the Low Visibility condition versus 15.0% in the High Visibility condition). The differences are not statistically signif-

icant when considering all Consumers or only high prosocial Consumers ($p > 0.2$).⁴ Thus, we support Hypothesis 4.1B only for low prosocial Consumers. We summarize our findings as follows.

Result 4.1. *A significant fraction of Consumers distrust the message $Msg = 0$ sent by the Firm. This holds when considering all Consumers, only high prosocial Consumers, or only low prosocial Consumers. Regarding the effect of visibility on distrust:*

- (i) *For high prosocial Consumers, the level of distrust is independent of visibility.*
- (ii) *For low prosocial Consumers, the level of distrust is significantly lower under high visibility than under low visibility.*

4.4.2 Hypothesis 4.2: Is the Effect of Trust on Consumers' WTP Affected by Visibility?

Next, we study the effect that trust (or lack thereof) has on Consumers' WTP and how this effect is impacted by the level of visibility. Table 4.2 presents summary statistics of Consumers' WTP for all treatments when either $Msg = 0$ (Disclosure condition) or $Clue = 0$ (Control condition). We make two initial observations. First, the average WTP is significantly lower than 99 tokens in all four experimental conditions (one-sided Wilcoxon rank-sum test, $p < 0.001$). This indicates that, on average, Consumers take into account the fact that the Worker's payoff may be reduced if the product is sold. However, it is worth noting that 22% of the observations correspond to the self-interested optimal WTP of 99 or 100 tokens; and that 13.5% of Consumers state $WTP \geq 99$ in all rounds, regardless of the information they observe about Y . The

⁴Each Consumer in the Disclosure condition received the message $Msg = 0$ between 0 and 5 times, depending on the decisions made by the Firm players they were matched with. To account for the disparity in the number of observations from each Consumer, we ran the following additional analysis. For each Consumer, we defined his average distrust as the fraction of observations where he received $Msg = 0$ and believed that $Clue = Y_{Low}$. Then, we compared the distributions of average distrust between visibility conditions. This analysis confirms our results that the average distrust is significantly greater than 0 in all treatments (ones-sided Wilcoxon signed-rank test, $p < 0.05$) and that higher visibility significantly reduces the level of distrust only among low prosocial Consumers (one-sided Wilcoxon rank-sum test, $p < 0.1$). Finally, we note that only 6 out of 114 Consumers never received the message $Msg = 0$.

second observation we make is that, inline with our results in §4.4.1 and as we will discuss below, low prosocial Consumers' WTP appears to be more sensitive to our experimental manipulations than that of high prosocial Consumers.

Table 4.2: Summary Statistics of Consumers' WTP in the CTG

Visibility	Condition	All Consumers			High Prosocial			Low Prosocial		
		N	Mean	St. Dev.	N	Mean	St. Dev.	N	Mean	St. Dev.
Low	Disclosure	25	69.59	26.29	13	68.90	23.57	12	70.36	29.34
Low	Control	25	66.73	24.40	17	60.86	25.45	8	78.16	17.78
High	Disclosure	31	72.32	22.82	17	63.32	23.21	14	83.58	16.67
High	Control	27	66.14	22.31	14	67.79	25.12	13	64.24	18.77

Note. These values correspond to observations where $\text{Msg} = 0$ (Disclosure condition) or $\text{Clue} = 0$ (Control condition). N = number of Consumer players in each condition.

To formally test Hypothesis 4.2, we estimate the following random-effects regression model:

$$\begin{aligned} \text{WTP}_{jk} = & \text{Intercept} + \beta_C \cdot \text{Control} + \beta_H \cdot \text{High} + \beta_{CH} \cdot \text{Control} \times \text{High} \\ & + \beta_Y \cdot Y_{Low} + \beta_R \cdot \text{Round} + \gamma \cdot D + \delta_j + \epsilon_{jk}. \end{aligned} \quad (4.1)$$

The subscript j indexes a Consumer participant and k indexes each of the five observations per participant. The dummy variables Control and High indicate the Control condition and High Visibility conditions, respectively. Similarly, the variables Y_{Low} and Round represent the value of Y_{Low} and the round in which that value was observed by Consumer j . We control for demographic variables through the vector D , which contains information about gender and level of education. For ease of exposition, we only discuss these demographic effects in the Appendix. The term δ_j represents the individual-specific error and ϵ_{jk} is the independent error across WTP decisions. Finally, similar to our analysis in §4.4.1, Hypothesis 4.2 is only related to the case where Consumers observe $\text{Msg} = 0$ in the Disclosure condition or $\text{Clue} = 0$ in the Control condition. Therefore, we run the regression in Equation (4.1) only with these observations, which account for 53.2% of the total sample.

We use a random-effects model for the following reason. Since we run the CTG

for five rounds, we have five WTP observations from each Consumer. We accommodate such repeated measures with the random-effects approach; i.e., we include the individual-specific error δ_j . This choice of regression model is well established in the literature (e.g., List 2006, Bolton et al. 2013). We also note that since 22% of our observations correspond to a WTP equal to 99 or 100 tokens, a Tobit regression model that accounts for these corner solution outcomes could also be desirable (Wooldridge 2002 pp. 517–542). To simplify the discussion, we do not follow this regression approach, but we confirm that our results remain qualitatively the same if we consider instead a one-sided Tobit random-effects model.

The regression estimates from Equation (4.1) can be interpreted as follows. First, a positive value of the dummy variable Control indicates that Consumers are willing to pay more when Clue = 0 in the Control condition under low visibility than when Msg = 0 in the Disclosure condition under low visibility. Following the discussion in §4.3, this would capture an effect of distrust in the Firm’s message on Consumers’ WTP, i.e., $\Delta(v = 0.6) > 0$. Similarly, we also use the regression estimates to compute the average marginal effect of the variable Control under high visibility: a positive value captures an effect of distrust on WTP in the High Visibility condition; i.e., $\Delta(v = 0.9) > 0$. We note that in our analysis the magnitude of this marginal effect is equal to the sum of the regression estimates for the variables Control and Control×High. Finally, a negative value of the interaction term Control×High would imply that the difference in WTP between the Control and Disclosure conditions is lower under high visibility than under low visibility. In other words, it would indicate a lower level of distrust in the High Visibility condition than in the Low Visibility condition, or equivalently, $\Delta(v = 0.9) < \Delta(v = 0.6)$.

The columns in Table 4.3 summarize our regression results when considering observations from all Consumers, only high prosocial Consumers, and only low prosocial Consumers. First, following the discussion in the previous paragraph, we find no evidence of trust significantly affecting WTP when considering either all Consumers or only high prosocial Consumers. Regardless of the visibility condition, the differences between the Disclosure and Control conditions are not significant, as can be seen from

Table 4.3: Regression Estimates: Equation (4.1)

	All Consumers	High Prosocial	Low Prosocial
Intercept	74.33*** (7.38)	75.40*** (9.32)	88.06*** (12.48)
Control	-0.87 (6.16)	-6.49 (8.01)	13.27* (8.95)
High	3.67 (5.83)	-4.57 (7.85)	12.80** (7.78)
Control×High	-3.52 (8.50)	11.03 (11.18)	-26.79** (11.82)
Y_{Low}	0.04* (0.03)	0.06* (0.04)	0.01 (0.03)
Round	0.83* (0.63)	0.71 (0.93)	0.82 (0.76)
Observations	295	166	129
Log Likelihood	-1,273.363	-719.503	-515.136

Note. Standard errors in parentheses.

***: $p < 0.01$; **: $p < 0.05$; *: $p < 0.1$, based on one-sided t -tests.

the regression estimate of the variable Control in Table 4.3, columns 2 and 3, and from the fact that the marginal effect of the dummy variable Control evaluated under high visibility is also not significant in these cases (one-sided t test, $p > 0.1$). Similarly, we do not observe a significant effect of the interaction term Control×High among all Consumers or high prosocial Consumers. Instead, they care only about the value of Y_{Low} , stating a lower value of WTP as Y_{Low} decreases; i.e., they are less willing to earn the consumption utility of 100 tokens at the expense of the Worker's payoff. Therefore, we do not find support for Hypotheses 4.2A and 4.2B when considering all Consumers or only high prosocial Consumers.

Conversely, as can be seen in the last column of Table 4.3, low prosocial Consumers' WTP is significantly affected by our experimental conditions. First, the dummy variable Control is positive and statistically significant, which indicates an effect of distrust on WTP in the Low Visibility condition; i.e., $\Delta(v = 0.6) > 0$. This trust gap, however, is significantly lower under high visibility, as evidenced by the negative and significant coefficient estimate of the interaction term Control×High. Furthermore, in the High Visibility condition the trust gap completely disappears, as the marginal effect of Control under high visibility is equal to $\Delta(v = 0.9) = -13.52 < 0$ (one-sided t test, $p < 0.1$). Thus, among low prosocial Consumers we find support for Hypothesis 4.2A only under low visibility and we also find support for Hypoth-

esis 4.2B. Interestingly, and in contrast to high prosocial Consumers, low prosocial Consumers' WTP is not significantly affected by the value of Y_{Low} .

We summarize the effect of distrust on Consumers' WTP in the following result:

Result 4.2. *Distrust in the Firm's message $Msg = 0$ has no effect on WTP and is not affected by the level of visibility when considering all Consumer participants or only high prosocial Consumers. However, among low prosocial Consumers:*

- (i) *Under low visibility, distrust on the Firm's message significantly reduces their WTP.*
- (ii) *Under high visibility, the effect of distrust on WTP is significantly reduced and it completely disappears.*

Figure 4-3: Effect of Experimental Conditions among low prosocial Consumers when $Msg = 0$ or $Clue = 0$

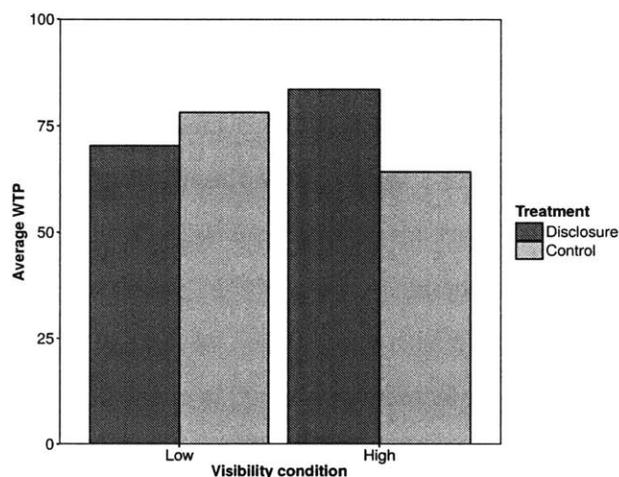


Figure 4-3 illustrates the effect that the experimental conditions of the CTG have on low prosocial Consumers. Consistent with our regression results, under low visibility, low prosocial Consumers state a significantly lower average WTP in the Disclosure condition than in the Control condition ($p < 0.01$). Conversely, under high visibility this is no longer the case. Instead, in this case low prosocial Consumers are willing to pay significantly more in the Disclosure condition than in the Control condition

($p < 0.01$).⁵ This last result also coincides with our regression analysis, where the marginal effect of the dummy variable Control in the High Visibility condition is negative and significant. This confirms that the lack of trust in the Firm’s message negatively impacts low prosocial Consumers’ WTP only under low visibility. Next, we discuss in more detail our results and what may influence the differences we observe between prosocial types.

Table 4.4 below summarizes our results about Consumers’ distrust, the effect of Consumers’ distrust on WTP, and how these depend on the level of visibility.

Table 4.4: Summary of Experimental Results

	All Consumers	High Prosocial	Low Prosocial
Distrust	Yes	Yes	Yes
– Effect of Visibility	No	No	Yes
Distrust: Impact on WTP	No	No	Yes
– Effect of Visibility	No	No	Yes

4.4.3 Discussion

We next discuss in more detail some of the key questions that arise from our experimental results.

Why does distrust not always translate into a significant effect on WTP?

This is the case among low prosocial Consumers under high visibility and among high prosocial Consumers in both visibility conditions. For the case of low prosocial Consumers, we note that the level of distrust is positive but small: in the High Visibility condition, low prosocial Consumers believe that $Clue = Y_{Low}$ when $Msg = 0$ in only 15% of our observations. This small portion of Consumers may not be enough

⁵We compare low prosocial Consumers’ WTP decisions between the Disclosure and Control conditions in two aspects: (i) the proportion of WTP values greater than or equal to 99 tokens (one-sided χ^2 test) and (ii) the distribution of WTP values less than 99 tokens (one-sided Wilcoxon rank-sum test). This approach is appropriate because a high percentage of WTP decisions by low prosocial Consumers are greater than or equal to 99 tokens (30.4% versus only 15.6% by high prosocial Consumers; see Lachenbruch 2002, Delucchi and Bostrom 2004). In the Low Visibility condition, only the distribution of WTP observations less than 99 tokens significantly differs between the two conditions. In the High Visibility condition, both (i) and (ii) are significantly different. The p values reported correspond to the significant results.

to significantly impact the average WTP. The case of high prosocial Consumers is more puzzling. As discussed in §4.4.1, in 35% (36%) of the observations in the Low Visibility (High Visibility) condition, high prosocial Consumers distrust the message $\text{Msg} = 0$ sent by the Firm in the Disclosure condition. Yet we do not observe a significant effect of distrust on WTP among high prosocial Consumers. We conjecture that these Consumers' high level of prosociality drives them to state a relatively low WTP for the product even when they believe (or know, in the Control condition) that the clue is equal to zero. This is because there always exists a positive probability that the clue is wrong and that $Y = Y_{Low} < 0$, which would hurt the Worker's payoff if the product was sold. Thus, distrust in the Firm's message does not affect high prosocial Consumers' WTP. Conversely, low prosocial Consumers care less about the Worker than their high prosocial counterparts and are therefore willing to increase their WTP when they believe that $\text{Clue} = 0$. This is because a greater WTP also increases the Consumer's own expected payoff, at the expense of potentially lowering the Worker's payoff. The high prosocial Consumers' concern for the Worker and the low prosocial Consumers' lack of concern can be further seen when we compare their WTPs. Specifically, low prosocial Consumers are willing to pay significantly more for the product than high prosocial Consumers (74.4 versus 64.8 tokens; one-sided Wilcoxon rank-sum test, $p < 0.001$).

Why do low prosocial Consumers trust the message more under high visibility? Related to the above question, we next examine why a higher level of visibility leads to greater trust among low prosocial Consumers. One aspect that may help answer this question comes from the post-experiment survey. In it, all Consumer participants who played the Disclosure condition in the CTG answered the following 5-point Likert question: "I felt that it was OK for Player A [the Firm] to send me a message different from the clue that he/she observed because Player A did not know the exact value of Y ." Among high prosocial Consumers, we find no significant differences in the proportion of respondents who agree with the statement when comparing between visibility conditions. However, low prosocial Consumers agree significantly more with it in the Low Visibility condition than in the High Visibility condition

(91.7% versus 57.1%, respectively, one-sided χ^2 test, $p < 0.1$). In other words, low prosocial Consumers believe that it is less acceptable for the Firm to send $\text{Msg} \neq \text{Clue}$ under high visibility because the Firm has greater certainty about the true value of Y . This in turn helps explain why they trust the message more in the High Visibility condition and is also inline with the main reason why we expect Hypothesis 4.1B to hold (see §4.3).

Nevertheless, it is interesting to note that the above differences in expected trustworthiness between visibility conditions do not match the actual trustworthiness behavior of the Firms in the CTG. When Firms receive a clue $\text{Clue} = Y_{Low}$, we do not observe significant differences in the proportion of Firms that send $\text{Msg} = \text{Clue}$ between the Low Visibility and High Visibility conditions. The only difference we observe is among high prosocial Firms (defined in the same way as high prosocial Consumers), who report $\text{Msg} = \text{Clue} = Y_{Low}$ in 70% of the cases under high visibility, versus in only 53.2% of the cases under low visibility. This difference, however, is not statistically significant (one-sided χ^2 test, $p = 0.11$). The same comparison among low prosocial Firms yields a difference of less than 3% and in the opposite direction.

Under high visibility, why are low prosocial Consumers willing to pay more in the Disclosure condition than in the Control condition? If the Consumer fully trusts the Firm's message, then he should be willing to pay the same for the product in both conditions, all else equal. We conjecture that low prosocial Consumers *willingness to pay more* is due to different levels of responsibility that they perceive for themselves in the Disclosure condition versus the Control condition. In the Control condition, the Consumer is the *only* player who makes an active decision to influence the final payoffs of all three players. Conversely, in the Disclosure condition, both the Firm's *and* the Consumer's decisions determine the final payoffs. Thus, the Consumer is likely to feel less responsible for the Worker's well-being in the Disclosure condition. With a decreased sense of responsibility for the final outcomes, the Consumer's low prosociality motivates him to state a higher WTP in the Disclosure condition than in the Control condition under high visibility. The phenomenon of responsibility alleviation has been documented in other games examining social preferences (e.g.,

Charness 2000, Dana et al. 2007, Kraft et al. 2017).

We indeed find evidence in the postexperiment survey that low prosocial Consumers perceive a greater level of responsibility for the Worker's payoff in the Control condition than in the Disclosure condition. Under low visibility, 88.9% of low prosocial Consumers in the Control condition felt that they were more responsible for the Worker's payoff than the Firm was (compared to the Firm being more responsible and to both players being equally responsible), versus 75% in the Disclosure condition. Consistent with our results, in the High Visibility condition this difference in perceived responsibility was even larger: 69.2% in the Control condition versus only 50% in the Disclosure condition. Even though these differences are not statistically significant, they lend support to our responsibility-alleviation conjecture.

4.5 Conclusions and Managerial Insights

Consumers increasingly want to know more about where and how the products they purchase are being made. Yet, consumers' information about SR is often limited to what companies communicate to them (e.g., through packaging), and hence there may exist a lack of consumer trust in companies' communications. In this paper, we employ an incentivized controlled laboratory experiment to investigate how supply chain visibility affects consumer trust in companies' SR communication. In particular, we study whether a company that has better visibility into its supply chain (i) enjoys greater consumer trust when claiming good SR practices and (ii) whether such an increase in trust translates into a greater willingness-to-pay for the company's products. Our work adds to the empirical SR literature, which has mainly focused on the benefits that companies can derive from disclosing SR information to consumers, by investigating the under-studied dimension of supply chain visibility.

Our results highlight the importance of companies taking into account consumer heterogeneity when designing their SR communication strategies. Specifically, consumers who are more driven by self-interest (i.e., low prosocial consumers) respond differently to information about SR practices than those who naturally care about

others' well-being (i.e., high prosocial consumers).⁶ In what follows, we conclude the paper by discussing two insights on how supply chain visibility affects low prosocial and high prosocial consumers' trust of companies' SR communications and how trust, in turn, affects their WTP.

Insight 4.1. *Low prosocial consumers trust companies' SR communications more when companies have greater visibility into their supply chains. This results in a greater willingness-to-pay for a company's products when the company claims good SR practices.*

The disclosure of SR information to consumers has been identified in practice as having a positive effect on consumer trust (MacLean and Rebernak 2007, Kang and Hustvedt 2014e.g.). Our first finding adds to this by showing that another key component of supply chain transparency – a company's visibility into the SR practices of its suppliers – can also help to improve consumer trust in SR communications. This finding thus highlights a potential market benefit for companies to improve their supply chain visibility. We find that the increase in trust among low prosocial consumers is at least partly driven by the belief that it is less acceptable for a company to lie about SR practices when it has better information about these practices in its supply chain (i.e., higher visibility).

Insight 4.2. *High prosocial consumers' trust in companies' SR communications is not affected by supply chain visibility. Furthermore, high prosocial consumers' willingness-to-pay for the product is not impacted by trust in companies' communications.*

In other words, supply chain visibility does not affect high prosocial consumers' trust in companies' SR claims and trust, in turn, does not affect their WTP. Instead, we find that high prosocial consumers' WTP is only significantly affected by the magnitude of the potential negative externality if the product is sold. Therefore,

⁶In practice, companies can develop an understanding of the heterogeneity of their consumers (e.g., their different prosociality) through psychographic segmentation, a method commonly used in marketing to segment consumers based on psychological dimensions "including activities, interests, opinions, needs, values, attitudes, and personality traits" (Wells 1975p. 197). For example, leading market research firms conduct consumer surveys to segment consumers based on their attitudes toward ethics and the environment to help companies develop better strategies for targeting specific segments (e.g., Cone Communications/Echo 2013, Mintel Group 2015).

companies targeting this group of consumers are better served by improving working conditions in the supply chain, so as to reduce the negative impact associated with poor SR practices.

Taken together, our insights thus identify when demonstrating high visibility into the SR practices in the supply chain can help improve consumers' trust. Specifically, a company with a higher level of visibility can benefit from greater trust, and consequently a greater willingness-to-pay, only if the company targets low prosocial consumers. Recall that in our experiment, we capture a situation where good SR practices correspond to the absence of a negative externality (e.g., sweatshop conditions) on workers. In this context, low prosocial consumers are more prone to believing a company's claim that the SR practices in its supply chain are good: this makes it easier for them to focus on their self-interest, which is best served by purchasing the product. Conversely, high prosocial consumers are more concerned for the workers' well-being. As a result, they are more reluctant to believe the company's claim and are also more reluctant to pay a high price for the product (even when the company has a high level of supply chain visibility).

To conclude, our research yields valuable insights into the impact of supply chain visibility on consumers' trust in companies' SR communications. Our findings highlight the importance of accounting for consumers' heterogeneity when studying consumers' responses to SR communications. Related to our work, a valuable direction for future research would be to give the Firm player in our experiment the ability to (costly) improve the level of visibility into the Worker's payoff. We expect that such an extension would yield an even greater potential for visibility to help increase Consumers' trust, as the investment could signal the Firm's intrinsic care for the Worker's well-being. Similarly, studying a context where SR practices yield a positive externality for workers (rather than negative) would be a valuable future research direction.

Chapter 5

Conclusions

Building transparent supply chains that allow companies to observe, improve, and communicate social responsibility information to consumers is an emerging business challenge for companies. Most companies do not have good visibility into the social responsibility practices in their supply chains and improving such visibility is both costly and time-consuming. This thesis, however, identifies several ways in which companies can benefit from increased supply chain transparency. Based on results from incentivized laboratory experiments and a game-theoretic model, we highlight some key managerial insights.

First, regarding consumers' valuations of social responsibility (Chapter 2), we find that consumers are willing to pay more for greater visibility into a company's social responsibility practices in its upstream supply chain. This is particularly true if visibility is currently low and if consumers exhibit a self-serving bias – i.e., if they assume that poor social responsibility practices are more likely when the information conveyed by the company is not precise. Furthermore, information regarding the social responsibility efforts made by the company can induce higher valuations from consumers when the outcomes of those efforts are very clear. This is especially the case for consumers who are mostly driven by self-interest (i.e., low prosocial consumers). These results not only provide evidence for consumers' willingness-to-pay for greater visibility, but they also help companies identify what social responsibility information best resonates with different consumer types.

Similarly, we identify important benefits from supply chain visibility for a company trying to improve the social responsibility performance of its supply chain (Chapter 3). When a company invests in a supplier's social responsibility capabilities, greater visibility into the supply chain helps the company be more efficient with its investment. In addition, if the company decides to voluntarily disclose social responsibility information to consumers, then we find that it should more aggressively invest in the supplier than if it were not to disclose. This in turn translates into better social responsibility practices by the supplier when the company discloses. Furthermore, greater visibility also motivates the company to disclose more truthful information about the social responsibility practices in its supply chain. Finally, we find that voluntary disclosure of social responsibility information is least effective in improving social responsibility practices when the supplier has only average social responsibility practices. This last result is particularly valuable for NGOs, government agencies, and other stakeholders whose objective is the improvement of social responsibility along the supply chain.

Finally, we study how visibility affects consumers' trust in social responsibility communications (Chapter 4). We find that the effect of visibility highly depends on consumers' prosociality. With greater visibility, low prosocial consumers trust companies' communications more and this results in a greater willingness-to-pay for the company's product. Conversely, high prosocial consumers' trust in companies' social responsibility communications is not affected by visibility. This result is partly explained by the fact that in our experiment, good social responsibility practices capture the lack of a negative externality to workers in the supply chain. Thus, low prosocial consumers are more prone to believing the company's claims, as this makes it easier for them to focus on their self-interest and purchase the company's product.

This thesis contributes to the growing literature on supply chain transparency and social responsibility. We address not only the disclosure of social responsibility information to consumers, but also an equally-important yet understudied dimension of transparency; namely, the level of visibility that a company has into the social responsibility practices of its suppliers. The field of operations management is uniquely

positioned to address the challenges that arise from managing (and improving) social responsibility practices in a complex supply chain environment. Similarly, behavioral sciences and experimental methods are essential to help us understand how consumers and other stakeholders may respond to social responsibility information. We believe that the intersection of these fields of study constitutes a key area for future research.

Appendix A

Additional Material Chapter 2

A.1 Regression Estimates and Marginal Effects of Main Results

Table A.1 presents the regression results for Equation (2.1) using data from all Consumer participants (column 2) and by Consumers' prosocial type (columns 3 and 4). Tables A.2 and A.3 summarize the marginal effects when analyzing the effects of visibility and reciprocity on WTP by Consumers' prosocial type.

A.2 Robustness Experiments

A.2.1 Implementing the Strategy Method with a Random-Order Presentation

To address potential priming effects that the strategy method may induce when all possible values of e are presented simultaneously in a table, we test a variation of our original design by employing a random-order presentation of the strategy method; i.e., we present to the Consumer each possible value of e one at a time and on separate screens in a random order. We conduct sessions of this variation for the Decision condition under both low and high visibility with 60 participants in total.

We continue to observe in these new sessions that the Consumers' WTP decisions are increasing in e (the correlations between WTP and e are 0.46 and 0.48; t tests, $p < 0.001$). This eliminates the concern of priming effects influencing our results. The robustness analyses conducted in §A.2.2 and §A.2.3 also employ the strategy method with a random-order presentation.

Table A.1: Regression Results: Equation (2.1)

Variable	All Consumers	High prosocial	Low prosocial
Intercept	-0.13 (8.15)	16.66 (8.02)**	-7.84 (11.19)
V_M	-0.72 (12.04)	-4.81 (10.82)	-2.08 (17.76)
V_H	-1.68 (11.72)	10.94 (11.90)	-6.30 (15.87)
D	-14.15 (12.58)	-5.80 (11.02)	-27.90 (19.30)*
$D \cdot V_M$	6.88 (18.17)	19.01 (15.67)	5.15 (27.69)
$D \cdot V_H$	21.31 (18.21)	10.39 (16.85)	30.97 (26.61)
e_k	0.34 (0.03)***	0.47 (0.05)***	0.26 (0.04)***
$V_M \cdot e_k$	0.06 (0.05)	0.09 (0.07)*	-0.06 (0.07)
$V_H \cdot e_k$	-0.06 (0.05)*	-0.09 (0.08)	-0.04 (0.06)
$D \cdot e_k$	0.05 (0.05)	-0.02 (0.07)	0.01 (0.08)
$D \cdot V_M \cdot e_k$	0.04 (0.08)	-0.09 (0.10)	0.29 (0.11)***
$D \cdot V_H \cdot e_k$	0.10 (0.08)	0.06 (0.11)	0.18 (0.11)**
σ_δ	48.10 (3.12)***	24.56 (2.32)***	54.73 (4.91)***
σ_ϵ	19.43 (0.48)***	18.64 (0.69)***	19.27 (0.64)***

Table A.2: Marginal Effects of V_H , V_M , and Their Differences, Δ_V by Prosocial Type

Effort	High prosocial			Low prosocial		
	V_H	V_M	Δ_V	V_H	V_M	Δ_V
0	16.04 (9.20)**	10.19 (8.23)	5.85 (10.30)	8.32 (7.26)	0.84 (5.80)	7.48 (7.16)
20	17.32 (9.59)**	11.55 (8.76)*	5.78 (10.58)	10.46 (7.99)*	2.44 (6.60)	8.01 (8.11)
40	18.12 (9.88)**	12.62 (9.17)*	5.50 (10.75)	12.85 (8.73)**	4.49 (7.46)	8.36 (9.09)
60	18.45 (10.12)**	13.39 (9.50)*	5.06 (10.88)	15.47 (9.49)*	6.99 (8.37)	8.48 (10.08)
80	18.36 (10.36)**	13.85 (9.81)*	4.51 (11.00)	18.28 (10.27)**	9.92 (9.32)	8.36 (11.07)
100	17.89 (10.58)**	14.01 (10.11)*	3.88 (11.10)	21.23 (11.05)**	13.26 (10.29)*	7.97 (12.01)
120	17.05 (10.75)*	13.86 (10.36)*	3.19 (11.13)	24.27 (11.83)**	16.92 (11.27)*	7.34 (12.90)

Table A.3: Marginal Effects of $D \cdot e_k$ by Prosocial Type

Effort	High prosocial			Low prosocial		
	High	Medium	Low	High	Medium	Low
0	0.05 (0.06)	-0.01 (0.06)	-0.05 (0.06)	0.09 (0.04)**	0.06 (0.05)	-0.04 (0.03)
20	0.05 (0.07)	-0.04 (0.07)	-0.05 (0.06)	0.10 (0.05)**	0.08 (0.05)*	-0.04 (0.04)
40	0.05 (0.08)	-0.07 (0.07)	-0.04 (0.07)	0.12 (0.05)***	0.11 (0.05)**	-0.04 (0.04)
60	0.04 (0.08)	-0.09 (0.07)*	-0.04 (0.07)	0.13 (0.05)***	0.13 (0.06)***	-0.04 (0.04)
80	0.04 (0.08)	-0.10 (0.07)*	-0.03 (0.07)	0.14 (0.05)***	0.16 (0.06)***	-0.04 (0.05)
100	0.03 (0.07)	-0.10 (0.07)*	-0.02 (0.07)	0.15 (0.05)***	0.18 (0.06)***	-0.04 (0.05)
120	0.02 (0.07)	-0.10 (0.06)*	-0.01 (0.06)	0.16 (0.05)***	0.20 (0.06)***	-0.03 (0.05)

A.2.2 Minimizing Efficiency Considerations

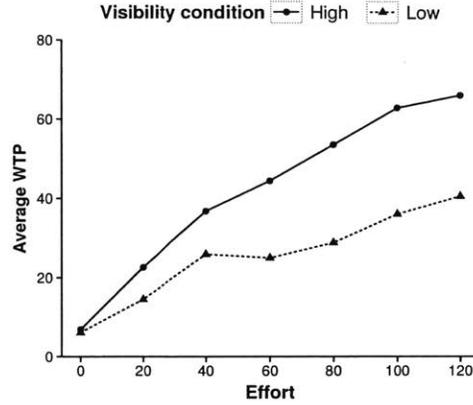
The average WTP being significantly higher than zero even when $e = 0$ may be due to Consumers' efficiency considerations; i.e., they prefer the product being sold to increase the total payoff for the group. The presence of efficiency considerations is well documented in the social preferences literature (e.g., Charness and Rabin 2002, Engelmann and Strobel 2004). To examine whether efficiency considerations influence our results, we perform two robustness checks. We first test a design where the Firm player starts with $160 + 120 = 280$ tokens; i.e., the 120 tokens are no longer provisional. All other design aspects remain the same as in the original experiment. We conducted sessions of this variation for both the Decision and Random conditions under low and high visibility with 117 participants in total. We confirm that our original results regarding the effects of visibility and reciprocity continue to hold (see Figure A-1). In particular, Consumers' WTP in the Decision condition is significantly higher under high visibility than under low visibility. Also, reciprocity exists and has a significant, positive effect on WTP under high visibility (one-sided Wilcoxon tests, $p < 0.05$ for $e \in \{20, 40, 60, 80, 100\}$; $p < 0.1$ for $e = 120$); we find no evidence of reciprocity under low visibility. Tables A.4 and A.5 summarize our results.

As a second analysis, we exclude from the original data those Consumers who state a positive WTP when $e = 0$ and reestimate our regression. The remaining data consists of decisions from 98 Consumers. Table A.6 summarizes the marginal effects. We again confirm the robustness of our original results: (i) Consumers state significantly higher WTP under high or medium visibility than under low visibility; (ii) reciprocity exists under high visibility, it diminishes under medium visibility, and there is no evidence of reciprocity under low visibility.

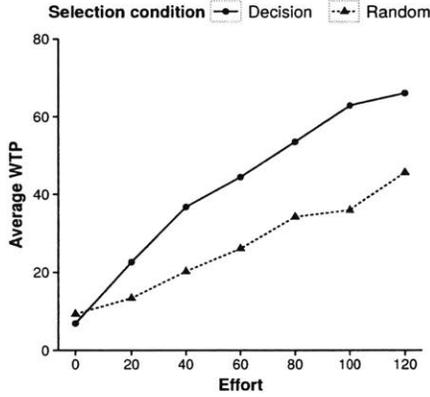
A.2.3 Varying the Wealth Status of the Disadvantaged Party

To examine the robustness of our results with respect to the disadvantaged party's wealth status, we study Consumers' WTP in the CPG when the Worker is less disadvantaged by increasing the Worker's initial endowment from 20 to 50 or 80 tokens

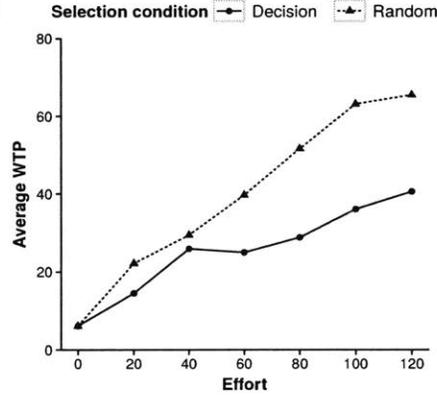
Figure A-1: Summary of Results in the Variation with No Provisional Endowment



(a) Effect of Visibility (Decision Condition)



(b) Indirect Reciprocity (High Visibility)



(c) Indirect Reciprocity (Low Visibility)

(Worker-50 or Worker-80 condition). Everything else remains the same. We conducted sessions of these variations for the Decision condition under high and low visibility. In total, we had 66 Consumers in the Worker-50 condition, equally split between high and low visibility; 57 Consumers in the Worker-80 condition with 29 (28) under high (low) visibility.

Table A.4: Regression Estimates: No Provisional Endowment

Variable	Regression estimates	Variable	Regression estimates
Intercept	-3.27 (7.61)	$V_H \cdot e_k$	-0.24 (0.06)***
V_H	-1.62 (10.81)	$D \cdot e_k$	-0.27 (0.06)***
D	0.67 (10.38)	$D \cdot V_H \cdot e_k$	0.42 (0.09)***
$D \cdot V_H$	12.94 (14.64)	σ_δ	28.52 (2.79)***
e_k	0.62 (0.05)***	σ_ϵ	18.35 (0.68)***

Table A.5: Marginal Effects of V_H and $D \cdot e_k$ (by Visibility Condition): No Provisional Endowment

Effort	Marginal effects of V_H	High	Low
0	6.06 (5.29)	0.15 (0.05)***	-0.12 (0.06)**
20	9.53 (6.03)*	0.18 (0.05)***	-0.18 (0.06)***
40	13.55 (6.68)**	0.19 (0.05)***	-0.23 (0.06)***
60	17.90 (7.25)***	0.20 (0.06)***	-0.27 (0.06)***
80	22.32 (7.75)***	0.19 (0.06)***	-0.29 (0.06)***
100	26.58 (8.21)***	0.18 (0.06)***	-0.28 (0.06)***
120	30.44 (8.60)***	0.15 (0.06)***	-0.26 (0.06)***

Table A.6: Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Excluding Consumers Stating $WTP > 0$ for $e = 0$

Effort	V_H	V_M	Δv	High	Medium	Low
0	4.74 (4.15)	9.05 (4.41)**	-4.31 (5.67)	0.06 (0.07)	0.03 (0.07)	-0.11 (0.04)***
20	7.05 (5.51)	12.54 (5.49)**	-5.49 (7.31)	0.09 (0.08)	0.05 (0.07)	-0.13 (0.05)***
40	10.05 (7.07)*	16.76 (6.65)***	-6.71 (9.12)	0.13 (0.09)*	0.07 (0.08)	-0.14 (0.06)***
60	13.79 (8.79)*	21.67 (7.87)***	-7.88 (11.02)	0.17 (0.10)**	0.09 (0.08)	-0.15 (0.06)**
80	18.20 (10.62)**	27.09 (9.11)***	-8.89 (12.91)	0.20 (0.10)**	0.10 (0.08)*	-0.15 (0.07)**
100	23.16 (12.46)**	32.78 (10.36)***	-9.63 (14.65)	0.24 (0.10)**	0.11 (0.07)*	-0.14 (0.08)**
120	28.41 (14.25)**	38.43 (11.58)***	-10.02 (16.12)	0.26 (0.10)***	0.11 (0.07)*	-0.13 (0.09)*

Table A.7: Marginal Effects of V_H : Worker-50 & Worker-80 by Consumer's Self-Serving Bias

Effort	50	80	50	80
0	10.86 (5.68)**	6.93 (5.16)*	-4.12 (6.14)	-7.82 (6.91)
20	14.49 (6.50)**	8.64 (6.22)*	-4.13 (6.88)	-9.83 (7.91)
40	18.56 (7.31)***	10.33 (7.30)*	-3.85 (7.52)	-11.75 (8.76)*
60	22.89 (8.12)***	11.88 (8.36)*	-3.31 (8.05)	-13.40 (9.44)*
80	27.27 (8.93)***	13.17 (9.38)*	-2.56 (8.48)	-14.63 (9.94)*
100	31.44 (9.74)***	14.08 (10.33)*	-1.67 (8.79)	-15.33 (10.23)*
120	35.15 (10.54)***	14.53 (11.15)*	-0.73 (8.94)	-15.44 (10.29)*

Table A.8: Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Equation (2.1) When Low Prosocial Consumers Are Those Who Transfer Exactly 0 in the Dictator Game

Effort	V_H	V_M	Δv	High	Medium	Low
0	1.02 (3.58)	1.40 (3.24)	-0.37 (3.74)	0.02 (0.05)	0.02 (0.05)	-0.12 (0.03)***
20	2.24 (4.35)	2.92 (3.90)	-0.68 (4.99)	0.03 (0.06)	0.03 (0.06)	-0.12 (0.03)***
40	3.87 (5.35)	5.00 (4.78)	-1.12 (6.51)	0.05 (0.07)	0.06 (0.07)	-0.13 (0.03)***
60	5.99 (6.59)	7.73 (5.87)*	-1.73 (8.29)	0.08 (0.08)	0.09 (0.08)	-0.14 (0.03)***
80	8.68 (8.05)	11.21 (7.15)*	-2.53 (10.33)	0.11 (0.09)	0.12 (0.09)*	-0.15 (0.03)***
100	12.00 (9.72)	15.52 (8.58)**	-3.52 (12.61)	0.14 (0.10)*	0.16 (0.10)**	-0.16 (0.04)***
120	15.99 (11.57)*	20.68 (10.12)**	-4.69 (15.08)	0.17 (0.11)*	0.19 (0.10)**	-0.16 (0.04)***

We follow a similar approach as in §2.4.1 to examine Consumers' WTP decisions in the Worker-50 and Worker-80 conditions, for Consumers who demonstrate a self-serving bias and those who do not. For the Worker-50 (Worker-80) condition, 13 out of

33 (15 out of 28) Consumers in the Low Visibility condition demonstrate a self-serving bias. We estimate the following random-effects two-sided Tobit model with data from these two groups of Consumers: $WTP_{jk}^* = \text{Intercept} + \alpha_{V_H} \cdot V_H + \alpha_{W_{80}} \cdot W_{80} + \alpha_{W_{80}V_H} \cdot W_{80} \cdot V_H + \beta \cdot e_k + \beta_{V_H} \cdot V_H \cdot e_k + \beta_{W_{80}} \cdot W_{80} \cdot e_k + \beta_{W_{80}V_H} \cdot W_{80} \cdot V_H \cdot e_k + \text{Female} + \delta_j + \epsilon_{jk}$. The variables WTP_{jk}^* , V_H , e_k , δ_j , and ϵ_{jk} are defined the same as in Equation (2.1). The dummy variable W_{80} indicates the Worker-80 conditions, and the variable Female is a dummy variable that takes the value of 1 if participant j is a female. We introduce this dummy variable to control for differences in gender composition across the four treatments. Table A.7 presents our results.

As before, a positive marginal effect of V_H indicates that WTP is higher under high visibility than under low visibility at the given effort level and Worker Endowment condition. Table A.7 summarizes the marginal effects of V_H for Consumers who demonstrate a self-serving bias (columns 2–3) and those who do not (columns 4–5). We observe that among Consumers who demonstrate a self-serving bias, their WTP is significantly higher under high visibility than under low visibility at almost all effort levels for both the Worker-50 and Worker-80 conditions (except when $e_k = 120$ in the Worker-80 condition), although the increases in WTP are smaller and less significant in the Worker-80 condition. In stark contrast, we do not observe a higher WTP under high visibility for Consumers who do not exhibit a self-serving bias in either Worker Endowment condition. In fact, we observe a higher WTP under low visibility at $e_k \geq 40$ for the Worker-80 condition. We conjecture that this unusual pattern may be due to the fact that among Consumers who do not exhibit a self-serving bias, those in the Worker-80 condition appear to be more prosocial than those in the Worker-20 or Worker-50 condition. Therefore, the higher prosociality of non-biased Consumers in the Worker-80 condition may lead to a significantly higher WTP under low visibility for some effort levels, since non-biased Consumers focus more on the average/maximum (as opposed to the minimum) payment to the Worker.

A.2.4 Analysis on Alternative Classifications of Prosocial Types

We examine two alternative classifications of prosocial types. The first alternative is to define low (high) prosocial Consumers as those whose dictator decisions are less than or equal to (greater than or equal to) the lower (upper) third of the distribution of dictator decisions from all Consumers. Under this definition, low (high) prosocial Consumers transfer exactly 0 (40 or more) tokens to the recipients in the dictator game. Because high prosocial Consumers are defined the same as in §2.4.3, we reestimate Equation (2.1) with the data from the above new group of low prosocial Consumers only. Table A.8 summarizes the marginal effects. We observe very similar results as in §2.4.3. Specifically, low prosocial Consumers state a significantly higher WTP under high and medium visibility than under low visibility at high effort levels. In addition, reciprocity exists under high and medium visibility, whereas it disappears under low visibility. Finally, we again observe a significantly lower WTP in the Decision (versus Random) condition under low visibility (one-sided Wilcoxon rank-sum test, $p < 0.1$ for $e = 0$; $p < 0.01$ for all $e > 0$).

The second alternative classification that we examine defines Consumers' prosocial types based on their responses in the attitudinal survey. We follow an approach similar to psychographic segmentation to classify Consumers (e.g., Hofstede et al. 1999, Straughan and Roberts 1999). First, we employ a factor analysis to identify the relevant attitudinal questions that measure the participants' prosociality. We use parallel analysis (Fabrigar et al. 1999, Hayton et al. 2004) to identify five factors, with the first (last) factor explaining the most (least) variance in the data. We then regress the participants' dictator decisions on their scores for these five factors. Only the first factor is significantly correlated with the dictator decisions ($p = 0.001$). Thus, we use the first factor alone to classify prosocial types. This procedure allows us to identify the set of relevant attitudinal questions whose responses are internally consistent (the Cronbach's α equals 0.9) and reflect a participant's prosociality. The set of relevant questions identified are available upon request.

Second, because the relevant questions contain both binary-choice (Yes or No)

and 5-point Likert-scale questions, we code the binary choices into 0/1 scores (1 meaning Yes) and the 5-point scale responses into scores ranging from 0 to 1 with an increment of 0.25. We then take the average score among all relevant questions for a participant as the participant’s final score. Third, we use the median value of the distribution of final scores among all Consumers as the threshold to classify prosocial types: Consumers with a final score higher than or equal to the median are classified as high prosocial (101 out of 196 or 52%), whereas Consumers with a final score lower than the median are classified as low prosocial (95 out of 196 or 48%). The above procedure follows from methods commonly used in attitudinal studies in the social science disciplines (e.g., Henry and Sears 2002, Zemack-Rugar et al. 2012).

Based on this alternative classification, we reestimate Equation (2.1) with the data from the two groups separately. Table A.9 summarizes the relevant marginal effects. We observe that our insights regarding high and low prosocial Consumers discussed in §2.4.3 are in general robust to this alternative classification. First, both high and low prosocial Consumers state higher WTP under high visibility than under low visibility. Second, the presence and the effect of reciprocity on both types of Consumers’ WTP remain very similar to the observations summarized in Result 2.3. The only exception is that high prosocial Consumers exhibit some reciprocity but with a small magnitude that has no significant effect on their WTP. Specifically, the marginal effects of $D \cdot e_k$ are smaller than 0.13 at all effort levels for high prosocial Consumers, compared to ranging from 0.10 to 0.29 for low prosocial Consumers.

A.3 Supplementary Analyses

When analyzing the effect of reciprocity on low prosocial Consumers’ WTP decisions, we compare decisions between the Decision and Random conditions for a given effort in two aspects: (i) the proportion of zero values and (ii) the distribution of non-zero values. This approach is appropriate because a high percentage of low prosocial Consumers state a zero WTP (38.8% versus only 9% by high prosocial Consumers; see Lachenbruch 2002, Delucchi and Bostrom 2004). Under medium and high visibility,

the proportion of zero values is not significantly different, but the distribution of positive WTP decisions significantly differs between the Decision and Random conditions. Under low visibility, the proportion of zero values is significantly larger for the Decision condition, but the distribution of positive WTP decisions is not significantly different. The p values reported in §2.4.3 correspond to the significant results.

Table A.9: Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Equation (2.1) When Prosocial Types Are Classified Based on the Attitudinal Survey

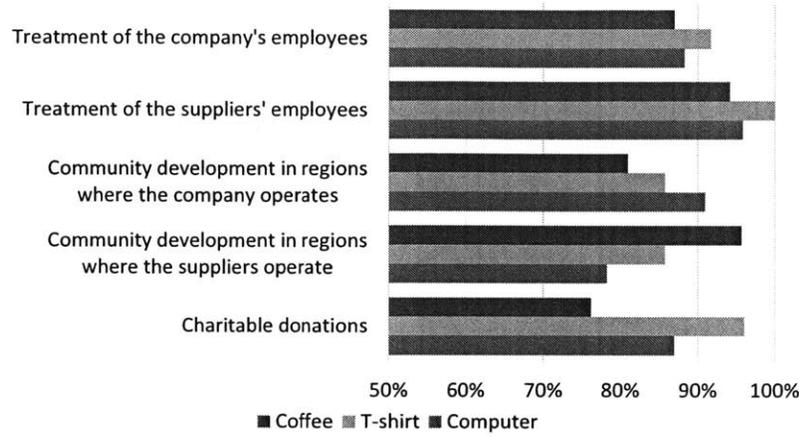
Effort	High prosocial			Low prosocial		
	V_H	V_M	Δ_V	V_H	V_M	Δ_V
0	10.20 (8.69)	8.86 (8.17)	1.34 (10.07)	7.00 (10.39)	-6.05 (8.18)	13.04 (8.65)*
20	11.91 (9.56)	10.03 (9.01)	1.88 (10.97)	8.98 (11.21)	-4.08 (9.23)	13.06 (9.81)*
40	13.60 (10.37)*	11.14 (9.82)	2.47 (11.79)	11.18 (12.04)	-1.16 (10.39)	12.34 (11.03)
60	15.20 (11.12)*	12.12 (10.58)	3.08 (12.51)	13.55 (12.86)	2.75 (11.63)	10.80 (12.28)
80	16.65 (11.80)*	12.95 (11.30)	3.70 (13.12)	16.07 (13.68)	7.61 (12.89)	8.46 (13.49)
100	17.87 (12.40)*	13.59 (11.95)	4.28 (13.60)	18.68 (14.49)*	13.27 (14.10)	5.40 (14.58)
120	18.80 (12.90)*	14.00 (12.52)	4.80 (13.94)	21.33 (15.28)*	19.53 (15.18)*	1.80 (15.48)
Effort	High	Medium	Low	High	Medium	Low
0	0.06 (0.06)	-0.01 (0.06)	0.03 (0.05)	0.10 (0.05)**	0.04 (0.07)	-0.08 (0.05)**
20	0.07 (0.06)	-0.02 (0.06)	0.05 (0.05)	0.12 (0.05)***	0.08 (0.07)	-0.10 (0.05)**
40	0.08 (0.06)*	-0.04 (0.06)	0.07 (0.05)	0.13 (0.05)***	0.13 (0.08)**	-0.11 (0.05)**
60	0.09 (0.06)*	-0.05 (0.06)	0.09 (0.06)*	0.14 (0.05)***	0.18 (0.08)**	-0.13 (0.06)**
80	0.09 (0.06)*	-0.06 (0.06)	0.10 (0.06)**	0.15 (0.05)***	0.23 (0.08)***	-0.14 (0.06)***
100	0.09 (0.06)*	-0.07 (0.06)	0.12 (0.06)**	0.15 (0.05)***	0.26 (0.08)***	-0.15 (0.06)***
120	0.08 (0.06)*	-0.07 (0.06)	0.13 (0.06)**	0.16 (0.05)***	0.29 (0.07)***	-0.15 (0.06)***

Figure A-2 summarizes Consumers' preferences in the product choice study. For Figure A-2b, if a Consumer always prefers the product with vague information or he/she prefers the product with precise information only when the prices are the same, then we treat the premium he/she is willing to pay as zero.

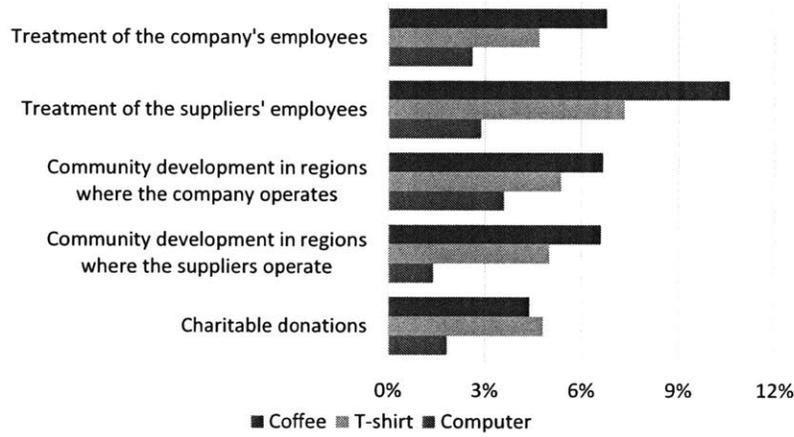
A.3.1 The Effects of Demographics on Consumers' WTP

We analyze whether demographic factors impact Consumers' WTP and our experimental findings. Specifically, we examine the effects of their gender, age, and income level on our regression results. Since age and income level are highly correlated in our sample, we perform two separate regression analyses: One controls for gender and age and their interaction term, and the other controls for gender and income and their interaction term. Tables A.10 and A.11 summarize the regression estimates and marginal effects for both analyses. We confirm that our main results regarding

Figure A-2: Consumers' Preferences in the Product Choice Study



(a) Fraction of Consumers Preferring Precise Information with No Premium



(b) Average Premiums (%) Consumers Are Willing to Pay for Precise Information

the effects of visibility and reciprocity on Consumers' WTP remain unchanged. In addition, female Consumers state higher WTP in general, whereas neither age nor income has any significant effect on Consumers' WTP.

A.3.2 Analysis on Order Effects in the CPG

We test for the potential existence of order effects in the CPG by comparing the distributions of Consumers' WTP between Round 1 and Round 2. For each combination of (i) Selection condition, (ii) Visibility condition, and (iii) effort level, we conduct

two-sided Wilcoxon rank-sum tests. The only significant difference we observe is that the Consumers' WTP is lower in Round 2 for High Visibility and low effort levels ($e \leq 60$ in the Decision condition and $e \leq 20$ in the Random condition, $p < 0.1$). For all other experimental conditions and effort levels, the Consumers' WTP decisions are not significantly different between the two rounds. We also conduct a regression analysis where we control for round as an independent variable and verify that our results on the effects of visibility and indirect reciprocity on WTP remain unchanged. Tables A.12 and A.13 summarize the regression estimates and relevant marginal effects from this analysis.

Table A.10: Regression Estimates: Controls for Demographics

Variable	Gender and Age	Gender and Income
Intercept	-18.18 (22.77)	-24.27 (13.75)**
Gender (Female = 1)	38.69 (32.44)	35.70 (14.83)***
Age/Income	-0.11 (0.95)	0.56 (1.40)
Gender·Age/Income	-0.38 (1.45)	-0.67 (1.75)
V_M	-2.29 (11.64)	-2.36 (11.61)
V_H	4.00 (11.36)	3.29 (11.40)
D	-11.04 (12.31)	-11.55 (12.29)
$D \cdot V_M$	9.44 (17.71)	9.42 (17.74)
$D \cdot V_H$	10.18 (17.74)	10.32 (17.73)
e_k	0.34 (0.03)***	0.34 (0.03)***
$V_M \cdot e_k$	0.05 (0.05)	0.05 (0.05)
$V_H \cdot e_k$	-0.06 (0.05)*	-0.06 (0.05)*
$D \cdot e_k$	0.03 (0.05)	0.03 (0.05)
$D \cdot V_M \cdot e_k$	0.08 (0.08)	0.08 (0.08)
$D \cdot V_H \cdot e_k$	0.11 (0.08)*	0.11 (0.08)*
σ_δ	45.97 (3.01)***	46.00 (3.01)***
σ_ϵ	19.35 (0.48)***	19.35 (0.48)***

Table A.11: Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Controls for Demographics

Effort	Gender and Age			Gender and Income		
	V_H	V_M	Δ_V	V_H	V_M	Δ_V
0	6.39 (6.18)	3.05 (5.74)	3.34 (6.51)	6.10 (6.15)	3.00 (5.70)	3.10 (6.50)
20	7.69 (6.78)	4.78 (6.43)	2.91 (7.24)	7.37 (6.77)	4.73 (6.40)	2.64 (7.24)
40	9.06 (7.38)	6.82 (7.11)	2.24 (7.93)	8.71 (7.37)	6.76 (7.07)	1.95 (7.94)
60	10.47 (7.97)*	9.11 (7.77)	1.36 (8.58)	10.09 (7.96)	9.04 (7.73)	1.04 (8.59)
80	11.87 (8.53)*	11.57 (8.39)*	0.30 (9.16)	11.47 (8.53)*	11.50 (8.34)*	-0.03 (9.18)
100	13.22 (9.07)*	14.11 (8.96)*	-0.89 (9.66)	12.80 (9.08)*	14.04 (8.91)*	-1.24 (9.68)
120	14.47 (9.58)*	16.61 (9.46)**	-2.14 (10.05)	14.04 (9.60)*	16.53 (9.41)**	-2.50 (10.09)
Effort	High	Medium	Low	High	Medium	Low
0	0.07 (0.04)**	0.04 (0.04)	-0.01 (0.03)	0.07 (0.04)**	0.04 (0.04)	-0.02 (0.03)
20	0.08 (0.04)**	0.06 (0.04)*	-0.01 (0.04)	0.08 (0.04)**	0.06 (0.04)	-0.01 (0.04)
40	0.10 (0.04)***	0.07 (0.05)*	-0.01 (0.04)	0.09 (0.04)***	0.07 (0.05)*	-0.01 (0.04)
60	0.11 (0.04)***	0.08 (0.05)**	-0.00 (0.04)	0.10 (0.04)***	0.08 (0.05)**	-0.00 (0.04)
80	0.11 (0.04)***	0.09 (0.05)**	0.00 (0.04)	0.11 (0.04)***	0.09 (0.05)**	0.00 (0.04)
100	0.12 (0.04)***	0.09 (0.05)**	0.01 (0.04)	0.12 (0.04)***	0.09 (0.05)**	0.01 (0.04)
120	0.12 (0.04)***	0.09 (0.04)**	0.01 (0.04)	0.12 (0.04)***	0.09 (0.04)**	0.01 (0.04)

Table A.12: Regression Estimates: Including Round as an Independent Variable

Variable	Regression estimates
Intercept	7.96 (8.66)
Round = 2	-17.19 (7.04)***
V_M	-0.16 (11.87)
V_H	-2.19 (11.56)
D	-14.40 (12.40)
$D \cdot V_M$	6.28 (17.91)
$D \cdot V_H$	23.90 (17.99)*
e_k	0.34 (0.03)***
$V_M \cdot e_k$	0.06 (0.05)
$V_H \cdot e_k$	-0.06 (0.05)*
$D \cdot e_k$	0.05 (0.05)
$D \cdot V_M \cdot e_k$	0.04 (0.08)
$D \cdot V_H \cdot e_k$	0.10 (0.08)
σ_δ	47.30 (3.07)***
σ_ϵ	19.43 (0.48)***

Table A.13: Marginal Effects of Visibility Variables and $D \cdot e_k$ (by Visibility Condition): Including Round as an Independent Variable

Effort	V_H	V_M	ΔV	High	Medium	Low
0	10.03 (6.49)*	2.48 (5.45)	7.54 (6.74)	0.10 (0.04)***	0.01 (0.04)	-0.02 (0.03)
20	11.61 (7.10)*	3.77 (6.16)	7.84 (7.44)	0.11 (0.04)***	0.03 (0.04)	-0.01 (0.04)
40	13.21 (7.69)**	5.29 (6.86)	7.92 (8.11)	0.12 (0.04)***	0.04 (0.05)	-0.01 (0.04)
60	14.78 (8.26)**	7.01 (7.55)	7.77 (8.74)	0.13 (0.04)***	0.05 (0.05)	0.00 (0.04)
80	16.25 (8.80)**	8.86 (8.22)	7.39 (9.30)	0.13 (0.04)***	0.06 (0.05)	0.01 (0.04)
100	17.58 (9.31)**	10.79 (8.84)	6.80 (9.79)	0.13 (0.04)***	0.06 (0.05)*	0.01 (0.04)
120	18.72 (9.76)**	12.70 (9.41)*	6.02 (10.18)	0.13 (0.04)***	0.07 (0.04)*	0.02 (0.04)

A.3.3 Visibility Results for the Random Condition

The seemingly higher WTP in the Medium Visibility, Random condition than in the High Visibility, Random condition and the Low Visibility, Random condition observed in Table 2.1 is in fact not statistically significant. This nonsignificant result is confirmed by both the Wilcoxon rank-sum test at each effort level ($p > 0.1$) and the marginal effects of V_H , V_M , and ΔV in our regression for the Random condition (i.e., replicating Table 2.2 for the Random condition: taking $D = 0$ in Equation (2.1)). Table A.14 summarizes these marginal effects evaluated at each effort level.

A.4 The Role of Self-Serving Bias in Affecting Consumers' WTP

We estimate the following random-effects two-sided Tobit model with the data from Consumers who exhibit a self-serving bias and those who do not separately: $WTP_{jk}^* =$

Intercept + $\alpha_{V_H} \cdot V_H + \beta \cdot e_k + \beta_{V_H} \cdot V_H \cdot e_k + \delta_j + \epsilon_{jk}$. The variables are defined the same as in Equation (2.1). In this analysis, we only use the data from the Decision condition, and the data under low visibility are compared to the aggregate data under high visibility. Consumers in the High Visibility treatments are not separated into two groups because we do not ask them questions related to a self-serving bias. Table A.15 summarizes the regression estimates and relevant marginal effects.

Table A.14: Marginal Effects of V_H , V_M , and Δ_V for the Random Condition

Effort	V_H	V_M	Δ_V
0	-0.81 (5.65)	-0.35 (5.85)	-0.46 (5.88)
20	-1.54 (6.05)	0.22 (6.36)	-1.76 (6.35)
40	-2.38 (6.46)	0.89 (6.86)	-3.27 (6.82)
60	-3.3 (6.87)	1.67 (7.35)	-4.97 (7.28)
80	-4.3 (7.29)	2.53 (7.82)	-6.83 (7.74)
100	-5.36 (7.71)	3.44 (8.27)	-8.8 (8.19)
120	-6.45 (8.14)	4.38 (8.68)	-10.84 (8.62)

Note. Values shown are the marginal effects evaluated at the Random condition; i.e., at $D = 0$ in Equation (2.1).

Table A.15: Regression Estimates and Marginal Effects for Consumers Who Exhibit a Self-Serving Bias and Those Who Do Not

Variable	Self-serving bias	No self-serving bias
	Regression estimates	
Intercept	-36.48 (16.37)**	7.97 (15.08)
V_H	41.24 (19.72)**	-3.12 (18.69)
e_k	0.35 (0.06)***	0.40 (0.05)***
$V_H \cdot e_k$	0.08 (0.07)	0.02 (0.07)
σ_δ	54.91 (8.24)***	53.25 (7.49)***
σ_ϵ	18.10 (1.02)***	18.64 (1.01)***
Effort	Marginal effects of V_H	
0	15.81 (7.15)**	-1.63 (9.83)
20	18.38 (7.81)***	-1.52 (10.56)
40	21.06 (8.48)***	-1.36 (11.24)
60	23.78 (9.15)***	-1.15 (11.86)
80	26.50 (9.83)***	-0.89 (12.40)
100	29.12 (10.50)***	-0.60 (12.85)
120	31.59 (11.17)***	-0.29 (13.19)

Note. Values in parentheses are the standard errors.

***: $p < 0.01$; **: $p < 0.05$; p values are derived from one-sided t tests.

Appendix B

Additional Material Chapter 3

B.1 Additional Analytical Results

Theorem 5. For given β_D and s_D , the supplier's optimal SR decision is given by $s^*(s_0, \beta_D) = \max(s_0, s_{uD}^*(\beta_D))$. Furthermore, $\frac{\partial s^*(s_0, \beta_D)}{\partial \beta_D} \geq 0$, $\frac{\partial s^*(s_0, \beta_D)}{\partial q} \geq 0$, and $s_{uD}^*(\beta_D) \leq s_D, \forall \beta_D \in [0, 1]$.

Theorem 6. For a given level of visibility v , there exist $\beta_{LD}(v)$ and $\beta_{HD}(v)$ such that $0 \leq \beta_{LD}(v) \leq \beta_{HD}(v) < 1$. The manufacturer's optimal investment, $\beta_D^*(v, \tilde{s})$, is defined as follows:

a) If $\beta_{LD}(v) \geq \hat{\beta}_D$, then $\beta_D^*(v, \tilde{s}) = \hat{\beta}_D$ for all \tilde{s} .

b) If $\beta_{LD}(v) < \hat{\beta}_D$, then there exists a threshold $\tau_D(v) \in (m, \hat{s}_D)$ such that:

i) If $\tilde{s} \leq \tau_D(v)$, then $\beta_D^*(v, \tilde{s}) = \bar{\beta}_{HD}(v) \equiv \min\{\beta_{HD}(v), \hat{\beta}_D\}$;

ii) If $\tilde{s} > \tau_D(v)$, then $\beta_D^*(v, \tilde{s}) = \beta_{LD}(v)$.

Proposition 8. (i) β_{LD} is strictly decreasing in v and β_{HD} is strictly increasing in v . (ii) The threshold τ_D is increasing in v if and only if

$$\delta(\beta_{HD}) - \delta(\beta_{LD}) > \mathbb{E}_{s_0} \left[\Theta(\beta_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) \mid \tilde{s} \neq s_0 \right], \quad (\text{B.1})$$

where $\Theta(\beta_D, s_D)$ corresponds to all the components of the manufacturer's profits except for $\delta(\beta_D)$ and s_{D_H} (s_{D_L}) is equal to the optimal value of s_D given a signal immediately to the left (immediately to the right) of $\tilde{s} = \tau_D(v)$.

Proposition 9. *The threshold τ_D is strictly increasing in q .*

Proposition 10. $\frac{ds^*(s_0, \beta_D^*(\tilde{s}, q), q)}{dq} \geq 0$ for all $s_0, \tilde{s} \in [m, M]$.

Theorem 7. *If $\tau_H < \hat{s}$, then for any level of visibility v , $\Delta_\Pi(v, \tilde{s})$ is continuous and*

- (i) *constant in \tilde{s} for $\tilde{s} \in [m, \tau]$;*
- (ii) *decreasing in \tilde{s} for $\tilde{s} \in [\tau, \tau_D]$;*
- (iii) *increasing in \tilde{s} for $\tilde{s} \in [\tau_D, \tau_H]$;*
- (iv) *decreasing in \tilde{s} for $\tilde{s} \in [\tau_H, \hat{s}]$;*
- (v) *constant in \tilde{s} for $\tilde{s} \in [\hat{s}, M]$.*

Furthermore, there exists $\bar{\tau} \in [m, \tau_D)$ such that $\Delta_\Pi(v, \tilde{s}) < 0$ for all $\tilde{s} \geq \bar{\tau}$.

Theorem 7 implies that if $\tau_H < \hat{s}$, then there exists some signal $\bar{\tau} \geq m$ such that the manufacturer prefers *not* to voluntarily disclose SR information to the consumers for any $\tilde{s} \geq \bar{\tau}$. In fact, it is likely that the manufacturer prefers not to disclose for any signal; i.e., $\bar{\tau} = m$ is likely. This is because the conditions for $\bar{\tau} > m$ are generally hard to satisfy. In particular, (a) the manufacturer must prefer disclosing a higher value of SR for the lowest possible signals (i.e., close to m) than for the highest possible signals (i.e., close to M); and (b) the manufacturer's optimal disclosure decision must be strictly smaller than \hat{s} even when $\tilde{s} = M$.

Proposition 11. *There exists a range of supplier types $s_0 \in (m, M)$ such that, for high enough levels of visibility v , the supplier's expected level of SR, $ES(v, s_0)$, is strictly decreasing in v .*

Proposition 12. *For any level of visibility v , the manufacturer's best estimate of the supplier's final level of SR, $s^*(\tilde{s}, \beta^*(v, \tilde{s}))$ under No Manufacturer Disclosure and $s^*(\tilde{s}, \beta_D^*(v, \tilde{s}))$ under Manufacturer Disclosure, can be characterized as follows:*

(i) If $\tilde{s} \leq \tau(v)$, then $s^*(\tilde{s}, \beta^*(v, \tilde{s})) = s^*(\tilde{s}, \bar{\beta}_H) = s_u^*(\bar{\beta}_H)$;

(ii) If $\tilde{s} > \tau(v)$, then $s^*(\tilde{s}, \beta^*(v, \tilde{s})) = s^*(\tilde{s}, \beta_L) = \tilde{s}$;

(iii) If $\tilde{s} \leq \tau_D(v)$, then $s^*(\tilde{s}, \beta_D^*(v, \tilde{s})) = s^*(\tilde{s}, \bar{\beta}_{HD}) = s_{uD}^*(\bar{\beta}_{HD})$;

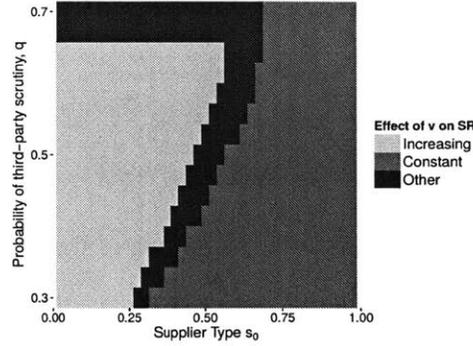
(iv) If $\tilde{s} > \tau_D(v)$, then $s^*(\tilde{s}, \beta_D^*(v, \tilde{s})) = s^*(\tilde{s}, \beta_{LD}) = \tilde{s}$.

Furthermore, $\lim_{\tilde{s} \uparrow \tau(v)} s^*(\tilde{s}, \beta^*) > \lim_{\tilde{s} \downarrow \tau(v)} s^*(\tilde{s}, \beta^*)$ and $\lim_{\tilde{s} \uparrow \tau_D(v)} s^*(\tilde{s}, \beta_D^*) > \lim_{\tilde{s} \downarrow \tau_D(v)} s^*(\tilde{s}, \beta_D^*)$. That is, the best estimates are discontinuous at $\tau(v)$ under No Manufacturer Disclosure and at $\tau_D(v)$ under Manufacturer Disclosure.

B.2 Examining the Expected Final SR Level

We discuss here more details regarding the effect of visibility on the expected final level of SR given the supplier's true type s_0 , with the expectation taken over all possible signals that the manufacturer may observe. We focus on the No Manufacturer Disclosure setting, noting that the results in the Manufacturer Disclosure setting are qualitatively the same. Figure B-1 illustrates the results with respect to the supplier's type (s_0) and the probability of third-party scrutiny (q). We identify three possible scenarios: (i) the expected final SR level is *increasing* in visibility; (ii) the expected final SR level is *constant* in visibility; and (iii) for at least some ranges of v , the expected final SR level is *decreasing* in visibility (denoted as "Other" in Figure B-1). We first observe from Figure B-1 that increased visibility primarily has a nonnegative effect on the expected final SR level (i.e., scenarios (i) and (ii)). For scenario (i), when the supplier's initial SR level is low, greater visibility increases the expected final SR level for a wide range of q values. This is due to the positive effects of greater visibility on both the accuracy of the signal and the manufacturer's investment. Specifically, greater visibility makes the manufacturer more certain that the supplier's current practice are poor. Hence, the manufacturer invests a high amount ($\bar{\beta}_H$). The manufacturer's high investment increases with visibility (Proposition 1(i)) and results in an improved SR level by the supplier ($s^*(s_0, \bar{\beta}_H) = s_u^*(\bar{\beta}_H)$ for $s_0 \leq \tau$; see Proposition 12).

Figure B-1: Effect of Visibility on the Expected Final Level of SR



Parameters: $\alpha = 0.4$, $r = 0.5$, $w = 0.15$, $c = 0$. Functions: $\rho(s) = e^{1.55s}$, $\gamma(s) = 0.5 - e^s$, $\delta(\beta) = \frac{0.7\beta^2}{1-\beta}$. The supplier's type s_0 is distributed uniformly between $m = 0$ and $M = 1$.

Conversely, for scenario (ii) the supplier's initial SR level is already high, and thus, visibility has no impact on the expected final SR level because the manufacturer invests β_L and $s^*(s_0, \beta_L) = s_0$ for $s_0 > \tau$ (Proposition 12). That is, the supplier does not change his SR level in response to the manufacturer's low investment. As the probability of third-party scrutiny increases, the constant region shrinks because (a) the manufacturer invests a high amount for a wider range of supplier types due to the increased chance of the consumers observing the supplier's final SR level (Proposition 2); and (b) the supplier increases his SR level in response to the manufacturer's high investment.

For scenario (iii), Figure B-1 highlights two regions where an increase in visibility can lead to a decrease in the expected final SR level: (I) when s_0 takes intermediate values, and (II) when s_0 is low and q is high. Region (I) is discussed in §3.3.1. In region (II), the probability of third-party scrutiny is so high that the manufacturer invests the highest possible amount $\hat{\beta}$ even when visibility is low. Consequently, an increase in visibility does not yield an increase in the manufacturer's high investment. It does, however, decrease her low investment β_L (Proposition 1(i)). As a result, the expected final SR level decreases with visibility.

B.3 The Effect of q on β_L and β_H

The key element that determines whether an increase in q has a positive or negative effect on β_L and β_H is how the rates of increase of the functions $\rho(s)$ and $\gamma(s)$ – captured by $R_\rho(s)$ and $R_\gamma(s)$ – change with s . If $\rho(s)$ increases extremely fast in s for high values of s and/or $\gamma(s)$ decreases extremely fast as s decreases for low values of s (e.g., if $R'_\rho(s) \ll 0$ and/or $R'_\gamma(s) \gg 0$), then the manufacturer's optimal investment in the supplier's capabilities is likely to *decrease* with q . In this case, as q increases, the supplier already has a strong incentive to improve SR. Thus, attempting to further motivate the supplier to increase s is too costly for the manufacturer. Conversely, if the rates of increase of $\rho(s)$ and $\gamma(s)$ are similar across all values of s (e.g., if $R'_\rho(s) = R'_\gamma(s) = 0$; or equivalently, if $\rho(s)$ and $\gamma(s)$ are exponential in s), then β is more likely to *increase* with q . In this case, the manufacturer complements an increase in q with a greater β , as the corresponding increase in s is worth the additional cost.

Appendix C

Proofs Chapter 3

We next provide the proofs for the supplier's and manufacturer's optimal strategies and the comparisons between the No Manufacturer Disclosure and Manufacturer Disclosure settings. Note that the No Manufacturer Disclosure setting is a particular case of the Manufacturer Disclosure setting, where (i) the value of s_D is fixed and equal to \hat{s} ; and (ii) there is no penalty associated with the disclosure of SR practices; i.e., $p = 0$. Therefore, we present proofs only for the Manufacturer Disclosure setting and then indicate where appropriate how they can be used to prove the results from the No Manufacturer Disclosure setting. For our proofs, we also note that Assumptions 2 and 3 are equivalent to the following:

Assumption 2. *The functions $\gamma(s)$ and $\rho(s)$ satisfy $\rho'(m) < \alpha q(w - c)\gamma'(m)$.*

Assumption 3. *$\delta'(\beta_M)(1 - \beta_M) > (\alpha q(r - w)\gamma'(M) + pq)(R_\rho(M) - R_\gamma(M))^{-1}$, where β_M is defined as the value of β_D such that, if the third party disclosed the value of s to consumers regardless of how it compared to s_D (or \hat{s} in the No Manufacturer Disclosure setting), then $s^*(s_0 = m, \beta_M) = M$.*

For ease of notation, we often write $\beta_{LD}(v)$, $\beta_{HD}(v)$, $\bar{\beta}_{HD}(v)$, and $\tau_D(v)$ simply as β_{LD} , β_{HD} , $\bar{\beta}_{HD}$, and τ_D , respectively, and we also define $R(s) \equiv (R_\rho(s) - R_\gamma(s))^{-1}$. Similarly, throughout this Appendix we use the optimal solution in Lemma 1, $\bar{s}(\beta_D)$, which we often write simply as \bar{s} . Finally, regarding the supplier's best response we also use the fact that $s^*(s_0, \beta_D, s_D) = \max\{s_0, \min\{s_D, \bar{s}\}\}$, as noted in Appendix C.1.

C.1 The Supplier's Best Response

In this section, we demonstrate our results regarding the supplier's best response to the manufacturer's choice of s_D and β_D . In Lemma 1, we define a modified objective function for the supplier where: (i) the supplier can select any value of s , regardless of s_0 and (ii) any value of s is communicated by the third party to the consumers with probability q . We then use this result to prove Theorem 5 (and Theorem 1 in the No Manufacturer Disclosure setting).

Lemma 1. *For any s_D , let us define*

$$F(s) \equiv \left(1 - r + \alpha(1 - q)\gamma(s_D) + \alpha q\gamma(s)\right)(w - c) - (1 - \beta_D)(\rho(s) - \rho(s_0)).$$

$F(s)$ is strictly concave in s and $\max_s F(s)$ admits a unique optimal solution \check{s} . Furthermore, $\frac{\partial \check{s}}{\partial \beta_D} > 0$, $\frac{\partial \check{s}}{\partial q} > 0$, and \check{s} is independent of v , p , and \tilde{s} .

Proof. First, assume $\beta_D < 1$. The first order condition (FOC) of the optimization problem $\max_s F(s)$ can be written as

$$F'(s) = \alpha q(w - c)\gamma'(s) - (1 - \beta_D)\rho'(s) = 0. \quad (\text{C.1})$$

Similarly, the second order condition (SOC) is given by

$$F''(s) = \alpha q(w - c)\gamma''(s) - (1 - \beta_D)\rho''(s) < 0, \quad (\text{C.2})$$

where the last inequality follows directly from (i) $\gamma(s)$ is strictly concave in s ; and (ii) $\rho(s)$ is convex in s . Thus, the FOC in Equation (C.1) is sufficient for optimality.

To show that the optimization problem admits a unique solution, we note that (i) $F'(m) \geq \alpha q(w - c)\gamma'(m) - \rho'(m) > 0$, where the second inequality follows directly from Assumption 2; (ii) $\lim_{s \rightarrow \infty} F'(s) = -\infty$, as $\rho(s)$ is strictly increasing in s and $\gamma(s)$ is strictly increasing and bounded above by r ; and (iii) $F'(s)$ is continuous in s because $\gamma(s)$ and $\rho(s)$ are continuously differentiable. Therefore, $\exists! \check{s} > m$ such that $F'(\check{s}) = 0$ and thus \check{s} is the unique optimal solution to $\max_s F(s)$.

If instead $\beta_D = 1$, then the supplier does not incur a cost for increasing its SR level. In addition, $\gamma(s)$ is strictly increasing in s and hence the solution to the supplier's problem $\max_s F(s)$ becomes $\check{s} = \infty$.

Finally, we note that for any s , $\frac{\partial^2 F}{\partial s \partial \beta_D} = \rho'(s) > 0$ and $\frac{\partial^2 F}{\partial s \partial q} = \alpha(w - c)\gamma'(s) > 0$, as both $\rho(s)$ and $\gamma(s)$ are strictly increasing in s . Therefore, by Topkis's theorem, it follows that $\frac{\partial \check{s}}{\partial \beta_D} > 0$ and $\frac{\partial \check{s}}{\partial q} > 0$. Similarly, the FOC in Equation (C.1) is independent of v , p , and \check{s} , and hence so is any solution to it. \square

Proof of Theorem 5. We first note that the function $F(s)$ defined in Lemma 1 corresponds to the supplier's profit function if any value of s was communicated by the third party to the consumers. However, because the third party only communicates $\min\{s, s_D\}$, the supplier has no incentive to select $s > s_D$. Specifically, the supplier's profit function is equal to $F(s)$ for any $s \leq s_D$; and equal to $(1 - r + \alpha\gamma(s_D))(w - c) - (1 - \beta_D)(\rho(s) - \rho(s_0))$ for any $s > s_D$. Therefore, it is strictly decreasing in s if $s > s_D$. As a result, using the concavity of $F(s)$, the unconstrained optimal solution to Equation (3.9) is equal to $s_{uD}^*(\beta_D, s_D) = \min\{s_D, \check{s}\}$. Thus, $s_{uD}^*(\beta_D, s_D) \leq s_D$.

Similarly, by the concavity of $F(s)$, we can incorporate the constraint $s \geq s_0$ by noting that if $s_{uD}^*(\beta_D, s_D) < s_0$, then it is optimal for the supplier to select $s^* = s_0$.

To summarize, we can write the supplier's constrained optimal solution as $s^*(s_0, \beta_D, s_D) = \max\{s_0, s_{uD}^*(\beta_D, s_D)\} = \max\{s_0, \min\{s_D, \check{s}\}\}$. Furthermore, from Lemma 1, it follows that $\frac{\partial s^*(s_0, \beta_D, s_D)}{\partial \beta_D} \geq 0$ and $\frac{\partial s^*(s_0, \beta_D, s_D)}{\partial q} \geq 0$. \square

Note: No Manufacturer Disclosure setting. If we replace s_D by \hat{s} and β_D by β , the above proof also applies to the supplier's optimal strategy in Theorem 1.

C.2 The Manufacturer's Optimal Strategy

In this section, we show our results concerning the manufacturer's optimal strategy. First, Lemma 2 shows how the solution to the modified supplier's objective function

in Lemma 1 depends on β_D . In Lemmas 3 and 4, we define two new functions which we then use in Lemma 6 to rewrite – based on Lemma 5 – the manufacturer’s expected profits. This alternative way of writing the manufacturer’s profit function simplifies our analysis of her optimal strategy. Theorem 13 then identifies the optimal choice of β_D given s_D . Using this result, Corollary 4 simplifies the analysis of the supplier’s best response to the optimal β_D and s_D by noting that it can be written only as a function of β_D .

We then analyze the manufacturer’s optimal choice of disclosure s_D . Proposition 14 then identifies the manufacturer’s optimal choice of s_D for a modified objective function and shows that, if $\check{s}(\beta_D) \leq s_D$, then the manufacturer’s expected profits are unimodal in s_D . Lemma 7 shows that there is at most one candidate solution (i.e., a local maximum) that satisfies $\check{s}(\beta_D) = s_D$. Using all of the above results, Lemma 8 identifies the three possible solutions to the manufacturer’s optimization problem and when they should be considered as candidate solutions. This allows us to then show in Lemma 9 that there exists an upper bound to the optimal β_D .

Finally, we show the manufacturer’s optimal choice of (β_D, s_D) , summarized in the proof of Theorems 3 and 6. We conclude this section by identifying some additional results regarding the threshold τ_D in Theorem 6 and, given the manufacturer’s optimal strategy, what is the resulting value of her best estimate of the supplier’s final level of SR.

Lemma 2. $\frac{d\check{s}}{d\beta_D} = \frac{R(\check{s})}{1 - \beta_D}$.

Proof. By Lemma EC.1 we know that the FOC of the optimization problem $\max_s F(s)$ (see Equation (C.1)) admits a unique solution \check{s} . Using the Implicit Function Theorem and taking the derivative of Equation (C.1) with respect to β_D , and by simple algebraic manipulation, we obtain:

$$\frac{d}{d\beta_D} \left(\alpha q(w - c)\gamma'(s) - (1 - \beta_D)\rho'(s) \right) = 0 \quad (\text{C.3})$$

$$\alpha q(w - c)\gamma''(\tilde{s})\frac{d\tilde{s}}{d\beta_D} - (1 - \beta_D)\rho''(s)\frac{d\tilde{s}}{d\beta_D} + \rho'(\tilde{s}) = 0 \quad (\text{C.4})$$

$$\left(\frac{\alpha q(w - c)}{\rho'(\tilde{s})}\gamma''(\tilde{s}) - (1 - \beta_D)\frac{\rho''(\tilde{s})}{\rho'(\tilde{s})} \right) \frac{d\tilde{s}}{d\beta_D} + 1 = 0 \quad (\text{C.5})$$

$$\left((1 - \beta_D)\frac{\gamma''(\tilde{s})}{\gamma'(\tilde{s})} - (1 - \beta_D)\frac{\rho''(\tilde{s})}{\rho'(\tilde{s})} \right) \frac{d\tilde{s}}{d\beta_D} = -1 \quad (\text{C.6})$$

$$\frac{d\tilde{s}}{d\beta_D} = \left((1 - \beta_D) (R_\rho(\tilde{s}) - R_\gamma(\tilde{s})) \right)^{-1} \quad (\text{C.7})$$

$$\frac{d\tilde{s}}{d\beta_D} = \frac{R(\tilde{s})}{1 - \beta_D}, \quad (\text{C.8})$$

where Equation (C.6) follows from the equality given by the FOC and Equation (C.8) uses the definitions of $R_\gamma(s)$ and $R_\rho(s)$. \square

Lemma 3. *Let us define*

$$\begin{aligned} G_L(\beta_D, s_D, \tilde{s}) &\equiv (r - w) \left[1 - r + \alpha(1 - q)\gamma(s_D) + \alpha q \left(v\gamma(\min\{\tilde{s}, s_D\}) \right. \right. \\ &\quad \left. \left. + (1 - v)(\gamma(\tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} \gamma(x)\phi(x)dx + \gamma(s_D)(1 - \Phi(s_D))) \right) \right] - \delta(\beta_D) \\ &\quad - pq \left(v(s_D - \min\{\tilde{s}, s_D\}) + (1 - v)((s_D - \tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x)\phi(x)dx) \right) \end{aligned} \quad (\text{C.9})$$

For any s_D and \tilde{s} , $\exists! \beta_D$ which we refer to as β_{LD} that maximizes $G_L(\beta_D, s_D, \tilde{s})$ and it satisfies the first order condition

$$G'_L(\beta_D, s_D, \tilde{s}) = (\alpha q(r - w)\gamma'(\tilde{s}) + pq) \frac{R(\tilde{s})}{1 - \beta_D} (1 - v)\Phi(\tilde{s}) - \delta'(\beta_D) = 0. \quad (\text{C.10})$$

Furthermore, $G'_L(\beta_D, s_D, \tilde{s})$ and β_{LD} are independent of s_D and \tilde{s} .

Proof. We first study the partial derivative $\frac{\partial G_L(\beta_D, s_D, \tilde{s})}{\partial \beta_D}$, which for ease of notation we

write as $G'_L(\beta_D)$. Using Leibniz's rule and the chain rule:

$$\begin{aligned} G'_L(\beta_D) &= \alpha q (r - w) (1 - v) \left(\gamma(\check{s}) \phi(\check{s}) + \gamma'(\check{s}) \frac{d\check{s}}{d\beta_D} \Phi(\check{s}) - \gamma(\check{s}) \phi(\check{s}) \right) \\ &\quad - \delta'(\beta_D) - p q (1 - v) \left(-\Phi(\check{s}) + (s_D - \check{s}) \phi(\check{s}) - (s_D - \check{s}) \phi(\check{s}) \right) \frac{d\check{s}}{d\beta_D} = 0. \end{aligned} \quad (\text{C.11})$$

Rearranging terms and applying Lemma 2, this simplifies to

$$G'_L(\beta_D) = (\alpha q (r - w) \gamma'(\check{s}) + p q) \frac{R(\check{s})}{1 - \beta_D} (1 - v) \Phi(\check{s}) - \delta'(\beta_D) = 0, \quad (\text{C.12})$$

which represents the first order condition that the optimal β_D^* must satisfy. We note that this condition is independent of s_D and \check{s} , which proves that the optimal β_D is also independent of these variables.

We first show that a solution to the FOC in Equation (C.12) exists. At $\beta_D = 0$, by definition we know that $\delta'(\beta_D = 0) = 0$. Furthermore, all the other terms in $G'_L(\beta_D)$ in Equation (C.12) are positive: $\rho(s)$ is strictly increasing in s and $R(s) = (R_\rho(s) - R_\gamma(s))^{-1}$ is positive because $\rho(s)$ is convex and $\gamma(s)$ is strictly concave in s . Thus, it follows that $G'_L(\beta_D = 0) > 0$. In addition, we know that $\lim_{\beta_D \rightarrow 1} \check{s} = \infty$ and thus we can define $\beta_M < 1$ such that $\check{s}(\beta_M) = M$ and hence $\Phi(\check{s}(\beta_M)) = 1$. Therefore, we know that

$$\begin{aligned} G'_L(\beta_D = \beta_M) &= (\alpha q (r - w) \gamma'(M) + p q) \frac{R(M)}{1 - \beta_M} (1 - v) - \delta'(\beta_M) \\ &\leq (\alpha q (r - w) \gamma'(M) + p q) \frac{R(M)}{1 - \beta_M} - \delta'(\beta_M) \\ &= \frac{1}{1 - \beta_M} \left((\alpha q (r - w) \gamma'(M) + p q) R(M) - \delta'(\beta_M) (1 - \beta_M) \right) < 0, \end{aligned}$$

where the last inequality follows directly from Assumption 3. Finally, we note that by our assumptions, $\rho(s)$, $\gamma(s)$ and $\delta(\beta_D)$ are twice-continuously differentiable and $\Phi(s)$ is continuous in s . Therefore, $G'_L(s)$ is continuous in s and since $G'_L(0) > 0$ and $G'_L(\beta_M) < 0$, there exists β_D that satisfies the FOC in Equation (C.10).

Next, let us define β_{LD} as a solution to the FOC in Equation (C.10) and $s_L =$

$\bar{s}(\beta_{LD})$. To show that this solution is unique, we study the second derivative of $G_L(\beta_D = \beta_{LD})$ with respect to β_D :

$$\begin{aligned} G_L''(\beta_{LD}) = & \left(\alpha q(r-w)\gamma''(s_L) \frac{R(s_L)}{1-\beta_{LD}} (1-v)\Phi(s_L) \right. \\ & + (\alpha q(r-w)\gamma'(s_L) + pq) \frac{R'(s_L)}{1-\beta_{LD}} (1-v)\Phi(s_L) \\ & \left. + (\alpha q(r-w)\gamma'(s_L) + pq) \frac{R(s_L)}{1-\beta_{LD}} (1-v)\phi(s_L) \right) \frac{ds_L}{d\beta_D} \\ & + (\alpha q(r-w)\gamma'(s_L) + pq) \frac{R(s_L)}{(1-\beta_{LD})^2} (1-v)\Phi(s_L) - \delta''(\beta_{LD}). \end{aligned}$$

Because of Assumption 4, $R'(s) \leq 0$ for any s . Furthermore, since $\gamma(s)$ is strictly concave in s , $\gamma''(s_L) < 0$ and hence

$$\begin{aligned} G_L''(\beta_{LD}) < & (\alpha q(r-w)\gamma'(s_L) + pq) \frac{R(s_L)}{1-\beta_{LD}} (1-v)\phi(s_L) \frac{R(s_L)}{1-\beta_{LD}} \\ & + (\alpha q(r-w)\gamma'(s_L) + pq) \frac{R(s_L)}{(1-\beta_{LD})^2} (1-v)\Phi(s_L) - \delta''(\beta_{LD}), \end{aligned}$$

or equivalently,

$$\begin{aligned} G_L''(\beta_{LD}) < & \left((\alpha q(r-w)\gamma'(s_L) + pq) \frac{R(s_L)}{1-\beta_{LD}} (1-v)\Phi(s_L) \right) \\ & \times \left(\frac{\phi(s_L)}{\Phi(s_L)} \frac{R(s_L)}{1-\beta_{LD}} + \frac{1}{1-\beta_{LD}} \right) - \delta''(\beta_{LD}). \end{aligned}$$

We note that first term in parentheses in the right-hand side corresponds to the FOC in Equation (C.10) minus the derivative of the cost, $\delta'(\beta_D)$. Thus, since β_{LD} satisfies the FOC, we can simplify this equation as:

$$G_L''(\beta_{LD}) \leq \delta'(\beta_{LD}) \left(\frac{\phi(s_L)}{\Phi(s_L)} \frac{R(s_L)}{1-\beta_{LD}} + \frac{1}{1-\beta_{LD}} \right) - \delta''(\beta_{LD}).$$

Finally, from simple algebraic manipulation, we obtain

$$G_L''(\beta_{LD}) \leq \frac{\delta'(\beta_{LD})}{1-\beta_{LD}} \left(1 + \frac{\phi(s_L)}{\Phi(s_L)} R(s_L) - \frac{\delta''(\beta_{LD})}{\delta'(\beta_{LD})} (1-\beta_{LD}) \right) < 0,$$

where the last inequality follows directly from the fact that the expression inside the parenthesis is negative by Assumption 5.

Therefore, we have shown that any β_{LD} that satisfies Equation (C.10) must be unique, as any other critical point $\beta_D^* \neq \beta_{LD}$ would also need to be a local maximum. Since $G'_L(\beta_D)$ is continuous in β_D , such an additional critical point cannot exist. \square

Lemma 4. *Let us define*

$$\begin{aligned} G_H(\beta_D, s_D) \equiv & (r - w) \left[1 - r + \alpha(1 - q)\gamma(s_D) + \alpha q \left(v\gamma(\tilde{s}) \right. \right. \\ & \left. \left. + (1 - v)(\gamma(\tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} \gamma(x)\phi(x)dx + \gamma(s_D)(1 - \Phi(s_D))) \right) \right] - \delta(\beta_D) \\ & - pq \left(v(s_D - \tilde{s}) + (1 - v)((s_D - \tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x)\phi(x)dx \right), \quad (\text{C.13}) \end{aligned}$$

which is independent of the signal \tilde{s} . For any s_D , $\exists! \beta_D$ which we refer to as β_{HD} that maximizes $G_H(\beta_D, s_D)$ and it satisfies the first order condition

$$\frac{\partial G_H(\beta_D, s_D)}{\partial \beta_D} = (\alpha q (r - w) \gamma'(\tilde{s}) + pq) \frac{R(\tilde{s})}{1 - \beta_D} (v + (1 - v) \Phi(\tilde{s})) - \delta'(\beta_D) = 0. \quad (\text{C.14})$$

Furthermore, $G'_H(\beta_D, s_D)$ and β_{HD} are independent of s_D .

Proof. The proof of this lemma follows the same logic as that of Lemma 3 and is therefore omitted. \square

Lemma 5. *The manufacturer's posterior distribution of s_0 given signal \tilde{s} is of a mixed type, with $\mathbb{P}[s_0 = \tilde{s} | \tilde{s}] = v$ and pdf $\psi(x | \tilde{s}) = (1 - v) \phi(x)$ for any $x \in [m, M] \neq \tilde{s}$.*

Proof. Let us define $\Psi(x | \tilde{s}) = \mathbb{P}[s_0 \leq x | \tilde{s}]$, the posterior cdf of s_0 given \tilde{s} . Using the law of total probability, we know that

$$\Psi(x | \tilde{s}) = \mathbb{P}[\tilde{s} = s_0] \mathbb{P}[s_0 \leq x | \tilde{s} \cap \tilde{s} = s_0] + \mathbb{P}[\tilde{s} \neq s_0] \mathbb{P}[s_0 \leq x | \tilde{s} \cap \tilde{s} \neq s_0].$$

If the signal is correct (i.e., if $\tilde{s} = s_0$), then the probability that s_0 is less than or

equal to a given value x can be rewritten as

$$\mathbb{P}[s_0 \leq x \mid \tilde{s} \cap \tilde{s} = s_0] = \mathbb{P}[\tilde{s} \leq x].$$

Thus, its value depends on the relationship between \tilde{s} and x : if $\tilde{s} \leq x$, then this probability is equal to 1; otherwise, it is equal to 0. This, as both \tilde{s} and x are known. Conversely, if the signal is not correct (i.e., if $\tilde{s} \neq s_0$), then the signal is uninformative and we obtain that $\mathbb{P}[s_0 \leq x \mid \tilde{s} \cap \tilde{s} \neq s_0]$ is simply equal to $\mathbb{P}[s_0 \leq x]$. Therefore, the above equation simplifies to

$$\begin{aligned} \Psi(x \mid \tilde{s}) &= \mathbb{P}[\tilde{s} = s_0] \mathbb{1}_{\tilde{s} \leq x} + (1 - \mathbb{P}[\tilde{s} = s_0]) \Phi(x) \\ \Psi(x \mid \tilde{s}) &= v \mathbb{1}_{\tilde{s} \leq x} + (1 - v) \Phi(x), \end{aligned}$$

as the prior cdf of s_0 is equal to $\Phi(x)$ and the probability that the signal is correct is equal to v . Thus, the pdf of s_0 given \tilde{s} is continuous and equal to $\psi(x \mid \tilde{s}) = (1 - v) \phi(x)$ a.e. in $[m, M]$. The only exception is $x = \tilde{s}$, where $\mathbb{P}[s_0 = \tilde{s} \mid \tilde{s}] = v$. \square

Lemma 6. *Given a signal \tilde{s} , the manufacturer's expected profit function as a function of β_D and s_D can be written as*

$$\begin{aligned} \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) \mid \tilde{s}] &= \left(G_L(\beta_D, s_D, \tilde{s}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + G_H(\beta_D, s_D) \mathbb{1}_{\tilde{s} > \tilde{s}} \right) \mathbb{1}_{\tilde{s} \leq s_D} \\ &\quad + \left((1 - r + \alpha \gamma(s_D)) (r - w) - \delta(\beta_D) \right) \mathbb{1}_{\tilde{s} > s_D}. \end{aligned} \quad (\text{C.15})$$

Proof. This result follows directly from the definitions of G_L (Lemma 3) and G_H (Lemma 4), the posterior distribution of s_0 given \tilde{s} , and algebraic manipulation. First, we compute the term corresponding to the consumers' SR utility in Equation (3.9). For simplicity, we write $\mathbb{E}_{s_0} [\gamma(\min\{s^*(s_0, \beta_D, s_D), s_D\}) \mid \tilde{s}]$ simply as $\mathbb{E}[\gamma \mid \tilde{s}]$. Using

the posterior distributions of s_0 given \tilde{s} from Lemma 5:

$$\begin{aligned}
\mathbb{E}[\gamma | \tilde{s}] &= v \gamma(\min\{s^*(s_0 = \tilde{s}, \beta_D, s_D), s_D\}) \\
&\quad + (1 - v) \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta_D, s_D), s_D\}) \mid s_0 \neq \tilde{s} \right] \\
&= v \gamma(\min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v \gamma(\min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} > \tilde{s}} \\
&\quad + (1 - v) \left(\gamma(\tilde{s}) \Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} \gamma(x) \phi(x) dx + \gamma(s_D) (1 - \Phi(s_D)) \right) \mathbb{1}_{\tilde{s} \leq s_D} \\
&\quad + (1 - v) \gamma(s_D) \mathbb{1}_{\tilde{s} > s_D},
\end{aligned}$$

where the last step follows from the definition of \check{s} in Lemma 1 and the proof of Theorem 5, whereby $s^*(s_0, \beta_D, s_D) = \max\{s_0, \min\{s_D, \check{s}\}\}$. Finally, rearranging terms:

$$\begin{aligned}
\mathbb{E}[\gamma | \tilde{s}] &= \left(v \gamma(\min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v \gamma(\tilde{s}) \mathbb{1}_{\tilde{s} > \tilde{s}} + (1 - v) \left(\gamma(\tilde{s}) \Phi(\tilde{s}) \right. \right. \\
&\quad \left. \left. + \int_{\tilde{s}}^{s_D} \gamma(x) \phi(x) dx + \gamma(s_D) (1 - \Phi(s_D)) \right) \right) \mathbb{1}_{\tilde{s} \leq s_D} + \gamma(s_D) \mathbb{1}_{\tilde{s} > s_D}. \quad (\text{C.16})
\end{aligned}$$

We can use a similar approach to compute the posterior distribution of s_0 given \tilde{s} , $E_{s_0} \left[(s_D - s^*(s_0, \beta_D))^+ \mid \tilde{s} \right]$,¹ which we denote simply as $\mathbb{E}[\text{penalty} \mid \tilde{s}]$:

$$\mathbb{E}[\text{penalty} \mid \tilde{s}] = v (s_D - s^*(s_0 = \tilde{s}, \beta_D))^+ + (1 - v) E_{s_0} \left[(s_D - s^*(s_0, \beta_D))^+ \mid s_0 \neq \tilde{s} \right].$$

Or equivalently,

$$\begin{aligned}
\mathbb{E}[\text{penalty} \mid \tilde{s}] &= v (s_D - \min\{s^*(s_0 = \tilde{s}, \beta_D, s_D), s_D\}) \\
&\quad + (1 - v) E_{s_0} \left[(s_D - \min\{s^*(s_0, \beta_D, s_D), s_D\}) \mid s_0 \neq \tilde{s} \right] \\
\mathbb{E}[\text{penalty} \mid \tilde{s}] &= v (s_D - \min\{s^*(s_0 = \tilde{s}, \beta_D, s_D), s_D\}) \\
&\quad + (1 - v) \left((s_D - \tilde{s}) \Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x) \phi(x) dx \right) \mathbb{1}_{\tilde{s} \leq s_D} \\
\mathbb{E}[\text{penalty} \mid \tilde{s}] &= v (s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v (s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} > \tilde{s}} \\
&\quad + (1 - v) \left((s_D - \tilde{s}) \Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x) \phi(x) dx \right) \mathbb{1}_{\tilde{s} \leq s_D}.
\end{aligned}$$

¹The positive part of a quantity x , denoted as x^+ , is defined as $x^+ = \max\{x, 0\}$.

Finally, we note that both of the terms that are multiplied by v , i.e. $(s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}}$ and $(s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} > \tilde{s}}$, are exactly equal to 0 if $\tilde{s} > s_D$; and they are equal to $(s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}}$ and $(s_D - \tilde{s})$, respectively, if $\tilde{s} \leq s_D$. Thus, we can rearrange terms to multiply all terms in the above equation by the indicator function $\mathbb{1}_{\tilde{s} > \tilde{s}}$ as follows:

$$\begin{aligned} \mathbb{E}[\text{penalty} | \tilde{s}] &= \left(v (s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v (s_D - \tilde{s}) \mathbb{1}_{\tilde{s} > \tilde{s}} \right. \\ &\quad \left. + (1 - v) \left((s_D - \tilde{s}) \Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x) \phi(x) dx \right) \right) \mathbb{1}_{\tilde{s} \leq s_D}. \quad (\text{C.17}) \end{aligned}$$

Thus, updating the manufacturer's expected profits in Equation (3.9) with the results in Equations (C.16) and (C.17), we can derive the following expression:

$$\begin{aligned} \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \tilde{s}] &= \left\{ \left(1 - r + \alpha(1 - q)\gamma(s_D) + \alpha q \left(v\gamma(\min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v\gamma(\tilde{s}) \mathbb{1}_{\tilde{s} > \tilde{s}} \right. \right. \right. \\ &\quad \left. \left. + (1 - v)(\gamma(\tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} \gamma(x)\phi(x)dx + \gamma(s_D)(1 - \Phi(s_D))) \right) \right) (r - w) \\ &\quad - \delta(\beta_D) - pq \left(v(s_D - \min\{\tilde{s}, s_D\}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + v(s_D - \tilde{s}) \mathbb{1}_{\tilde{s} > \tilde{s}} \right. \\ &\quad \left. + (1 - v) \left((s_D - \tilde{s})\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} (s_D - x)\phi(x)dx \right) \right) \left. \right\} \mathbb{1}_{\tilde{s} \leq s_D} \\ &\quad + \left\{ (1 - r + \alpha\gamma(s_D))(r - w) - \delta(\beta_D) \right\} \mathbb{1}_{\tilde{s} > s_D}. \end{aligned}$$

Using the definitions of $G_L(\beta_D, s_D, \tilde{s})$ and $G_H(\beta_D, s_D)$ in Lemmas 3 and 4, respectively, the above expression simplifies to

$$\begin{aligned} \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \tilde{s}] &= \left(G_L(\beta_D, s_D, \tilde{s}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + G_H(\beta_D, s_D) \mathbb{1}_{\tilde{s} > \tilde{s}} \right) \mathbb{1}_{\tilde{s} \leq s_D} \\ &\quad + \left((1 - r + \alpha\gamma(s_D))(r - w) - \delta(\beta_D) \right) \mathbb{1}_{\tilde{s} > s_D}. \quad \square \end{aligned}$$

Proposition 13. *Given any s_D , let us define $\hat{\beta}_{s_D}$ such that $\check{s}(\hat{\beta}_{s_D}) = s_D$. For any level of visibility v and for any fixed s_D , there exist $0 \leq \beta_{LD} \leq \beta_{HD} < 1$ such that the manufacturer's optimal choice of β_D can be characterized as follows:*

(a) If $\beta_{LD} \geq \hat{\beta}$, then $\beta_D^*(v, \tilde{s}) = \hat{\beta}_{s_D}$ for all \tilde{s} .

(b) If $\beta_{LD} < \hat{\beta}_{s_D}$, then there exists a threshold $\tau_{s_D} \in (m, \hat{s})$ such that:

(i) If $\tilde{s} \leq \tau_{s_D}$, then $\beta_D^*(v, \tilde{s}) = \min \{\beta_{HD}, \hat{\beta}_{s_D}\}$;

(ii) If $\tilde{s} > \tau_{s_D}$, then $\beta_D^*(v, \tilde{s}) = \beta_{LD}$.

β_{LD} is the unique solution to $\max_{\beta_D} G_L(\beta_D, s_D, \tilde{s})$ in Lemma 3 and β_{HD} is the unique solution to $\max_{\beta_D} G_H(\beta_D, s_D)$ in Lemma 4. Furthermore, (i) $\beta_{LD} = \beta_{HD}$ if and only if $v = 0$ and (ii) $\beta_{LD} = 0$ if and only if $v = 1$.

Proof. First, we show that the unique solution to $\max_{\beta_D} G_L(\beta_D, s_D, \tilde{s})$, which we denote as β_{LD} , is less than or equal than the unique solution to $\max_{\beta_D} G_H(\beta_D, s_D)$, which we denote as β_{HD} . For ease of notation, we write the partial derivative $\frac{\partial G_H(\beta_D, s_D)}{\partial \beta_D}$ simply as $G'_H(\beta_D)$, and the partial derivative $\frac{\partial G_L(\beta_D, s_D, \tilde{s})}{\partial \beta_D}$ simply as $G'_L(\beta_D)$, as they are both independent of s_D and \tilde{s} . From their definitions in Equations (C.14) and C.10 (Lemmas 4 and 3, respectively), for any β_D we obtain the following result:

$$\begin{aligned} G'_H(\beta_D) &= (\alpha q (r - w) \gamma'(\tilde{s}) + p q) \frac{R(\tilde{s})}{1 - \beta_D} (v + (1 - v) \Phi(\tilde{s})) - \delta'(\beta_D) \\ &\geq (\alpha q (r - w) \gamma'(\tilde{s}) + p q) \frac{R(\tilde{s})}{1 - \beta_D} (1 - v) \Phi(\tilde{s}) - \delta'(\beta_D) \\ &= G'_L(\beta_D), \end{aligned} \tag{C.18}$$

where the inequality follows from $v \geq 0$ and the fact that, except for $-\delta'(\beta_D)$, all other terms are non-negative. In particular, Equation (C.18) applies to $\beta_D = \beta_{LD}$, i.e., $G'_H(\beta_{LD}) \geq G'_L(\beta_{LD}) = 0$, where the equality follows from β_{LD} satisfying the FOC in Equation (C.10). Thus, since $G_H(\beta_D, s_D)$ is unimodal in β_D and $G'_H(\beta_{HD}) = 0$ (Equation (C.14)), it must be that $\beta_{HD} \geq \beta_{LD}$. From this analysis, it also follows that $\beta_{HD} = \beta_{LD}$ if and only if $v = 0$, as it is the only case in which $G'_H(\beta_D) = G'_L(\beta_D)$. Similarly, since \tilde{s} is strictly increasing in β_D , we also have that $\tilde{s}(\beta_{HD}) \geq \tilde{s}(\beta_{LD})$, with the inequality being strict if $v > 0$.

To show that $\beta_{LD} = 0$ if and only if $v = 1$, we know from Equation (C.10) in

Lemma 3 that β_{LD} satisfies the FOC

$$\frac{\partial G_L}{\partial \beta_D} = (\alpha q (r - w) \gamma'(\tilde{s}) + p q) \frac{R(\tilde{s})}{1 - \beta_{LD}} (1 - v) \Phi(\tilde{s}) - \delta'(\beta_{LD}) = 0,$$

where for ease of notation we have omitted the arguments of the function G_L . If $v = 1$, then this condition simplifies to $\delta'(\beta_{LD}) = 0$. By definition, for $\delta'(\beta_{LD}) = 0$, β_{LD} must be equal to 0.

Next, we define $s_L = \check{s}(\beta_{LD})$, $s_H = \check{s}(\beta_{HD})$ and $\tilde{\beta}$ as the value of β_D such that $\check{s}(\beta_D = \tilde{\beta}) = \tilde{s}$. We note that from Lemma 6 and from the definitions of $G_L(\beta_D, s_D, \tilde{s})$ and $G_H(\beta_D, s_D)$, the manufacturer's expected profit function is continuous in β_D . However, its derivative has two discontinuity points at $\beta_D = \tilde{\beta}$ and $\beta_D = \hat{\beta}_{s_D}$. In addition, since the expected profit is equal to $(1 - r + \alpha \gamma(s_D))(r - w) - \delta(\beta_D)$ if $\tilde{s} > s_D$, it is strictly decreasing in β_D for any $\beta_D > \hat{\beta}_{s_D}$. As a result and since β_{LD} (β_{HD}) is the unique local optimum when the manufacturer's expected profit is equal to $G_L(\beta_D, s_D, \tilde{s})$ ($G_H(\beta_D, s_D)$), there are only four possible solutions to the manufacturer's optimization problem given s_D and \tilde{s} : (i) $\beta_D^*(v, \tilde{s}) = \beta_{LD}$; (ii) $\beta_D^*(v, \tilde{s}) = \tilde{\beta}$; (iii) $\beta_D^*(v, \tilde{s}) = \beta_{HD}$; or (iv) $\beta_D^*(v, \tilde{s}) = \hat{\beta}_{s_D}$.

We first rule out $\tilde{\beta}$ as the manufacturer's optimal β_D . From Lemma 6 and assuming that $\tilde{s} \leq s_D$, we know that $\Pi_M(\beta_D) = G_L(\beta_D)$ if $\tilde{s} \leq \tilde{s}$ and $\Pi_M(\beta_D) = G_H(\beta_D)$ if $\tilde{s} > \tilde{s}$, where for ease of notation $\Pi_M(\beta_D)$ represents the manufacturer's expected profit. This in turn implies that $\lim_{\beta_D \downarrow \tilde{\beta}} \Pi'_M(\beta_D) = \lim_{\beta_D \downarrow \tilde{\beta}} G'_L(\beta_D)$ and that $\lim_{\beta_D \uparrow \tilde{\beta}} \Pi'_M(\beta_D) = \lim_{\beta_D \uparrow \tilde{\beta}} G'_H(\beta_D)$. In addition, from Equation (C.18) we know that $G'_H(\tilde{\beta}) \geq G'_L(\tilde{\beta})$. Thus, it follows that $\lim_{\beta_D \downarrow \tilde{\beta}} \Pi'_M(\beta_D) \geq \lim_{\beta_D \uparrow \tilde{\beta}} \Pi'_M(\beta_D)$. As a result, the manufacturer's expected profit cannot be strictly increasing immediately to the left of $\tilde{\beta}$ and strictly decreasing immediately to the right of $\tilde{\beta}$. Therefore, $\beta_D = \tilde{\beta}$ cannot be a local maximum.

Regarding the other possible optimal solutions, we note from the manufacturer's expected profit in Equation (C.15) and from the unimodality of $G_L(\beta_D, s_D, \tilde{s})$ and $G_H(\beta_D, s_D)$ that: (i) β_{LD} is a candidate solution iff $s_L \leq \tilde{s}$ and $s_L \leq s_D$; (ii) β_{HD} is a candidate solution iff $\tilde{s} < s_H \leq s_D$; and (iii) $\hat{\beta}_{s_D}$ is a candidate solution iff $\tilde{s} <$

$s_D < s_H$, or if $\tilde{s} \geq s_D$ and $s_L > s_D$. Case (iii) follows from the fact that $\tilde{s} < s_D < s_H$ guarantees that $\Pi_M(\beta_D = \hat{\beta}_{s_D}) = G_H(\beta_D = \hat{\beta}_{s_D}, s_D)$ and that $G_H(\beta_D = \hat{\beta}_{s_D}, s_D)$ is strictly increasing in β_D , as $s_H > s_D$ is equivalent to $\beta_{HD} > \hat{\beta}_{s_D}$ (\tilde{s} is strictly increasing in β_D) and G_H is unimodal in β_D . The same logic can be applied to show why $\hat{\beta}_{s_D}$ is a candidate solution in case (iii) when $\tilde{s} \geq s_D$ and $s_L > s_D$, as this guarantees that $\Pi_M(\beta_D = \hat{\beta}_{s_D}) = G_L(\beta_D = \hat{\beta}_{s_D}, s_D, \tilde{s})$ and that $G_L(\beta_D = \hat{\beta}_{s_D}, s_D, \tilde{s})$ is strictly increasing in β_D . In other words, case (iii) guarantees that the manufacturer's expected profits are left-increasing at $\beta_D = \hat{\beta}_{s_D}$, which in turn means that $\hat{\beta}_{s_D}$ is a candidate solution. Thus, we can identify the following cases:

- (i) If $s_L > s_D$, then either $\tilde{s} < s_D$ (and hence $\tilde{s} < s_D < s_H$, since $s_L \leq s_H$) or $\tilde{s} \geq s_D$. In both cases, the only solution is $\beta_D^*(v, \tilde{s}) = \hat{\beta}_{s_D}$ for any \tilde{s} .
- (ii) If $s_L \leq s_D$ and $s_L > \tilde{s}$, then β_{LD} is not a candidate solution. In this case, the only candidate solution is β_{HD} , if $s_H \leq s_D$; or $\hat{\beta}_{s_D}$, if $s_H > s_D$. In other words, $\beta_D^*(v, \tilde{s}) = \min\{\beta_{HD}, \hat{\beta}_{s_D}\}$.
- (iii) If $s_L \leq s_D$ and $\tilde{s} > \min\{s_H, s_D\}$, then neither β_{HD} nor $\hat{\beta}_{s_D}$ are candidate solutions. Therefore, $\beta_D^*(v, \tilde{s}) = \beta_{LD}$.
- (iv) If $s_L \leq s_D$ and $s_L \leq \tilde{s} < \min\{s_H, s_D\}$, then both β_{LD} and $\min\{\beta_{HD}, \hat{\beta}_{s_D}\}$ are candidate solutions.

Therefore, only in case (iv) when $s_L \leq s_D$ and $\tilde{s} \in (s_L, \min\{s_H, s_D\}]$ do we have two possible solutions to the manufacturer's optimal choice of β_D . In what follows, we assume that $s_L \leq s_D$ and identify for what values of \tilde{s} the manufacturer prefers $\beta_D = \beta_{LD}$ and for what values of \tilde{s} she instead prefers $\beta_D = \bar{\beta}_{HD} \equiv \min\{s_H, s_D\}$.

First, we note from Lemma 6 that in this range of signals, $\Pi_M(\beta_D = \beta_{LD}) = G_L(\beta_{LD}, s_D, \tilde{s})$ and $\Pi_M(\beta_D = \bar{\beta}_{HD}) = G_H(\bar{\beta}_{HD}, s_D)$. This, as the condition $\check{s} \leq \tilde{s}$ is satisfied only by $s_L = \check{s}(\beta_{LD})$, while the condition $\check{s} > \tilde{s}$ is satisfied only by $s_H = \check{s}(\bar{\beta}_{HD})$. In addition, from the definition of $G_L(\beta_D, s_D, \tilde{s})$ in Lemma 3 we know that if $s_L \leq s_D$, then $G'_L(\beta_{LD}) = \alpha q v \gamma'(\tilde{s})(r - w) + p q v > 0$; i.e., $G_L(\beta_{LD}, s_D, \tilde{s})$ is strictly

increasing in the signal \tilde{s} . Conversely, from Lemma 4 we know that $G_H(\beta_D, s_D)$ is independent of the signal \tilde{s} for any β_D , and thus for $\beta_D = \bar{\beta}_{HD}$ in particular.

Thus, combining the results from the above paragraph we obtain that $\Pi_M(\beta_{LD})$ is continuous and strictly increasing in the signal \tilde{s} in the range $[s_L, \min\{s_H, s_D\}]$; and $\Pi_M(\bar{\beta}_{HD})$ is independent of the signal \tilde{s} . In addition, we also know from case (iii) above that if $\tilde{s} = \min\{s_H, s_D\}$, then the manufacturer strictly prefers β_{LD} over $\bar{\beta}_{HD}$. Therefore, there must exist a unique threshold $\tau_{s_D} \in [s_L, \min\{s_H, s_D\}]$ such that (i) if $\tilde{s} < \tau_{s_D}$, then $\Pi_M(\bar{\beta}_{HD}) > \Pi_M(\beta_{LD})$; (ii) if $\tilde{s} = \tau_{s_D}$, then $\Pi_M(\bar{\beta}_{HD}) = \Pi_M(\beta_{LD})$; and (iii) if $\tilde{s} > \tau_{s_D}$, then $\Pi_M(\beta_{HD}) < \Pi_M(\beta_{LD})$.

To summarize, and assuming that the manufacturer selects $\beta_D = \bar{\beta}_{HD}$ when she is indifferent between β_{LD} and $\bar{\beta}_{HD}$, we can identify the following two scenarios: (1) If $\beta_{LD} \geq \hat{\beta}_{s_D}$, then $\forall \tilde{s}$, $\beta_D^*(v, \tilde{s}) = \hat{\beta}_{s_D}$; (2) If $\beta_{LD} < \hat{\beta}_{s_D}$, then (i) if $\tilde{s} \leq \tau_{s_D}$, then $\beta_D^*(v, \tilde{s}) = \beta_{LD}$; or (ii) if $\tilde{s} > \tau_{s_D}$, then $\beta_D^*(v, \tilde{s}) = \bar{\beta}_{HD}$. Case (1) and the condition $\beta_{LD} < \hat{\beta}_{s_D}$ for case (2) follow from the fact that \check{s} is strictly increasing in β_D and thus, from the definition of $\hat{\beta}_{s_D}$ as the value of β_D such that $\check{s}(\beta_D) = s_D$, the condition $s_L > s_D$ is equivalent to $\beta_{LD} > \hat{\beta}_{s_D}$. Similarly, if $s_L = s_D$, then the manufacturer also selects either $\beta_{LD} = \hat{\beta}_D$ or $\bar{\beta}_{HD} = \hat{\beta}_D$ for any signal \tilde{s} , which explains why the inequality in case (1) is not strict. \square

Note: No Manufacturer Disclosure setting. If we replace s_D by \hat{s} , β_D by β , $\hat{\beta}_{s_D}$ by $\hat{\beta}$, and if we consider $p = 0$, then the above proof demonstrates Theorem 2.

Corollary 4. *Given any s_D , $\beta_D^*(s_D)$ is always such that $\check{s}(\beta_D^*(s_D)) \leq s_D$. By the definition of $\hat{\beta}_{s_D}$ in Proposition 13, this means that $\beta_D^*(s_D) \leq \hat{\beta}_{s_D}$. Thus, the supplier's best response to $(\beta_D = \beta_D^*(s_D), s_D)$ is equal to $s^*(s_0, \beta_D^*(s_D), s_D) = s^*(s_0, \beta_D^*(s_D)) = \max\{s_0, \check{s}(\beta_D^*(s_D))\}$.*

Proof. The manufacturer's selection of $\beta_D^*(s_D) \leq \hat{\beta}_{s_D}$ follows directly from Proposition 13. Similarly, $s^*(s_0, \beta_D^*(s_D), s_D) = \max\{s_0, \check{s}(\beta_D^*(s_D))\}$ is a direct consequence of $s^*(s_0, \beta_D, s_D) = \max\{s_0, \min\{s_D, \check{s}(\beta_D)\}\}$ (see Appendix C.1) and the fact that $\check{s}(\beta_D^*(s_D)) \leq s_D$. \square

Proposition 14. *Let us define*

$$G(\beta_D, s_D, \tilde{s}) \equiv G_L(\beta_D, s_D, \tilde{s}) \mathbb{1}_{\tilde{s} \leq s_D} + G_H(\beta_D, s_D) \mathbb{1}_{\tilde{s} > s_D},$$

which is equal to the manufacturer's expected profits when $\tilde{s} \leq s_D$. The function $G(\beta_D, s_D, \tilde{s})$ is unimodal in s_D and there exist $m < \hat{\tau}_L \leq \hat{\tau}_H \leq M$ such that $z(\tilde{s}) \equiv \arg \max_{s_D} G(\beta_D, s_D, \tilde{s})$ is independent of β_D and can be characterized as follows:

(i) *If $\tilde{s} < \hat{\tau}_L$, then $z(\tilde{s}) = \hat{\tau}_L$;*

(ii) *If $\hat{\tau}_L \leq \tilde{s} < \hat{\tau}_H$, then $z(\tilde{s}) = \tilde{s}$;*

(iii) *If $\tilde{s} \geq \hat{\tau}_H$, then $z(\tilde{s}) = \hat{\tau}_H$.*

Furthermore, $\hat{\tau}_H > \hat{\tau}_L$ and they are defined as the unique solutions to

$$V_1(s_D = \hat{\tau}_H) \equiv \alpha(r - w) \gamma'(\hat{\tau}_H) (1 - q(1 - v) \Phi(\hat{\tau}_H)) - pq(1 - v) \Phi(\hat{\tau}_H) = 0 \quad (\text{C.19})$$

$$V_2(s_D = \hat{\tau}_L) \equiv \alpha(r - w) \gamma'(\hat{\tau}_L) (1 - q \Phi(\hat{\tau}_L)) - pq(v + (1 - v) \Phi(\hat{\tau}_L)) = 0. \quad (\text{C.20})$$

Proof. From the definitions of the functions G_L and G_H in Lemmas 3 and 4, and using Leibniz's rule we can write the partial derivative of G with respect to s_D as follows:

$$\begin{aligned} \frac{\partial G(\beta_D, s_D, \tilde{s})}{\partial s_D} &= \left(\alpha(1 - q) \gamma'(s_D) + \alpha q \left(v \gamma'(s_D) \mathbb{1}_{s_D \leq \tilde{s}} + (1 - v) \left(\gamma(s_D) \phi(s_D) \right. \right. \right. \\ &\quad \left. \left. \left. + \gamma'(s_D) (1 - \Phi(s_D)) - \gamma(s_D) \phi(s_D) \right) \right) \right) (r - w) \\ &\quad - pq \left(v \mathbb{1}_{s_D > \tilde{s}} + (1 - v) (\Phi(\tilde{s}) + \int_{\tilde{s}}^{s_D} \phi(x) dx) \right), \end{aligned}$$

or equivalently,

$$\begin{aligned} \frac{\partial G(\beta_D, s_D, \tilde{s})}{\partial s_D} &= \left(\alpha(1-q)\gamma'(s_D) + \alpha q \left(v\gamma'(s_D)\mathbb{1}_{s_D \leq \tilde{s}} + (1-v)\gamma'(s_D)(1-\Phi(s_D)) \right) \right) \\ &\quad \times (r-w) - pq(v\mathbb{1}_{s_D > \tilde{s}} + (1-v)\Phi(s_D)). \end{aligned} \quad (\text{C.21})$$

Equation (C.21) is independent of β_D and thus, s_D^* is also independent of β_D and only depends on \tilde{s} .

Next, to simplify our analysis of how $G(\beta_D, s_D, \tilde{s})$ changes with s_D , we rewrite Equation (C.21) as two separate components for when $s_D \leq \tilde{s}$ and when $s_D > \tilde{s}$. In particular, we define

$$\begin{aligned} V_1(s_D) &\equiv \left(\alpha(1-q)\gamma'(s_D) + \alpha q \left(v\gamma'(s_D) + (1-v)\gamma'(s_D)(1-\Phi(s_D)) \right) \right) (r-w) \\ &\quad - pq(1-v)\Phi(s_D) \\ &= \alpha(r-w)\gamma'(s_D)(1-q(1-v)\Phi(s_D)) - pq(1-v)\Phi(s_D), \end{aligned} \quad (\text{C.22})$$

and similarly,

$$\begin{aligned} V_2(s_D) &\equiv \left(\alpha(1-q)\gamma'(s_D) + \alpha q(1-v)\gamma'(s_D)(1-\Phi(s_D)) \right) (r-w) \\ &\quad - pq(v + (1-v)\Phi(s_D)) \\ &= \alpha(r-w)\gamma'(s_D)(1-q\Phi(s_D)) - pq(v + (1-v)\Phi(s_D)). \end{aligned} \quad (\text{C.23})$$

With these definitions, Equation (C.21) can be rewritten as

$$\frac{\partial G(\beta_D, s_D, \tilde{s})}{\partial s_D} = V_1(s_D)\mathbb{1}_{s_D \leq \tilde{s}} + V_2(s_D)\mathbb{1}_{s_D > \tilde{s}}. \quad (\text{C.24})$$

We first note that both $V_1(s_D)$ and $V_2(s_D)$ are strictly decreasing in s_D . In par-

ticular,

$$\begin{aligned} \frac{dV_1(s_D)}{ds_D} &= \alpha(r-w)\gamma''(s_D)(1-q(1-v)\Phi(s_D)) - \alpha(r-w)\gamma'(s_D)q(1-v)\phi(s_D) \\ &\quad - pq(1-v)\phi(s_D) < 0, \end{aligned}$$

where the last inequality follows from $\gamma(s)$ being strictly increasing and strictly concave in s . Similarly,

$$\begin{aligned} \frac{dV_2(s_D)}{ds_D} &= \alpha(r-w)\gamma''(s_D)(1-q\Phi(s_D)) - \alpha(r-w)\gamma'(s_D)q\phi(s_D) - pq(1-v)\phi(s_D) \\ &< 0. \end{aligned}$$

In addition, since $\Phi(m) = 0$, $V_1(m) > 0$ and $\lim_{s_D \rightarrow -\infty} V_2(s_D) > 0$. Similarly, since $\gamma(s)$ is strictly increasing in s and bounded above by r , it follows that $\lim_{s \rightarrow \infty} \gamma'(s) = 0$ and hence $\lim_{s_D \rightarrow \infty} V_1(s_D) < 0$ and $\lim_{s_D \rightarrow \infty} V_2(s_D) < 0$. Thus, there exists a unique $s_D = \hat{\tau}_H$ that satisfies $V_1(s_D) = 0$ and a unique $s_D = \hat{\tau}_L$ that satisfies $V_2(s_D) = 0$. In addition, by comparing Equations (C.22) and (C.23) we know that for any s_D , $V_1(s_D) \geq V_2(s_D)$ and hence $\hat{\tau}_H \geq \hat{\tau}_L$. Since $V_1(s_D) = V_2(s_D)$ only if $v = 0$, it follows that $\hat{\tau}_L = \hat{\tau}_H$ if and only if $v = 0$. Thus, for any $v > 0$, $\hat{\tau}_H > \hat{\tau}_L$.

In addition, the above also means that for any given signal \tilde{s} , $G(\beta_D, s_D, \tilde{s})$ is strictly concave in s_D and independent of β_D . Similarly, the derivative of $G(\beta_D, s_D, \tilde{s})$ with respect to s_D has a discontinuity at $s_D = \tilde{s}$, where $\lim_{s_D \uparrow \tilde{s}} \frac{\partial G(\beta_D, s_D, \tilde{s})}{\partial s_D} > \lim_{s_D \downarrow \tilde{s}} \frac{\partial G(\beta_D, s_D, \tilde{s})}{\partial s_D}$. Therefore, using our results regarding the functions $V_1(s_D)$ and $V_2(s_D)$ above and Equation (C.24), given any \tilde{s} the supplier's choice of s_D that maximizes $G(\beta_D, s_D, \tilde{s})$ is as follows:

- (i) If $\hat{\tau}_H \leq \tilde{s}$, then $z(\tilde{s}) = \hat{\tau}_H$;
- (ii) If $\hat{\tau}_L > \tilde{s}$, then $z(\tilde{s}) = \hat{\tau}_L$;
- (iii) Otherwise, if $\hat{\tau}_L < \tilde{s} \leq \hat{\tau}_H$, then $z(\tilde{s}) = \tilde{s}$. □

Lemma 7. *There exists at most one pair of values (β_D, s_D) that satisfy $s_D = \check{s}(\beta_D)$ and that represent a local maximum of the manufacturer's expected profits. That is, they satisfy the first order condition*

$$\alpha \gamma'(\check{s}) (r - w) \frac{R(\check{s})}{1 - \beta_D} - \delta'(\beta_D) = 0. \quad (\text{C.25})$$

Furthermore, both of these values are independent of the level of visibility v , the signal \tilde{s} , and the penalty p . We define $(\hat{\beta}_D, \hat{s}_D)$ as this pair of values, if they exist; and as (β_M, M) otherwise.

Proof. First, if no such pair of values exist, then Lemma 7 is trivially true. Thus, we assume that such a pair (β_D, s_D) exists. By Lemma 6 and from the continuity of the manufacturer's expected profit function, at any point $(\beta_D, s_D = \check{s}(\beta_D))$ we know that

$$\mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \tilde{s}] = (1 - r + \alpha \gamma(s_D)) (r - w) - \delta(\beta_D).$$

We know that $s_D = \check{s}(\beta_D)$, so we can rewrite the above expression exclusively as a function of β_D :

$$\mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D = \check{s}(\beta_D)) | \tilde{s}] = (1 - r + \alpha \gamma(\check{s}(\beta_D))) (r - w) - \delta(\beta_D).$$

For ease of notation, we hereafter write the manufacturer's expected profit function simply as $\Pi_M(\beta_D)$, and $\check{s}(\beta_D)$ as \check{s} .

From the above discussion, we note that the FOC of the manufacturer's profits with respect to β_D is given by

$$\Pi'_M(\beta_D) = \alpha \gamma'(\check{s}) (r - w) \frac{d\check{s}}{d\beta_D} - \delta'(\beta_D) = 0,$$

or equivalently

$$\Pi'_M(\beta_D) = \alpha \gamma'(\check{s}) (r - w) \frac{R(\check{s})}{1 - \beta_D} - \delta'(\beta_D) = 0, \quad (\text{C.26})$$

where we use the fact that $\frac{d\check{s}}{d\beta_D} = \frac{R(\check{s})}{1-\beta_D}$ (see Lemma 2). Thus, the Equation (C.26) is a necessary condition for β_D to maximize the manufacturer's expected profits. Since \check{s} is independent of v , \tilde{s} , and p (see Lemma 1), the above expression implies that any β_D that satisfies the above equation is also independent of these values.

Next, we show that the solution to the FOC in Equation (C.26) is unique. By taking the second derivative of the manufacturer's profits, we obtain

$$\begin{aligned} \Pi''_M(\beta_D) &= \left(\alpha \gamma''(\check{s}) (r-w) \frac{R(\check{s})}{1-\beta_D} + \alpha \gamma'(\check{s}) (r-w) \frac{R'(\check{s})}{1-\beta_D} \right) \frac{d\check{s}}{d\beta_D} \\ &\quad + \alpha \gamma'(\check{s}) (r-w) \frac{R(\check{s})}{(1-\beta_D)^2} - \delta''(\beta_D). \end{aligned}$$

Since $\gamma(s)$ is strictly increasing and strictly concave in s , $R(s)$ is positive and non-increasing in s (see Assumption 4), and \check{s} is strictly increasing in β_D , we know that the first term on the RHS is strictly negative. Thus, the second derivative of the manufacturer's profits satisfies

$$\Pi''_M(\beta_D) < \alpha \gamma'(\check{s}) (r-w) \frac{R(\check{s})}{(1-\beta_D)^2} - \delta''(\beta_D).$$

From the FOC, $\alpha \gamma'(\check{s}) (r-w) \frac{R(\check{s})}{(1-\beta_D)} = \delta'(\beta_D)$ and so the above expression simplifies to

$$\Pi''_M(\beta_D) < \frac{\delta'(\beta_D)}{(1-\beta_D)} - \delta''(\beta_D) < \frac{\delta'(\beta_D)}{1-\beta_D} \left(1 - \frac{\delta''(\beta_D)}{\delta'(\beta_D)} (1-\beta_D) \right) < 0,$$

where the last inequality follows directly from Assumption 5. Therefore, any β_D that satisfies the FOC also satisfies $\Pi''_M(\beta_D) < 0$. Therefore, since $\Pi_M(\beta_D)$ is continuous in β_D , it follows that if such a β_D exists, then it must be unique. Since we assume that $s_D = \check{s}(\beta_D)$ and $\check{s}(\beta_D)$ is strictly increasing in β_D (see Lemma 1), it follows that such a s_D is also unique. Thus, if there exists (β_D, s_D) that satisfy $s_D = \check{s}(\beta_D)$ and that represent a local maximum of the manufacturer's expected profits, then such a pair of values must be unique. \square

Lemma 8. *There are at most three candidate solutions to the manufacturer's problem*

in Equation (3.9):

(i) If $\check{s}(\beta_{LD}) \leq \tilde{s}$ and $\check{s}(\beta_{LD}) < z(\tilde{s})$, then $(\beta_D, s_D) = (\beta_{LD}, \min\{z(\tilde{s}), M\})$ is a candidate solution;

(ii) If $\tilde{s} < \check{s}(\beta_{HD}) < z(\tilde{s})$, then $(\beta_D, s_D) = (\beta_{HD}, \min\{z(\tilde{s}), M\})$ is a candidate solution;

(iii) If (a) $z(\tilde{s}) \leq \check{s}(\beta_{LD}) \leq \tilde{s}$ or (b) $\check{s}(\beta_{HD}) > \tilde{s}$ and $\check{s}(\beta_{HD}) \geq z(\tilde{s})$ then $(\beta_D, s_D) = (\hat{\beta}_D, \hat{s}_D)$ is a candidate solution,

where $z(\tilde{s})$ is the unique optimal solution to $\max_{s_D} G_L(\beta_D, s_D, \tilde{s}) \mathbb{1}_{\tilde{s} \leq \tilde{s}} + G_H(\beta_D, s_D) \mathbb{1}_{\tilde{s} > \tilde{s}}$ defined in Proposition 14; and β_{LD} and β_{HD} are the unique optimal solutions to $\max_{\beta_D} G_L(\beta_D, s_D, \tilde{s})$ and $\max_{\beta_D} G_H(\beta_D, s_D)$.

Proof. From the structure of the manufacturer's expected profits in Lemma 6, it is easy to see that any optimal solution must satisfy $\check{s}(\beta_D) \leq s_D$; any strategy where this is not satisfied is strictly dominated by reducing β_D to the point where $\check{s}(\beta_D) = s_D$. Therefore, in this proof we first study the case $\check{s}(\beta_D) < s_D$ and then we study the case $\check{s}(\beta_D) = s_D$.

We first note that any solution (β_D, s_D) that satisfies $\check{s}(\beta_D) < s_D$ must also satisfy $s_D = \min\{z(\tilde{s}), M\}$ and β_D equal to either β_{LD} or β_{HD} . This follows directly from the structure of the manufacturer's expected profits and from the fact that (i) if $\check{s}(\beta_D) < s_D$, then by Proposition 14, $G(\beta_D, s_D, \tilde{s})$ is equal to the manufacturer's profits; (ii) since $G(\beta_D, s_D, \tilde{s})$ is unimodal in s_D and $z(\tilde{s})$ is its unique optimal solution and is independent of β_D , then s_D must be equal to $\min\{z(\tilde{s}), M\}$ (since the manufacturer selects $s_D \in [m, M]$); and (iii) similarly, by the manufacturer's optimal choice of β_D given a fixed s_D in Theorem 3, $\check{s}(\beta_D) < s_D$ means that $\beta_D < \hat{\beta}_D$ and thus it must satisfy either $\beta_D = \beta_{LD}$ or $\beta_D = \beta_{HD}$, both of which are independent of s_D (see Lemmas 3 and 4).

Therefore, by the structure of the manufacturer's profits in Lemma 6 and the unimodality of the functions G_L and G_H in β_D , $(\beta_{LD}, z(\tilde{s}))$ is a candidate solution if $\check{s}(\beta_{LD}) \leq \tilde{s}$ and $\check{s}(\beta_{LD}) < z(\tilde{s})$ (i.e., part (i) of Lemma 8); and $(\beta_{HD}, z(\tilde{s}))$ is a

candidate solution if $\check{s}(\beta_{HD}) > \tilde{s}$ and $\check{s}(\beta_{HD}) < z(\tilde{s})$ (i.e., part (ii) of Lemma 8). These conditions guarantee that the manufacturer's expected profits indeed have a local optimum at $\beta_D = \beta_{LD}$ and $\beta_D = \beta_{HD}$, respectively. For example, if $\check{s}(\beta_{LD}) \leq \tilde{s}$ and $\check{s}(\beta_{LD}) < z(\tilde{s})$, then $\mathbb{E}_{s_0}[\Pi_M(\beta_D, s_D) | \tilde{s}] = G_L(\beta_{LD}, s_D, \tilde{s})$ and hence β_{LD} is a local optimum not only of the function G_L , but of the manufacturer's profits.

Finally, if a candidate solution is of the form $\check{s}(\beta_D) = s_D$, then we know from Lemma 7 that it must satisfy $(\beta_D, s_D) = (\hat{\beta}_D, \hat{s}_D)$. In particular, there are two cases in which the optimal choice of β_D given the supplier's choice of $s_D = z(\tilde{s})$ leads to $\check{s}(\beta_D) = z(\tilde{s})$: (i) if $\check{s}(\beta_{LD}) \leq \tilde{s}$ and $\check{s}(\beta_{LD}) \geq z(\tilde{s})$, then from the structure of the manufacturer's expected profits in Lemma 6 and the unimodality of G_L in β_D , the manufacturer's expected profits are increasing immediately to the left of β_D such that $\check{s}(\beta_D) = s_D = z(\tilde{s})$; or (ii) similarly, if $\check{s}(\beta_{HD}) > \tilde{s}$ and $\check{s}(\beta_{HD}) \geq z(\tilde{s})$, then the manufacturer's expected profits are also increasing immediately to the left of β_D such that $\check{s}(\beta_D) = s_D = z(\tilde{s})$. Since the manufacturer's profits are strictly decreasing in β_D immediately to the right of such critical value, it follows that in any of these two cases, $(\beta_D, s_D = z(\tilde{s}))$ is a local optimum. Thus, if any of these pairs of conditions is satisfied, then $(\beta_D, s_D) = (\hat{\beta}_D, \hat{s}_D)$ is a candidate solution. \square

Lemma 9. *Given any signal \tilde{s} , any candidate solution (β_D, s_D) to the manufacturer's optimization problem must satisfy $\beta_D \leq \hat{\beta}_D$, where $\hat{\beta}_D$ is defined in Lemma 7. Similarly, if a candidate solution is such that $\beta_D = \hat{\beta}_D$, then it must satisfy $s_D = \hat{s}_D$.*

Proof. First, we note that if $\hat{\beta}_D = \beta_M$, then the result is trivially satisfied. If instead $(\hat{\beta}_D, \hat{s}_D)$ is defined as a local optimum, then we prove this result by contradiction. Suppose that (β_D, s_D) is a local maximum (i.e., it is a candidate solution to the manufacturer's optimization problem) and that $\beta_D > \hat{\beta}_D$. From Lemma 8, we know that in this case, either $\beta_D = \beta_{LD}$ or $\beta_D = \beta_{HD}$. For the remainder of this proof, we consider the first case, as the proof is the same in both cases.

Also from Lemma 8, $\beta_D = \beta_{LD}$ being a candidate solution implies that $\hat{s}(\beta_{LD}) < z(\tilde{s})$. Because we assume that $\beta_D > \hat{\beta}_D$, this also implies that $\hat{s}_D = \check{s}(\hat{\beta}_D) < z(\tilde{s})$.

When $s_D > \check{s}(\beta_D)$, the manufacturer's expected profits are equal to the function $G(\beta_D, s_D, \tilde{s})$ (see Proposition 14). Given that the function G is unimodal in s_D and maximized at $s_D = z(\tilde{s})$, then we know that the manufacturer's expected profits when selecting $(\hat{\beta}_D, z(\tilde{s}))$ are strictly greater than when selecting $(\hat{\beta}_D, \hat{s}_D)$. Thus, the latter pair of values do not constitute a local optimum, which is a contradiction.

Finally, the same logic can be used to show that if a local maximum satisfies $\beta_D = \hat{\beta}_D$, then $s_D = \hat{s}_D$. From the structure of the manufacturer's optimal profits in Lemma 6, a local maximum cannot be such that $s_D < \check{s}(\hat{\beta}_D)$, which by definition satisfies $\check{s}(\hat{\beta}_D) = \hat{s}_D$. Conversely, if $s_D > \hat{s}_D$ and it is a local maximum, then by the same argument as in the above paragraph, $(\hat{\beta}_D, \hat{s}_D)$ would not be a local maximum, which contradicts its definition if $\hat{\beta}_D < \beta_M$. Finally, if $\hat{\beta}_D = \beta_M$, then $\check{s}(\hat{\beta}_D) = M$ and hence $s_D = \hat{s}_D$ is the only value of $s_D \in [m, M]$ that satisfies $s_D \geq \check{s}(\hat{\beta}_D)$. \square

Proof of Theorems 3 and 6. Next, we demonstrate the manufacturer's optimal strategy, (β_D^*, s_D^*) , for any given signal \tilde{s} and visibility v . For the remainder of this proof, we define the threshold $\tau_H \equiv \min\{\hat{\tau}_H, M\}$, where $\hat{\tau}_H$ is defined in Proposition 14.

First, from Lemma 9, we know that if $\beta_{LD} \geq \hat{\beta}_D$, then this also implies that $\beta_{HD} > \hat{\beta}_D$ and thus, $\beta_D^* = \hat{\beta}_D$ for any signal. Similarly, this means that $s_D^* = \hat{s}_D = \check{s}(\hat{\beta}_D)$ for any signal. From Assumption 2 and the definition of $\check{s}(\beta_D)$ in Lemma 1, it follows that $s^*(s_0 = m, \beta_D = \hat{\beta}_D) = \check{s}(\hat{\beta}_D)$, which completes the proof of case (a) in Theorems 3 and 6.

Second, by Proposition 14, if $\beta_{LD} < \hat{\beta}_D$ and $\hat{\tau}_L \geq M$, then for any signal \tilde{s} , we know that $z(\tilde{s}) \geq M$. From Assumption 3, we know that any β_D that is optimal for the manufacturer satisfies $\check{s}(\beta_D) < M$ and hence $\check{s}(\beta_{LD}) \leq \check{s}(\beta_{HD}) < M$. Thus, by Lemma 8 and since $\hat{\tau}_H \geq \hat{\tau}_L$, the only candidate solution for disclosure is $s_D = \min\{z(\tilde{s}), M\} = M$, regardless of the signal \tilde{s} . In addition, from the manufacturer's optimal choice of β_D given s_D in Proposition 13, for any s_D that satisfies $s_D \geq \check{s}(\beta_{HD})$ (or equivalently, $\hat{\beta}_{s_D} \geq \beta_{HD}$), we know that the structure of the optimal β_D is equal to $\beta_D^* = \beta_{HD}$ up to a certain signal-threshold τ_D ; and equal to $\beta_D^* = \beta_{LD}$ thereafter.

Using the definition of $\hat{\tau}_L$ in Equation (C.20) as well as the fact that the function $G(\beta_D, s_D, \tilde{s})$ in Proposition 14 is unimodal in s_D , it follows that $\hat{\tau}_L \geq M$ if and only if $\alpha(1-q)(r-w)\gamma'(M) \geq pq$. This completes the proof of case (b) in Theorem 3.

Third, if $\beta_{LD} < \hat{\beta}_D$ and $\hat{\tau}_L < M$ (or equivalently, using the above result, if $\alpha(1-q)(r-w)\gamma'(M) < pq$), using the definitions and properties in Proposition 14 we identify two additional cases. First, if $\check{s}(\beta_{HD}) < \hat{\tau}_L$, then we also have that $\check{s}(\beta_{LD}) < \hat{\tau}_L$. In addition, since $z(\tilde{s})$ is non-decreasing in \tilde{s} (see Proposition 14), it follows that $\check{s}(\beta_{LD}) \leq \check{s}(\beta_{HD}) < z(\tilde{s})$ for every signal \tilde{s} . Thus, by Lemma 8 it follows that for any signal \tilde{s} , the optimal disclosure strategy is $s_D^* = z(\tilde{s})$. In particular, s_D^* is constant and equal to $\hat{\tau}_L$ for $\tilde{s} \leq \hat{\tau}_L$. Similar to above, for a given s_D such that $s_D \geq \check{s}(\beta_{HD})$, the manufacturer's optimal choice of β_D is given by the structure in Proposition 13; i.e., it is equal to $\beta_D^* = \beta_{HD}$ for $\tilde{s} \leq \tau_D$; and equal to $\beta_D^* = \beta_{LD}$ for $\tilde{s} > \tau_D$. At $\tilde{s} = \tau_D$, the manufacturer is indifferent between the strategies $(\beta_{HD}, z(\tau_D))$ and $(\beta_{LD}, z(\tau_D))$, but for simplicity we assume that she selects the former. We also note that at $\tilde{s} = \hat{\tau}_L$, we cannot simultaneously have $\check{s}(\beta_{HD}) < \hat{\tau}_L = \tilde{s}$ and $\check{s}(\beta_{HD}) > \tilde{s}$. Thus, by Lemma 8, $(\beta_{HD}, z(\tilde{s}) = \hat{\tau}_L)$ cannot be a candidate solution and hence at $\tilde{s} = \hat{\tau}_L$, $(\beta_D^*, s_D^*) = (\beta_{LD}, \hat{\tau}_L)$. We can further define the threshold τ_L as $\tau_L \equiv \hat{\tau}_L$ such that $\tau_D < \tau_L$.

Next, we study the case where $\beta_{LD} < \hat{\beta}_D$, $\hat{\tau}_L < M$, and $\check{s}(\beta_{HD}) \geq \hat{\tau}_L$. By Lemma 8, we know that the last inequality implies that $(\beta_{HD}, z(\tilde{s}))$ is not a candidate solution in this case. In addition, if $\tilde{s} < \check{s}(\beta_{LD})$, then the only possible solution is $(\hat{\beta}_D, \hat{s}_D)$. That is, for very low signals, $(\beta_D^*, s_D^*) = (\hat{\beta}_D, \hat{s}_D)$. Conversely, for any signal such that $\tilde{s} \geq \check{s}(\beta_{LD})$, we know that either $(\hat{\beta}_D, \hat{s}_D)$ is the unique candidate solution (if $\check{s}(\beta_{LD}) \geq z(\tilde{s})$), or that we need to compare the alternatives $(\hat{\beta}_D, \hat{s}_D)$ and $(\beta_{LD}, z(\tilde{s}))$.

To compare the last two alternatives, we note the following. First, from the definitions of $\hat{\beta}_D$ and \hat{s}_D in Lemma 7 and the structure of the manufacturer's profits in Lemma 6, we know that

$$\mathbb{E}_{s_0} [\Pi_M(\hat{\beta}_D, \hat{s}_D) | \tilde{s}] = (1 - r + \alpha\gamma(\hat{s}_D))(r - w) - \delta(\hat{\beta}_D),$$

which is independent of \tilde{s} . Conversely, we can define the strategy $(\beta_D = \beta_{LD}, s_D = \tilde{s})$

such that for any signal $\tilde{s} \geq \check{s}(\beta_{LD})$,

$$\mathbb{E}_{s_0}[\Pi_M(\beta_D = \beta_{LD}, s_D = \tilde{s}) | \tilde{s}] = G_L(\beta_D = \beta_{LD}, s_D = \tilde{s}, \tilde{s}),$$

which from its definition in Lemma 3 is continuous in \tilde{s} and satisfies

$$\frac{dG_L(\beta_{LD}, \tilde{s}, \tilde{s})}{d\tilde{s}} = \alpha q v (r - w) \gamma'(\tilde{s}) + p q v > 0.$$

That is, the manufacturer's expected profits are constant in \tilde{s} when she selects the strategy $(\hat{\beta}_D, \hat{s}_D)$, and strictly increasing in \tilde{s} when she selects instead the strategy (β_{LD}, \tilde{s}) . Furthermore, at $\tilde{s} = \hat{s}_D$ by Lemma 6 we know that the expected profits are equal to $G_L(\beta_{LD}, s_D = \tilde{s}, \tilde{s})$ for any β_D that satisfies $\check{s}(\beta_D) \leq s_D$, which includes $\check{s}(\hat{\beta}_D) = \hat{s}_D = \tilde{s}$. Thus, since β_{LD} is the unique β_D that maximizes $G_L(\beta_D, s_D, \tilde{s})$ for any s_D and \tilde{s} , it follows that

$$G_L(\beta_D = \beta_{LD}, s_D = \hat{s}_D, \tilde{s} = \hat{s}_D) > (1 - r + \alpha \gamma(\hat{s}_D)) (r - w) - \delta(\hat{\beta}_D).$$

As a result, at $\tilde{s} = \hat{s}_D$ the strategy (β_{LD}, \tilde{s}) strictly dominates $(\hat{\beta}_D, \hat{s}_D)$. Thus, $(\hat{\beta}_D, \hat{s}_D)$ cannot be the optimal solution and since we have already ruled out $(\beta_{HD}, z(\tilde{s}))$ as a candidate solution, it follows that at $\tilde{s} = \hat{s}_D$, $(\beta_D^*, s_D^*) = (\beta_{LD}, z(\tilde{s}))$. Similarly, from the continuity of the profit function in \tilde{s} and since the manufacturer's expected profits are increasing in \tilde{s} when she selects $(\beta_{LD}, z(\tilde{s}))$, there must exist a threshold signal, $\tau_D < \hat{s}_D$, such that the manufacturer prefers $(\hat{\beta}_D, \hat{s}_D)$ for $\tilde{s} \leq \tau_D$ and $(\beta_{LD}, z(\tilde{s}))$ for $\tilde{s} > \tau_D$. At $\tilde{s} = \tau_D$, the manufacturer is indifferent between both strategies, but for simplicity we assume that she selects $(\hat{\beta}_D, \hat{s}_D)$.

Similarly, in this case (i.e., if $\beta_{LD} < \hat{\beta}_D$, $\hat{\tau}_L < M$, and $\check{s}(\beta_{HD}) \geq \hat{\tau}_L$), we can define $\tau_L \equiv \max\{\tau_D, \hat{\tau}_L\}$, such that (i) if $\tau_D < \hat{\tau}_L$, then $\tau_L = \hat{\tau}_L$ and the manufacturer's optimal strategy is equal to $(\hat{\beta}_D, \hat{s}_D)$ for every signal $\tilde{s} \leq \tau_D$ and equal to $(\beta_{LD}, z(\tilde{s}) = \tau_L)$ for every $\tilde{s} \in (\tau_D, \tau_L]$; and (ii) if $\tau_D \geq \hat{\tau}_L$, then $\tau_L = \tau_D$ and the manufacturer's optimal strategy is equal to $(\hat{\beta}_D, \hat{s}_D)$ for every signal $\tilde{s} \leq \tau_L$. Thus, by definition $\tau_D \leq \tau_L$ and the manufacturer's optimal disclosure strategy s_D^* is non-increasing and

piece-wise constant for $\tilde{s} \in [m, \tau_L]$.

Finally, using the definitions and properties in Proposition 14, we note that for $\tilde{s} \geq \hat{\tau}_H$, $z(\tilde{s}) = \hat{\tau}_H$ and thus the manufacturer's expected profits, equal to $G_L(\beta_D, s_D = \hat{\tau}_H, \tilde{s})$, are constant in \tilde{s} . As a result, the threshold τ_D must satisfy $\tau_D < \hat{\tau}_H$. Also by Proposition 14, $\hat{\tau}_L < \hat{\tau}_H$. Similarly, since in case (c) of Theorem 3 we assume that $\tau_L < M$, by definition $\tau_H = \min\{\hat{\tau}_H, M\}$ and hence $\tau_L < \tau_H$. In addition, using the definition of $\hat{\tau}_H$ in Equation (C.19) as well as the fact that the function $G(\beta_D, s_D, \tilde{s})$ in Proposition 14 is unimodal in s_D , it follows that $\hat{\tau}_H < M$ (and hence $\tau_H < M$) if and only if $\alpha(r-w)\gamma'(M)(1-q(1-v)) < pq(1-v)$.

To summarize, we have shown that if $\beta_{LD} < \hat{\beta}_D$, then (i) the manufacturer's optimal investment strategy is equal to $\beta^* = \bar{\beta}_{HD} = \min\beta_{HD}, \hat{\beta}_D$ for $\tilde{s} \leq \tau_D$ and $\beta^* = \beta_{LD}$ for $\tilde{s} > \tau_D$; (ii) if $\alpha(1-q)(r-w)\gamma'(M) < pq$, then the manufacturer's optimal disclosure strategy is equal to a piece-wise constant, non-increasing s_D^* for $\tilde{s} \leq \tau_L$, and strictly less than M ; and equal to $s_D^* = z(\tilde{s})$ for $\tilde{s} > \tau_L$. By Proposition 14, this means that if $\tilde{s} > \tau_L$, then $s_D^* = \tilde{s}$ if $\tilde{s} \leq \tau_H$; and $s_D^* = \tau_H$ if $\tilde{s} \in (\tau_H, M]$. Furthermore, we have also shown that $\tau_D \leq \tau_L < \tau_H \leq M$.

Knowing the manufacturer's optimal disclosure strategy s_D^* , we can finally compare it against her best estimate of the supplier's final level of SR. From the proof of the structure of s_D^* above, in Region (i) we can see that either (a) $s_D^* = \hat{s}_D = \check{s}(\bar{\beta}_{HD})$ for $\tilde{s} \leq \tau_D$ and $s_D^* = \tau_L$ for $\tilde{s} \in (\tau_D, \tau_L]$; or (b) $s_D^* = \tau_L > \check{s}(\bar{\beta}_{HD}) > \check{s}(\beta_{LD})$ for $\tilde{s} \leq \tau_L$. In both cases, we obtain that in Region (i), s_D^* is greater than or equal to $s^*(\tilde{s}, \beta_D^*)$; by Proposition 12 $s^*(\tilde{s}, \beta_D^*) = \check{s}(\bar{\beta}_{HD})$ for $\tilde{s} \leq \tau_D$ and $s^*(\tilde{s}, \beta_D^*) = \tilde{s} \leq \tau_L$ for $\tilde{s} \in (\tau_D, \tau_L]$. For $\tilde{s} > \tau_L$, since $\tau_L \geq \tau_D$, by Proposition 12 we know that $s^*(\tilde{s}, \beta_D^*) = \tilde{s}$. In Region (ii), the manufacturer discloses $s_D^* = \tilde{s}$ and thus both values coincide. Finally, if Region (iii) exists (i.e., if $\tau_H < M$), then the manufacturer discloses $s_D^* = \tau_H$ for all $\tilde{s} \in (\tau_H, M]$. Thus, in this region we get that $s_D^* = \tau_H < \tilde{s} = s^*(\tilde{s}, \beta_D^*)$. \square

Corollary 5. *If $\tilde{s} = \tau_D$, then the manufacturer is indifferent between the strategies $(\bar{\beta}_{HD}, s_{DH})$ and (β_{LD}, s_{DL}) , where s_{DH} (s_{DL}) is equal to the optimal value of s_D given a signal immediately to the left (immediately to the right) of $\tilde{s} = \tau_D$. Furthermore,*

τ_D satisfies $\check{s}(\beta_{LD}) < \tau_D < \check{s}(\bar{\beta}_{HD})$.

Proof. The manufacturer being indifferent between the strategies $(\bar{\beta}_{HD}, s_{DH})$ and (β_{LD}, s_{DL}) follows directly from the proof of Theorems 3 and 6, in the case where $\beta_{LD} < \hat{\beta}_D$. Finally, from the structure of the manufacturer's expected profit function in Lemma 6 and the unimodality of the functions $G_L(\beta_D, s_D, \tilde{s})$ and $G_H(\beta_D, s_D)$ in β_D , we know that $\mathbb{E}_{s_0}[\Pi_M(\bar{\beta}_{HD}(v), s_{DH}) | \tilde{s} = \tau_D]$ can only be equal to $\mathbb{E}_{s_0}[\Pi_M(\beta_{LD}(v), s_{DL}) | \tilde{s} = \tau_D]$ if τ_D satisfies $\check{s}(\beta_{LD}) < \tau_D < \check{s}(\bar{\beta}_{HD})$. \square

Proof of Proposition 12. First, from the structure of the optimal β_D in Theorem 6, we know that if $\tilde{s} \leq \tau_D$, then $\beta_D^* = \bar{\beta}_{HD}$. Thus, assuming that $s_0 = \tilde{s} \leq \tau_D$, we obtain $s^*(s_0 = \tilde{s}, \beta_D) = s^*(s_0 = \tilde{s}, \bar{\beta}_{HD}) = \check{s}(\bar{\beta}_{HD})$, where the last equality follows from $s^*(s_0, \beta_D) = \max\{s_0, \check{s}(\beta_D)\}$ (see Corollary 4) and the fact that $\check{s}(\beta_{LD}) < \tau_D < \check{s}(\bar{\beta}_{HD})$ (see Corollary 5).

Following the same steps, if $\tilde{s} > \tau_D$, then $\beta_D^* = \beta_{LD}$ and $s^*(s_0 = \tilde{s}, \beta_D) = s^*(s_0 = \tilde{s}, \beta_{LD}) = s_0$, since $s_0 = \tilde{s} > \tau_D$. Finally, the two results combined mean that the best estimate of SR is discontinuous at $\tilde{s} = \tau_D$, as $\lim_{\tilde{s} \uparrow \tau_D} s^*(s_0 = \tilde{s}, \beta_D) = \check{s}(\bar{\beta}_{HD}) > \tau_D = \lim_{\tilde{s} \downarrow \tau_D} s^*(s_0 = \tilde{s}, \beta_D)$, where the inequality follows directly from Corollary 5. \square

C.3 Comparative Statics

We next show how the level of visibility and the probability of third-party scrutiny affect the manufacturer's optimal strategy. First, we show how v affects the manufacturer's choice of s_D (Proposition 5). Then, we study how the manufacturer's best estimate of the supplier's SR level changes with q (Lemma 10) so that we can then demonstrate the effect of q on the difference between the manufacturer's disclosure and her best estimate of SR in Proposition 6.

We also demonstrate how v and q impact the manufacturer's optimal choice of β_D . First, Lemma 11 presents a result that is useful for studying the dependency of

the threshold τ_D on model parameters. We then show in Proposition 8 how visibility affects β_{LD} , β_{HD} , and τ_D . Finally, we show in Proposition 9 how τ_D changes with q .

Proof of Proposition 5. First, we show that τ_H is strictly increasing in v if $\tau_H < M$. From the proof of Theorems 3 and 6, and using the notation and properties from Proposition 14, we know that $\tau_H = \hat{\tau}_H$. Thus, assuming that there exists $\hat{\tau}_H < M$ that satisfies the condition $V_1(\hat{\tau}_H) = 0$ in Equation (C.19), we need to show that $\hat{\tau}_H$ is strictly increasing in v . Using the Implicit Function Theorem, we can study the dependence of $\tau_H = \hat{\tau}_H$ on v as follows:

$$\frac{dV_1(\hat{\tau}_H)}{dv} = \left[\alpha q(r-w) \left(\gamma''(\hat{\tau}_H)(1-q(1-v)\Phi(\hat{\tau}_H)) - \gamma'(\hat{\tau}_H)(1-v)\phi(\hat{\tau}_H) \right) - pq(1-v)\phi(\hat{\tau}_H) \right] \frac{d\hat{\tau}_H}{dv} \quad (\text{C.27})$$

$$+ \alpha q(r-w)\gamma'(\hat{\tau}_H)\Phi(\hat{\tau}_H) + pq\Phi(\hat{\tau}_H) = 0. \quad (\text{C.28})$$

Since the terms in brackets in (C.27) are all negative (in particular, because $\gamma(s)$ is strictly increasing and strictly concave in s), and the terms in (C.28) are all positive, it follows that $\frac{d\hat{\tau}_H}{dv} > 0$ – it is necessary to guarantee that $\frac{dV_1(\hat{\tau}_H)}{dv} = 0$. That is, $\tau_H = \hat{\tau}_H$ is strictly increasing in v .

Similarly, by Proposition 14 and using the proof of Theorems 3 and 6, we also know that (i) if $\check{s}(\beta_{HD}) < \hat{\tau}_L$, then $\tau_L = \hat{\tau}_L$ and $s_D^* = \tau_L$ for all $\tilde{s} \leq \tau_L$; or (ii) if $\check{s}(\beta_{HD}) \geq \hat{\tau}_L$, then $\tau_L = \max\{\tau_D, \hat{\tau}_L\}$, where τ_D is the signal $\tilde{s} = \tau_D < \hat{\tau}_H$ such that the manufacturer prefers the strategy $(\hat{\beta}_D, \hat{s}_D)$ for all $\tilde{s} \leq \tau_D$, and she prefers $(\beta_{LD}, s_D = z(\tilde{s}))$ for all $\tilde{s} > \tau_D$. Since $\tau_L = \tau_D$ only if $\tau_D > \hat{\tau}_L$, the latter strategy can be replaced by $(\beta_{LD}, s_D = \tilde{s})$. In the remainder of this proof, we first show that the threshold τ_L is strictly decreasing in v and then we show that s_D^* is non-increasing in v for all $\tilde{s} \leq \tau_L$.

For case (i) above, if $\check{s}(\beta_{HD}) < \hat{\tau}_L$, then $\tau_L = \hat{\tau}_L$. Thus, we can show that τ_L is strictly decreasing in v by showing that $\hat{\tau}_L$ is strictly decreasing in v , which follows directly from the Implicit Function Theorem applied to the condition $V_2(\hat{\tau}_L) = 0$ in Equation (C.20) (i.e., following the same steps as in Equation (C.27)–(C.28)). For case (ii) above, if $\check{s}(\beta_{HD}) \geq \hat{\tau}_L$, then $\tau_L = \max\{\tau_D, \hat{\tau}_L\}$. Since we have already shown

that $\hat{\tau}_L$ is decreasing in v , we need to show that τ_D is also decreasing in v when $\tau_D > \hat{\tau}_L$. We first note that in this case, from the proof of Theorems 3 and 6, τ_D corresponds to the signal \tilde{s} such that the manufacturer is indifferent between selecting the strategies $(\hat{\beta}_D, \hat{s}_D)$ and $(\beta_{LD}, s_{DL} = \tilde{s})$. Similarly, we also know from that proof that if the manufacturer selects $(\beta_{LD}, s_{DL} = \tilde{s})$, then her expected profits are equal to $G_L(\beta_{LD}, \tilde{s}, \tilde{s})$, which is strictly increasing in \tilde{s} . Finally, by Lemma 7 we know that $\hat{\beta}_D$ and \hat{s}_D are independent of v . Thus, for case (ii) we need to show that $\frac{dG_L(\beta_{LD}, \tilde{s}, \tilde{s})}{dv} > 0$: if this condition was satisfied, then a smaller value of τ_D would be required to make the manufacturer indifferent between the two strategies.

From the definition of the function G_L in Lemma 3 and using the Envelope Theorem, we know that

$$\begin{aligned} \frac{dG_L(\beta_{LD}, \tilde{s}, \tilde{s})}{dv} &= \frac{\partial G_L(\beta_{LD}, \tilde{s}, \tilde{s})}{\partial v} \\ &= \alpha q (r - w) \left(\gamma(\tilde{s}) \Phi(\tilde{s}) - \int_{\tilde{s}(\beta_{LD})}^{\tilde{s}} \gamma(x) \phi(x) dx - \gamma(\tilde{s}(\beta_{LD})) \Phi(\tilde{s}(\beta_{LD})) \right) \\ &\quad + p q \left((\tilde{s} - \tilde{s}(\beta_{LD})) \Phi(\tilde{s}(\beta_{LD})) + \int_{\tilde{s}(\beta_{LD})}^{\tilde{s}} (\tilde{s} - x) \phi(x) dx \right) > 0. \end{aligned}$$

Thus, if $\tau_L = \tau_D$, then τ_D is decreasing in v , which completes the proof for τ_L being strictly decreasing in v in all cases.

Finally, we study how s_D^* changes with v . For case (i), if $\check{s}(\beta_{HD}) < \hat{\tau}_L$, then $\tau_L = \hat{\tau}_L$ and thus $s_D^* = \tau_L$ for all $\tilde{s} \leq \tau_L$ and we have already shown above that $\hat{\tau}_L$ is decreasing in v . For case (ii), if $\check{s}(\beta_{HD}) \geq \hat{\tau}_L$, then two alternatives are possible: (i) if $\tau_L = \tau_D$, then the manufacturer selects $(\hat{\beta}_D, \hat{s}_D)$ for $\tilde{s} \leq \tau_L = \tau_D$, where \hat{s}_D is independent of v ; or (ii) if $\tau_L = \hat{\tau}_L > \tau_D$, then for $\tilde{s} \leq \tau_D$, the manufacturer selects $(\hat{\beta}_D, \hat{s}_D)$; if instead $\tilde{s} \in (\tau_D, \hat{\tau}_L]$, then she selects $(\beta_{LD}, s_D = z(\tilde{s}) = \hat{\tau}_L)$. Since both $\hat{\tau}_L$ and τ_D are decreasing in v and $\tau_L < \hat{s}_D$, it follows that for all $\tilde{s} \leq \tau_L$, $s_D^*(\tilde{s})$ is strictly decreasing in v . \square

Lemma 10. *The manufacturer's best estimate of the supplier's final level of SR, $s^*(s_0 = \tilde{s}, \beta_D = \beta_D^*(v, \tilde{s}))$, is strictly increasing in q .*

Proof. This result follows directly from: (i) the structure of the manufacturer's best estimate of SR in Proposition 12, whereby it is equal to $\tilde{s}(\bar{\beta}_{HD})$ for $\tilde{s} \leq \tau_D$ and equal to s_0 otherwise; (ii) by Proposition 3, $\tilde{s}(\bar{\beta}_{HD})$ is non-decreasing in q ; and (iii) by Proposition 9 τ_D is strictly increasing in q . \square

Proof of Proposition 6. The proof of this Lemma follows almost the same logic as that of Proposition 5, so we refer the reader to that proof for more details.

First, from the proof of Theorems 3 and 6, we know that $\tau_H = \hat{\tau}_H$, which is defined in Proposition 14. Thus, to show that τ_H is strictly decreasing in q , we study how the condition in Equation (C.19) varies with q . In particular, using the Implicit Function Theorem,

$$\frac{dV_1(\hat{\tau}_H)}{dq} = \left[\alpha q(r-w) \left(\gamma''(\hat{\tau}_H)(1-q(1-v)\Phi(\hat{\tau}_H)) - \gamma'(\hat{\tau}_H)(1-v)\phi(\hat{\tau}_H) \right) - pq(1-v)\phi(\hat{\tau}_H) \right] \frac{d\hat{\tau}_H}{dq} \quad (\text{C.29})$$

$$- \alpha(r-w)\gamma'(\hat{\tau}_H)(1-v)\Phi(\hat{\tau}_H) - p(1-v)\Phi(\hat{\tau}_H) = 0. \quad (\text{C.30})$$

Since the terms in brackets in (C.29) are all negative (in particular, because $\gamma(s)$ is strictly increasing and strictly concave in s), and the terms in (C.30) are also negative, it follows that $\frac{d\hat{\tau}_H}{dq} < 0$ – it is necessary to guarantee that $\frac{dV_1(\hat{\tau}_H)}{dq} = 0$. That is, assuming that there exists $\hat{\tau}_H < M$ that satisfies $V_1(\hat{\tau}_H) = 0$ (i.e., if Region (iii) from Theorem 3 exists), $\hat{\tau}_H$ is strictly decreasing in q . Similarly, by Theorem 3, in Region (iii) the optimal disclosure is $s_D^* = \tau_H$ and we also know that $\tau_H = \hat{\tau}_H$. Thus, if Region (iii) exists, then in this region s_D^* also decreases with q .

In addition, from Theorem 3 we also know that $\tau_L \geq \tau_D$ and thus in Region (iii), $\beta_D^* = \beta_{LD}$. Similarly, we also know from Proposition 12 that in this region the manufacturer's best estimate of the SR level is equal to \tilde{s} . Thus, $\Omega(v, \tilde{s}) = s_D^*(v, \tilde{s}) - s^*(s_0 = \tilde{s}, \beta_D^*(v, \tilde{s}))$ is strictly decreasing in q in Region (iii). Furthermore, Theorem 3 also states that in Region (ii), $s_D^* = \tilde{s}$. Using again the fact that $\tau_D \leq \tau_L$, it follows that in this region the manufacturer's best estimate of the SR level is exactly equal

to \tilde{s} . Thus, in Region (ii) we know that $\Omega(v, \tilde{s}) = 0$.

Finally, we need to show that $\Omega(v, \tilde{s})$ is non-increasing in q in Region (i). First, if $\tau_L = \tau_D$, then the manufacturer's optimal strategy in this region is $(\beta_D = \hat{\beta}_D, s_D = \hat{s}_D)$ (see the proof of Theorem 3) and thus her optimal disclosure and best estimate of SR are both equal to \hat{s}_D ; i.e., $\Omega(v, \tilde{s}) = 0$. Conversely, if $\tau_L = \hat{\tau}_L > \tau_D$, then we know that the manufacturer's optimal disclosure strategy is equal to (i) if $\tilde{s} \leq \tau_D$, then either $s_D^* = \hat{s}_D$ (independent of q , see Lemma 7) or $s_D^* = \hat{\tau}_L$; and (ii) if $\tilde{s} \in (\tau_D, \hat{\tau}_L]$, then $s_D^* = \hat{\tau}_L$. We note that (i) the manufacturer's best estimate of the supplier's SR level is increasing in q if $\tilde{s} \leq \tau_D$ (see Lemma 10); (ii) her best estimate of the SR level is equal to \tilde{s} if $\tilde{s} > \tau_D$ (Proposition 12); and (iii) τ_D is strictly increasing in q by Proposition 9. Thus, all we need to show is that $\hat{\tau}_L$ is strictly decreasing in q . This follows directly from applying the Implicit Function Theorem to the condition $V_2(\hat{\tau}_L) = 0$ in Equation (C.20), as we did in Equation (C.29)–(C.30). \square

Lemma 11. *Let us define*

$$\Delta(\tilde{s}, x) = \mathbb{E}_{s_0} [\Pi_M(\bar{\beta}_{HD}(v), s_{D_H}) | \tilde{s}] - \mathbb{E}_{s_0} [\Pi_M(\beta_{LD}(v), s_{D_L}) | \tilde{s}], \quad (\text{C.31})$$

where x represents any parameter of the manufacturer's expected profit function. The threshold $\tau_D(v)$ is strictly increasing in x if and only if $\frac{\partial \Delta(\tilde{s}=\tau_D(v), x)}{\partial x} > 0$.

Proof. First, we show that if $\Delta(\tilde{s} = \tau_D(x), x)$ is strictly increasing in x , then the threshold $\tilde{s}(x)$ increases in x . Then, we show that in order to determine whether $\Delta(\tilde{s} = \tau_D, x)$ is indeed increasing in x , we only need to study the partial derivative of $\Delta(\tilde{s} = \tau_D, x)$ with respect to x . Taken together, these two results imply the result in Lemma 11; i.e., that if $\frac{\partial \Delta(\tilde{s}=\tau_D(v), x)}{\partial x} > 0$, then the threshold $\tau_D(v)$ is strictly increasing in x . We show this result for any parameter x and we later use it for the particular cases where x equals v , q , or p .

In what follows, we use the notation $\tau_D(x)$ to make the dependence of τ_D on the parameter x explicit. Similarly, by Corollary 5 we know that $\Delta(\tilde{s} = \tau_D(x), x) = 0$. Suppose that $\Delta(\tilde{s} = \tau_D(x), x)$ is strictly increasing in x . From the definition of

$\Delta(\tilde{s}, x)$, this means that the manufacturer's expected profits at $\tilde{s} = \tau_D(x)$ increase faster in x if she selects $(\bar{\beta}_{HD}, s_{DH})$ than if she selects (β_{LD}, s_{DL}) . Thus, there exists $\bar{\epsilon} > 0$ such that for any $\epsilon \leq \bar{\epsilon}$, $\Delta(\tilde{s} = \tau_D(x), x + \epsilon) > 0$. In addition, the structure of the optimal β_D is always such that $\bar{\beta}_{HD}$ is selected by the manufacturer for every signal $\tilde{s} \leq \tau_D$; and β_{LD} is instead preferred for every signal $\tilde{s} > \tau_D$ (see Theorem 6). Equivalently, from the definition of $\Delta(\tilde{s}, x)$, this means that for any given x , $\Delta(\tilde{s}, x) \geq 0$ for all $\tilde{s} \leq \tau_D$ and $\Delta(\tilde{s}, x) < 0$ for all $\tilde{s} > \tau_D$. Therefore, $\Delta(\tilde{s} = \tau_D(x), x + \epsilon) > 0$ implies that a signal greater than $\tau_D(x)$ is required to make the manufacturer indifferent between the two possible strategies when the parameter value is instead $x + \epsilon$; i.e., $\tau_D(x + \epsilon) > \tau_D(x)$.

Next, we show that it is sufficient to study the partial derivative of $\Delta(\tilde{s}, x)$ with respect to x to determine whether this function is strictly increasing in x . The total derivative of the manufacturer's expected profits with respect to x can be written as:

$$\frac{d\mathbb{E}[\Pi_M]}{dx} = \frac{\partial\mathbb{E}[\Pi_M]}{\partial x} + \frac{\partial\mathbb{E}[\Pi_M]}{\partial\beta_D} \frac{d\beta_D}{dx} + \frac{\partial\mathbb{E}[\Pi_M]}{\partial s_D} \frac{ds_D}{dx}, \quad (\text{C.32})$$

where for ease of notation we write $\mathbb{E}_{s_0}[\Pi_M(\beta_D, s_D) | \tilde{s}]$ simply as $\mathbb{E}[\Pi_M]$. From Theorem 6, we know that β_D^* is equal to either β_{LD} , β_{HD} , or $\hat{\beta}_D$. In the first two cases, by Lemmas 3 and 4 and Lemma 6, it follows that $\frac{\partial\mathbb{E}[\Pi_M]}{\partial\beta_D} = 0$. Similarly, if $\beta_D^* \neq \hat{\beta}_D$, then from the proof of Theorems 3 and 6, s_D^* either satisfies the FOC $\frac{\partial\mathbb{E}[\Pi_M]}{\partial s_D} = 0$ or $s_D^* = \tilde{s}$, in which case $\frac{ds_D}{dx} = 0$. Therefore, if the optimal β_D is equal to either β_{LD} or β_{HD} , then $\frac{d\mathbb{E}[\Pi_M]}{dx} = \frac{\partial\mathbb{E}[\Pi_M]}{\partial x}$ and thus, from the definition of Δ in Equation (C.31), it follows that $\frac{d\Delta}{dx} = \frac{\partial\Delta}{\partial x}$.

In the last case, if $\beta_D^* = \hat{\beta}_D$, then by Lemma 8, $s_D^* = \hat{s}_D = \check{s}(\hat{\beta}_D)$ and the manufacturer's expected profit function can be written as:

$$\begin{aligned} \mathbb{E}[\Pi_M(\beta_D = \hat{\beta}_D)] &\equiv \mathbb{E}_{s_0}[\Pi_M(\beta_D = \hat{\beta}_D, s_D = \check{s}(\hat{\beta}_D)) | \tilde{s}] \\ &= (1 - r + \alpha \gamma(\check{s}(\hat{\beta}_D))) (r - w) - \delta(\hat{\beta}_D). \end{aligned}$$

This function is continuously differentiable in β_D and the model parameters (see the

definition of $\check{s}(\beta_D)$ in Lemma 1). Thus, using the Envelope Theorem we obtain that $\frac{d\mathbb{E}[\Pi_M(\beta_D=\hat{\beta}_D)]}{dx} = \frac{\partial\mathbb{E}[\Pi_M(\beta_D=\hat{\beta}_D)]}{\partial x}$. This proves that in the case where $\beta_D^* = \hat{\beta}_D$, we also have that $\frac{d\Delta}{dx} = \frac{\partial\Delta}{\partial x}$. Thus, studying the partial derivative of $\Delta(\check{s} = \tau_D(v), x)$ with respect to x is sufficient to determine whether this function is strictly increasing in x . \square

Proof of Proposition 8(i). First, we show that β_{LD} is strictly decreasing in v . Since $\beta_{LD}(v)$ satisfies the above FOC, we can compute the cross partial derivative of G_L with respect to β_D and v as follows:

$$\frac{\partial^2 G_L}{\partial \beta_D \partial v} = -(\alpha q (r - w) \gamma'(\check{s}) + p q) \frac{R(\check{s})}{1 - \beta_{LD}} \Phi(\check{s}) < 0.$$

The last inequality follows from: $\gamma(s)$ is strictly increasing in s ; $\beta_{LD} < 1$ (which follows directly from Assumption 3); by definition, $R(s) > 0$ because $\rho''(s) > 0$ and $\gamma''(s) < 0$; and $\Phi(\check{s}) > 0$ because for any $\beta_D \geq 0$, the lowest supplier types increase their level of SR; i.e., $\check{s} > m$ (Assumption 2). Thus, the function G_L is submodular in (β_D, v) and β_{LD} solves the unconstrained optimization problem $\max_{\beta_D} G_L(\beta_D, s_D, \check{s})$. By Topkis's theorem, this implies that β_{LD} is strictly decreasing in v .

The proof of β_{HD} being strictly increasing in v follows the same approach and is therefore omitted. \square

Note: No Manufacturer Disclosure setting. The above proof also demonstrates Proposition 1 for the specific case $p = 0$.

Proof of Proposition 8(ii). First, we note that using the law of total probability, the manufacturer's expected profits given signal \check{s} can be written as

$$\mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \check{s}] = v \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \check{s}, s_0 = \check{s}] + (1 - v) \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \check{s}, s_0 \neq \check{s}]. \quad (\text{C.33})$$

Similarly, since the cost of selecting β_D is known with certainty to the manufacturer, we can

write her expected profits as follows:

$$\mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \tilde{s}] = \mathbb{E}_{s_0} [\Theta(\beta_D, s_D) | \tilde{s}] - \delta(\beta_D),$$

where $\Theta(\beta_D, s_D)$ corresponds to all the components of the manufacturer's profits except for $\delta(\beta_D)$; from the definition of the manufacturer's expected profits in Equation (3.9), we know that $\Theta(\beta_D, s_D)$ is independent of v . Replacing this term in Equation (C.33), we obtain

$$\begin{aligned} \mathbb{E}_{s_0} [\Pi_M(\beta_D, s_D) | \tilde{s}] &= v \mathbb{E}_{s_0} [\Theta(\beta_D, s_D) | \tilde{s}, s_0 = \tilde{s}] \\ &\quad + (1 - v) \mathbb{E}_{s_0} [\Theta(\beta_D, s_D) | \tilde{s}, s_0 \neq \tilde{s}] - \delta(\beta_D). \end{aligned} \quad (\text{C.34})$$

Given Equation (C.34), the difference in expected profits between the strategies (β_{LD}, s_{DL}) and $(\bar{\beta}_{HD}, s_{DH})$ in Lemma 11, can be written as

$$\begin{aligned} \Delta(\tilde{s}, v) &= \delta(\beta_{LD}) - \delta(\bar{\beta}_{HD}) + \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) | \tilde{s}] - \mathbb{E}_{s_0} [\Theta(\beta_{LD}, s_{DL}) | \tilde{s}] \\ &= \delta(\beta_{LD}) - \delta(\bar{\beta}_{HD}) + v \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | \tilde{s}, s_0 = \tilde{s}] \\ &\quad + (1 - v) \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | \tilde{s}, s_0 \neq \tilde{s}]. \end{aligned} \quad (\text{C.35})$$

The partial derivative of $\Delta(\tilde{s}, v)$ with respect to v is then equal to

$$\begin{aligned} \frac{\partial \Delta(\tilde{s}, v)}{\partial v} &= \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | \tilde{s}, s_0 = \tilde{s}] \\ &\quad - \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | \tilde{s}, s_0 \neq \tilde{s}], \end{aligned} \quad (\text{C.36})$$

which corresponds to the difference in differences in expected profits between both strategies under consideration when the signal is correct versus when the signal is not correct.

Rewriting Equation (C.35) with the result from Equation (C.36), we find that

$$\Delta(\tilde{s}, v) = \delta(\beta_{LD}) - \delta(\bar{\beta}_{HD}) + v \frac{\partial \Delta(\tilde{s}, v)}{\partial v} + \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | \tilde{s}, s_0 \neq \tilde{s}].$$

By definition, τ_D is defined as the signal \tilde{s} where $\Delta(\tilde{s}, v) = 0$. In addition, we note that in our model if the signal \tilde{s} is not correct, then it is equal to any s_0 with p.d.f. $\phi(s_0)$. Thus, if the manufacturer knows that the signal is not correct, then it is uninformative and the expected difference in profits between both strategies given $s_0 \neq \tilde{s}$ (i.e., knowing that the

signal is not correct) does not depend on the value of \tilde{s} . The above equation evaluated at τ_D can therefore be written as

$$\begin{aligned} \Delta(\tilde{s} = \tau_D, v) &= \delta(\beta_{LD}) - \delta(\bar{\beta}_{HD}) + v \frac{\partial \Delta(\tau_D, v)}{\partial v} \\ &\quad + \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | s_0 \neq \tilde{s}] = 0, \end{aligned}$$

or equivalently

$$v \frac{\partial \Delta(\tau_D, v)}{\partial v} = \delta(\bar{\beta}_{HD}) - \delta(\beta_{LD}) - \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | s_0 \neq \tilde{s}].$$

Therefore, it follows that $\frac{\partial \Delta(\tau_D, v)}{\partial v} > 0$ if and only if

$$\delta(\bar{\beta}_{HD}) - \delta(\beta_{LD}) > \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | s_0 \neq \tilde{s}]. \quad (\text{C.37})$$

Finally, by Lemma 11 we know that for any parameter x , τ_D is strictly increasing in x if and only if $\frac{\partial \Delta(\tau_D, x)}{\partial v} > 0$. Therefore, τ_D is strictly increasing in v if and only if the condition in Equation (C.37) is satisfied. \square

Note: No Manufacturer Disclosure setting. Given $p = 0$, in this setting we have

$$\begin{aligned} \mathbb{E}_{s_0} [\Theta(\bar{\beta}_{HD}, s_{DH}) - \Theta(\beta_{LD}, s_{DL}) | s_0 \neq \tilde{s}] &= \\ \alpha q(r - w) \mathbb{E}_{s_0} [\gamma(\min\{s^*(s_0, \bar{\beta}_H), \hat{s}\}) - \gamma(\min\{s^*(s_0, \beta_L), \hat{s}\}) | \tilde{s} \neq s_0]. \end{aligned}$$

Thus, we can simplify the condition in Equation (C.37) as follows: $\frac{\partial \Delta(\tau_D, v)}{\partial v} > 0$ if and only if

$$\delta(\bar{\beta}_{HD}) - \delta(\beta_{LD}) > \alpha q(r - w) \mathbb{E}_{s_0} [\gamma(\min\{s^*(s_0, \bar{\beta}_H), \hat{s}\}) - \gamma(\min\{s^*(s_0, \beta_L), \hat{s}\}) | \tilde{s} \neq s_0]. \quad (\text{C.38})$$

Finally, by Lemma 11, τ_D is strictly increasing in v in the No Manufacturer Disclosure setting if and only if the condition in Equation (C.38) is satisfied. This proves Proposition 1(ii).

Proof of Proposition 9. Based on Lemma 11 and using the notation therein, it

is sufficient to show that $\frac{\partial \Delta(\tilde{s}=\tau_D, q)}{\partial q} > 0$. From the definition of the manufacturer's expected profits in Equation (3.9), we know that

$$\begin{aligned} \frac{\partial \Delta(\tau_D, q)}{\partial q} = & \\ & \alpha(r-w) \left(\mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \bar{\beta}_{HD}, s_{DH}), s_{DH}\}) \mid \tau_D \right] \right. \\ & \left. - \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta_{LD}, s_{DL}), s_{DL}\}) \mid \tau_D \right] \right) \end{aligned} \quad (\text{C.39})$$

$$+ p \left(\mathbb{E}_{s_0} \left[(s_{DL} - s^*(s_0, \beta_{LD}))^+ \mid \tau_D \right] - \mathbb{E}_{s_0} \left[(s_{DH} - s^*(s_0, \bar{\beta}_{HD}))^+ \mid \tau_D \right] \right), \quad (\text{C.40})$$

which we can decompose into (i) the difference in consumers' SR utility $\gamma(s)$ in (C.39); and (ii) the difference in expected penalty between the two alternative strategies, $(\bar{\beta}_{HD}, s_{DH})$ and (β_{LD}, s_{DL}) in (C.40).

First, we show that the difference in consumers' SR utility in (C.39) is strictly positive. From Theorem 6 we know the structure of the optimal β_D as a function of \tilde{s} , namely equal to $\bar{\beta}_{HD}$ for any $\tilde{s} \leq \tau_D$ and equal to β_{LD} for any $\tilde{s} > \tau_D$. In addition, from Proposition 8 we know that $\beta_{HD} > \beta_{LD}$ (since we assume that $v > 0$). Similarly, from the proof of Theorems 3 and 6, we know that $s_{DH} \geq s_{DL}$. Since the supplier's best-response is non-decreasing in both β_D and s_D , and strictly increasing for at least some s_0 (see Assumption 2 and Theorem 5), this implies that $\mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \bar{\beta}_{HD}, s_{DH}), s_{DH}\}) \mid \tau_D \right] > \mathbb{E}_{s_0} \left[\gamma(\min\{s^*(s_0, \beta_{LD}, s_{DL}), s_{DL}\}) \mid \tau_D \right]$. Thus, the first component of $\frac{\partial \Delta(\tau_D, q)}{\partial q}$ is strictly positive.

Next, we show that the difference in expected penalty in (C.40), between the alternative strategies $(\bar{\beta}_{HD}, s_{DH})$ and (β_{LD}, s_{DL}) evaluated at $\tilde{s} = \tau_D$, is also strictly positive. We first note that at τ_D , both strategies are candidate solutions (i.e., local optima, see the proof of Theorems 3 and 6). We also note that we can identify two cases. First, if $\bar{\beta}_{HD} = \min\{\beta_{HD}, \hat{\beta}_D\} = \hat{\beta}_D$, then it must be the case that $s_{DH} = \tilde{s}(\hat{\beta}_D)$ (see Lemma 8). As a result, $s^*(s_0, \hat{\beta}_D) = \max\{s_0, \min\{s_{DH}, \tilde{s}(\hat{\beta}_D)\}\} \geq s_{DH}$ and hence $\mathbb{E}_{s_0} \left[(s_{DH} - s^*(s_0, \bar{\beta}_{HD}))^+ \mid \tau_D \right] = 0$. Similarly, we know that β_{LD} must satisfy $\beta_{LD} < \hat{\beta}_D$ for τ_D to exist. Since s_{DL} is the disclosed SR level when selecting $\beta_D = \beta_{LD}$, by Lemma 7 and the fact that any optimal solution must satisfy $\tilde{s} \leq s_D$,

it follows that $\check{s}(\beta_{LD}) < s_{DL}$. Therefore, for any $s_0 \leq \check{s}(\beta_{LD})$, the supplier's optimal response is $s^*(s_0, \beta_{LD}) = \max\{s_0, \check{s}(\beta_{LD})\} = \check{s}(\beta_{LD}) < s_{DL}$, and hence, $E_{s_0} \left[(s_{DL} - s^*(s_0, \beta_{LD}))^+ | \tau_D \right] > 0$. Taken together, these results imply that if $\bar{\beta}_{HD} = \hat{\beta}_D$, then the difference in expected penalty is positive.

In the second case, if $\bar{\beta}_{HD} < \hat{\beta}_D$, then $\beta_H < \hat{\beta}_D$ and $\bar{\beta}_{HD} = \beta_{HD}$. Since at $\bar{s} = \tau_D$ both strategies (β_{HD}, s_{DH}) and (β_{LD}, s_{DL}) are candidate solutions, then by Lemma 8, cases (i) and (ii), it follows that $s_{DH} = s_{DL} = \min\{z(\tau_D), M\}$ and it must also be the case that $\check{s}(\beta_{LD}) < z(\tau_D)$ and $\check{s}(\beta_{HD}) < z(\tau_D)$. In addition, since $\beta_{LD} < \beta_{HD}$ for any $v > 0$ and \check{s} is strictly increasing in β_D , then $\check{s}(\beta_{LD}) < \check{s}(\beta_{HD})$ and hence $s^*(s_0, \beta_{LD}) \leq s^*(s_0, \bar{\beta}_{HD})$ for all s_0 . In particular, given the supplier's best response in Theorem 5, the last inequality is strict for any $s_0 \in [\check{s}(\beta_{LD}), \check{s}(\beta_{HD})]$. Combining these results, it follows that $E_{s_0} \left[(s_{DL} - s^*(s_0, \beta_{LD}))^+ | \tau_D \right]$ is strictly positive and strictly greater than $E_{s_0} \left[(s_{DH} - s^*(s_0, \bar{\beta}_{HD}))^+ | \tau_D \right]$. Thus, if $\bar{\beta}_{HD} < \hat{\beta}_D$, then we also obtain that the difference in expected penalty in (C.40) is positive.

Using the above results, we have shown that $\frac{\partial \Delta(\tau_D, q)}{\partial q} > 0$ and therefore, by Lemma 11, the threshold τ_D is strictly increasing in q . \square

Note: No Manufacturer Disclosure setting. The above proof also demonstrates Proposition 2.

C.4 When Does the Manufacturer Prefer to Voluntarily Disclose SR Information?

In this section, we show our results regarding the comparison between the Manufacturer Disclosure and No Manufacturer Disclosure settings. First, in Proposition 4 we show that for any signal and level of visibility, $\beta_D^*(v, \bar{s}) \geq \beta^*(v, \bar{s})$ and that the difference between them increases with p . Then, we show how the difference in the optimal profits between both settings changes with the signal \bar{s} (Theorem 4). Using this result, we then show the properties of the manufacturer's optimal setting, sum-

marized in Corollaries 1 and 2. Finally, Lemma 12 identifies a property that is useful for studying the comparison between both settings in the specific case when $\tau_H < \hat{s}$ (Theorem 7).

Proof of Proposition 4. First, we show that $\beta_{LD} > \beta_L$ and $\beta_{HD} > \beta_H$. As noted earlier, the No Manufacturer Disclosure setting is a specific case of the Manufacturer Disclosure setting where $\hat{s} = s_D$ and $p = 0$. In addition, by Proposition 13, β_{LD} and β_{HD} are the optimal solutions to $G_L(\beta_D, s_D, \tilde{s})$ and $G_H(\beta_D, s_D)$, respectively, and by Lemmas 3 and 4 they are independent of s_D . Thus, it suffices to show that β_{LD} and β_{HD} are strictly increasing in p .

From the first order conditions of β_{LD} and β_{HD} in Lemmas 3 and 4, we know that

$$\frac{\partial^2 G_L}{\partial \beta_D \partial p} = R(\tilde{s}) (1 - v) \Phi(\tilde{s}) \quad \frac{\partial^2 G_H}{\partial \beta_D \partial p} = R(\tilde{s}) (v + (1 - v) \Phi(\tilde{s})),$$

where for ease of notation, we have removed the arguments from the functions G_L and G_H . From the definition of $R(s)$ and from the fact that $\rho(s)$ is strictly increasing and convex in s and $\gamma(s)$ is strictly increasing and strictly concave in s , we know that $R(s) > 0$ for every s . Similarly, by Assumption 2, $\Phi(\tilde{s})$ is also strictly positive. Thus, by Topkis's theorem, it follows that β_{LD} and β_{HD} are strictly increasing in p and hence $\beta_{LD} > \beta_L$ and $\beta_{HD} > \beta_H$.

Next, to show that $\bar{\beta}_{HD} = \min\{\beta_{HD}, \hat{\beta}_D\} > \bar{\beta}_H = \min\{\beta_H, \hat{\beta}\}$, we also need to show that $\hat{\beta}_D > \beta_H$. To do this, we note that if $\beta_D^*(v, \tilde{s}) = \hat{\beta}_D$, then by Lemma 8 it must be that $s_D = \check{s}(\hat{\beta}_D) = \hat{s}_D$. Thus, from Lemma 7, $\hat{\beta}_D$ satisfies the FOC

$$\Pi'_M(\beta_D) = \alpha \gamma'(\check{s}(\beta_D)) (r - w) \frac{R(\check{s}(\beta_D))}{1 - \beta_D} - \delta'(\beta_D) = 0,$$

where $\Pi_M(\beta_D)$ is the manufacturer's profit function if $s_D = \check{s}(\beta_D)$. Evaluating the

above expression for $\beta_D = \hat{\beta}$, we can then derive

$$\begin{aligned}\Pi'_M(\beta_H) &= \alpha \gamma'(\check{s}(\beta_H)) (r - w) \frac{R(\check{s}(\beta_H))}{1 - \beta_H} - \delta'(\beta_H) \\ &> \left(\alpha q \gamma'(\check{s}(\beta_H)) (r - w) \right) (v + (1 - v) \Phi(\check{s}(\beta_H))) \frac{R(\check{s}(\beta_H))}{1 - \beta_H} - \delta'(\beta_H) = 0,\end{aligned}$$

where the last equality follows from the first order condition of $G_H(\beta_D, s_D)$ with respect to β_D in Equation (C.14) evaluated at $\beta_D = \beta_H$ (remember that in the No Manufacturer Disclosure setting, $p = 0$). Similarly, the inequality is given by the fact that $q \leq 1$ and $\Phi(\check{s}) < 1$ (Assumption 3). Thus, from the unimodality of $\Pi_M(\beta_D)$, it follows that $\hat{\beta}_D > \beta_H$ and so $\bar{\beta}_{HD} = \min\{\beta_{HD}, \hat{\beta}_D\} > \bar{\beta}_H = \min\{\beta_H, \hat{\beta}\}$. In addition, since $\hat{\beta}_D$ is independent of p , we also have that $\bar{\beta}_{HD} - \bar{\beta}_H$ is non-decreasing in p .

Finally, we show that $\tau_D > \tau$. As with β_L and β_H , it is enough to show that τ_D is increasing in the penalty factor, p . From Lemma 11 and using the notation therein, it is enough to show that $\frac{\partial \Delta(\tau_D, p)}{\partial p} > 0$. From the definition of $\Delta(\tau_D, x)$ for any model parameter x in Equation (C.31) and from the manufacturer's expected profits in Equation (3.9), we know that

$$\frac{\partial \Delta(\tau_D)}{\partial p} = q \left(E_{s_0} \left[(s_{DL} - s^*(s_0, \beta_{LD}))^+ | \tau_D \right] - E_{s_0} \left[(s_{DH} - s^*(s_0, \bar{\beta}_{HD}))^+ | \tau_D \right] \right). \quad (\text{C.41})$$

Equation (C.41) is almost exactly equal to the difference in expected penalty given by (C.40) in Proposition 2. The proof of $\frac{\partial \Delta(\tau_D, p)}{\partial p} > 0$ follows the same logic therein and is therefore omitted.

Using the above results and taking into account the structure of the manufacturer's optimal choice of β and β_D (see Theorems 2 and 6), we have shown that for any level of visibility, $\beta_D^*(v, \check{s}) > \beta^*(v, \check{s})$, where the inequality is strict since we assume that $v < 1$. \square

Proof of Theorem 4. First, we note that by Corollary 5, $\check{s}(\beta_L) < \tau < \check{s}(\bar{\beta}_H)$ and $\check{s}(\beta_{LD}) < \tau_D < \check{s}(\bar{\beta}_{HD})$. In addition, we know from Theorem 3 that $\tau_D > \tau$. Since the No Manufacturer Disclosure setting corresponds to a special case of the Manu-

facturer Disclosure setting where $p = 0$ and $s_D = \hat{s}$, we can use the structure of the manufacturer's expected profit function in Lemma 6, as well as the structure of the manufacturer's optimal strategy in Theorems 2, 3, and 6, to identify the following cases:

- (i) If $\tilde{s} \in [m, \tau)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_H(\bar{\beta}_{HD}, s_D)$ and $\Pi_M^*(v, \tilde{s}) = G_H(\bar{\beta}_H, s_D = \hat{s}, p = 0)$. Since the function G_H is independent of \tilde{s} , it follows that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0$.
- (ii) If $\tilde{s} \in [\tau, \min\{\tau_D, \hat{s}\})$, then $\Pi_{MD}^*(v, \tilde{s}) = G_H(\bar{\beta}_{HD}, s_D)$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Since the function G_H is independent of \tilde{s} , it follows that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0 - \frac{dG_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)}{d\tilde{s}} = -\alpha q v \gamma'(\tilde{s}) < 0$.
- (iii) If $\hat{s} < \tau_D$ and $\tilde{s} \in [\hat{s}, \tau_D)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_H(\bar{\beta}_{HD}, s_D)$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. In this case, both functions are independent of \tilde{s} , as the third party does not communicate s to consumers if $s > \hat{s}$ in the No Manufacturer Disclosure setting. Thus, $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0$.
- (iv) If $\hat{s} \geq \tau_D$ and $\tilde{s} \in [\tau_D, \hat{s})$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. We know that in the Manufacturer Disclosure setting $s_D \geq \tilde{s}$ because from Theorem 3, in Region (i), $s_D \geq s_0^*(s_0 = \tilde{s}) \geq \tilde{s}$; and in Region (ii), $s_D = \tilde{s}$. Therefore, we obtain that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = \frac{dG_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})}{d\tilde{s}} - \frac{dG_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)}{d\tilde{s}} = p q v > 0$.
- (v) If $\tilde{s} \in [\max\{\tau_D, \hat{s}\}, \tau_H)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Thus, $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = \frac{dG_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})}{d\tilde{s}} - 0 = \alpha q v \gamma'(\tilde{s}) + p q v$, which is strictly greater than 0.
- (vi) If $\tilde{s} \in [\tau_H, M]$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D = \tau_H, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Thus, $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0$.

Therefore, $\Delta_{\Pi}(v, \tilde{s})$ is constant in \tilde{s} for $\tilde{s} \in [m, \tau)$ (by case (i)); non-increasing in \tilde{s} for $\tilde{s} \in [\tau, \tau_D)$ (by cases (ii)–(iii)); and non-decreasing in \tilde{s} for $\tilde{s} \in [\tau_D, M]$ (by cases (iv)–(vi)). \square

Proof of Corollary 1. This result follows directly from Theorem 4. \square

Proof of Corollary 2. The strategies in this corollary follow directly from Theorem 4 and the continuity of the manufacturer's profit function in the signal \tilde{s} . In particular, since the minimum of the difference in expected profits between the two settings, $\Delta_{\Pi}(v, \tilde{s})$, occurs when $\tilde{s} = \tau_D$, we can define the manufacturer's optimal disclosure strategy as follows:

- (a) If $\Delta_{\Pi}(v, \tau_D) \geq 0$, then disclose for all \tilde{s} .
- (b) If $\max_{\tilde{s}} \Delta_{\Pi}(v, \tilde{s}) < 0$, then do not disclose for any \tilde{s} .
- (c) If $\Delta_{\Pi}(v, m) < 0$, $\Delta_{\Pi}(v, \tau_D) < 0$, and $\Delta_{\Pi}(v, M) \geq 0$, then there exists a threshold κ such that it is optimal to disclose only if $\tilde{s} > \kappa$; i.e., it is optimal not to disclose if $\tilde{s} \leq \kappa$. Because of the continuity of $\Delta_{\Pi}(v, \tilde{s})$ in \tilde{s} , such a threshold κ must be greater than τ_D .
- (d) If $\Delta_{\Pi}(v, m) \geq 0$, $\Delta_{\Pi}(v, \tau_D) < 0$, and $\Delta_{\Pi}(v, M) < 0$, then there exists a threshold κ' such that it is optimal to disclose only if $\tilde{s} \leq \kappa'$. Because of the continuity of $\Delta_{\Pi}(v, \tilde{s})$ in \tilde{s} , such a threshold κ' must be less than τ_D .
- (e) If $\Delta_{\Pi}(v, m) \geq 0$, $\Delta_{\Pi}(v, \tau_D) < 0$, and $\Delta_{\Pi}(v, M) \geq 0$, then there exist $\kappa_L < \tau_D$ and $\kappa_H > \tau_D$ such that it is optimal to disclose only if $\tilde{s} \in (\kappa_L, \kappa_H)$.

Finally, from the proof of Theorem 4, it is easy to see that the absolute value of $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}}$ is strictly increasing in v whenever $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} \neq 0$. That is, when the difference in expected profits increases or decreases in \tilde{s} , it does so at a faster rate for higher values of v . While this does not guarantee that the difference between the minimum and maximum values of $\Delta_{\Pi}(v, \tilde{s})$ increases with every increase of \tilde{s} (particularly since τ_D also depends on the value of \tilde{s}), it does ensure that strategies (i) and (ii) are more likely to be optimal under low levels of visibility; and strategies (iii)–(v) are more likely to be optimal under high levels of visibility. In particular, as $v \rightarrow 0$, $\beta_{HD} \rightarrow \beta_{LD}$ (see Proposition 8), and thus $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} \rightarrow 0$; i.e., the difference between both disclosure strategies becomes close to constant in \tilde{s} . \square

Lemma 12. *Let us define $\Pi_M^*(v, \tilde{s})$ and $\Pi_{MD}^*(v, \tilde{s})$ as the manufacturer's optimal expected profits in the No Manufacturer Disclosure and Manufacturer Disclosure settings, respectively. For a given signal \tilde{s} and visibility v , if $s_D^*(v, \tilde{s}) < \hat{s}$, then $\Pi_M^*(v, \tilde{s}) > \Pi_{MD}^*(v, \tilde{s})$.*

Proof. Since the No Manufacturer Disclosure setting corresponds to a special case of the Manufacturer Disclosure setting where $p = 0$ and $s_D = \hat{s}$, from the definition of the manufacturer's expected profits in Equation (3.4), we know that

$$\Pi_M(v, \tilde{s}, \beta) = \left(D_{SN} + \mathbb{E}_{s_0} \left[D_{SC}(s_D = \hat{s}, s^*(s_0, \beta)) \mid \tilde{s} \right] \right) (r - w) - \delta(\beta). \quad (\text{C.42})$$

Similarly, in the Manufacturer Disclosure case we know that if $s_D^*(v, \tilde{s}) < \hat{s}$, then

$$\begin{aligned} \Pi_{MD}^*(v, \tilde{s}) &= \left(D_{SN} + \mathbb{E}_{s_0} \left[D_{SC}(s_D^*, s^*(s_0, \beta_D^*)) \mid \tilde{s} \right] \right) (r - w) - \delta(\beta_D^*) \\ &\quad - p q \mathbb{E}_{s_0} \left[(s_D^* - s^*(s_0, \beta_D^*))^+ \mid \tilde{s} \right] \\ &< \left(D_{SN} + \mathbb{E}_{s_0} \left[D_{SC}(\hat{s}, s^*(s_0, \beta_D^*)) \mid \tilde{s} \right] \right) (r - w) - \delta(\beta_D^*) \\ &\quad - p q \mathbb{E}_{s_0} \left[(s_D^* - s^*(s_0, \beta_D^*, s_D^*))^+ \mid \tilde{s} \right] \end{aligned} \quad (\text{C.43})$$

$$< \Pi_M(v, \tilde{s}, \beta = \beta_D^*) - p q \mathbb{E}_{s_0} \left[(s_D^* - s^*(s_0, \beta_D^*, s_D^*))^+ \mid \tilde{s} \right] \quad (\text{C.44})$$

$$< \Pi_M^*(v, \tilde{s}), \quad (\text{C.45})$$

where $\Pi_M^*(v, \tilde{s}) = \Pi_M(v, \tilde{s}, \beta^*)$. Equation (C.43) follows from $s_D^* = s_D^*(v, \tilde{s}) < \hat{s}$ and the fact that the expected demand from SC consumers is strictly increasing in s_D and Equation (C.44) uses the result in Equation (C.42). Finally, the inequality in Equation (C.45) follows directly from the following observations: (i) β_D^* is a feasible solution to the manufacturer's optimization problem in the No Manufacturer Disclosure setting and hence $\Pi_M(v, \tilde{s}, \beta = \beta_D^*) \leq \Pi_M^*(v, \tilde{s})$, and (ii) $p > 0$. \square

Proof of Theorem 7. This proof follows the same logic as the proof of Theorem 4.

In particular, since $\tau_H < \hat{s}$ implies that $\tau_D < \hat{s}$, we can identify the following cases:

- (i) If $\tilde{s} \in [m, \tau)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_H(\bar{\beta}_{HD}, s_D)$ and $\Pi_M^*(v, \tilde{s}) = G_H(\bar{\beta}_H, s_D = \hat{s}, p = 0)$. Since the function G_H is independent of \tilde{s} , it follows that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0$.
- (ii) If $\tilde{s} \in [\tau, \tau_D)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_H(\bar{\beta}_{HD}, s_D)$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Since the function G_H is independent of \tilde{s} , it follows that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0 - \frac{dG_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)}{d\tilde{s}} = -\alpha q v \gamma'(\tilde{s}) < 0$.
- (iii) If $\tilde{s} \in [\tau_D, \tau_H)$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. We know that in the Manufacturer Disclosure setting $s_D \geq \tilde{s}$ because from Theorem 3, in Region (i), $s_D \geq s_0^*(s_0 = \tilde{s}) \geq \tilde{s}$; and in Region (ii), $s_D = \tilde{s}$. Therefore, we obtain that $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = \frac{dG_L(\beta_{LD}, s_D \geq \tilde{s}, \tilde{s})}{d\tilde{s}} - \frac{dG_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)}{d\tilde{s}} = p q v > 0$.
- (iv) If $\tilde{s} \in [\tau_H, \hat{s})$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D = \tau_H, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Thus, $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0 - \alpha q v \gamma'(\tilde{s}) < 0$.
- (v) $\tilde{s} \in [\hat{s}, M]$, then $\Pi_{MD}^*(v, \tilde{s}) = G_L(\beta_{LD}, s_D = \tau_H, \tilde{s})$ and $\Pi_M^*(v, \tilde{s}) = G_L(\beta_L, s_D = \hat{s}, \tilde{s}, p = 0)$. Thus, $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}} = 0$.

Therefore, $\Delta_{\Pi}(v, \tilde{s})$ is constant in \tilde{s} for $\tilde{s} \in [m, \tau)$ (by case (i)); decreasing in \tilde{s} for $\tilde{s} \in [\tau, \tau_D)$ (by case (ii)); increasing in \tilde{s} for $\tilde{s} \in [\tau_D, \tau_H)$ (by case (iii)); decreasing in \tilde{s} for $\tilde{s} \in [\tau_H, \hat{s})$ (by case (iv) above); and constant in \tilde{s} for $\tilde{s} \in [\hat{s}, M]$ (by case (v)).

In addition, since $\tau_H < \hat{s}$, we know from Lemma 12 that the manufacturer prefers not to disclose if $\tilde{s} = \tau_H$, as the optimal value of s_D at this signal value is $s_D = \tau_H$ (see Theorem 3). From the structure of $\frac{d\Delta_{\Pi}(v, \tilde{s})}{d\tilde{s}}$ outlined above, we also know that $\Delta_{\Pi}(v, \tilde{s})$ is maximized at $\tilde{s} = \tau_H$ if the signal is restricted to the interval $\tilde{s} \in [\tau_D, M]$. Thus, by the continuity of $\Delta_{\Pi}(v, \tilde{s})$ in \tilde{s} , there exists a threshold $\bar{\tau} \in [m, \tau_D]$ such that it is optimal for the manufacturer not to disclose if $\tilde{s} \geq \bar{\tau}$.

Similarly, if $s_D^*(v, \tilde{s} = m) < \hat{s}$, then by Lemma 12 it is optimal for the manufacturer not to disclose SR at $\tilde{s} = m$. From cases (i) and (ii) in the proof of Theorem 4, $\Delta_{\Pi}(v, \tilde{s})$ is maximized at $\tilde{s} = m$ if the signal is restricted to the interval $\tilde{s} \in [m, \tau_D)$.

Thus, combined with the result in the above paragraph, if $s_D^*(v, \tilde{s} = m) < \hat{s}$, then it is optimal for the manufacturer not to disclose for any signal \tilde{s} .

Therefore, the manufacturer's potential strategies are limited to either not disclose for any signal \tilde{s} ; or to disclose only if \tilde{s} is greater than or equal to a threshold signal $\bar{\tau}$. From the above discussion, we note that the latter strategy requires $s_D^*(v, \tilde{s} = m) \geq \hat{s} > \tau_H$; i.e., by Theorem 3 in the Manufacturer Disclosure setting the manufacturer would need to prefer disclosing a greater value of s_D^* if $\tilde{s} = m$ than if $\tilde{s} = M$. \square

C.5 The Effects on Social Responsibility

Finally, we show our results regarding the impact of visibility and third-party scrutiny on the supplier's SR level.

Proof of Proposition 3. First, we note from the proof of Theorem 5 that the supplier's optimal choice of SR level does not depend on s_D if $\check{s}(\beta_D) \leq s_D$. Also by Proposition 13, we know that the manufacturer never offers β_D such that $\check{s}(\beta_D) > s_D$. Thus, to study the effect of q on s , it is sufficient to study its direct effect and the effect that q has on β_D .

From Theorem 6, β_D^* can take one of three values: β_{LD} , β_{HD} , or $\hat{\beta}_D$. In the first case, by Lemma 3 we know that β_{LD} satisfies the first order condition

$$\begin{aligned} U\left(q, \beta_{LD}(q), \check{s}(q, \beta_{LD}(q))\right) &\equiv (\alpha q (r - w) \gamma'(\check{s}) + pq) \frac{R(\check{s})}{1 - \beta_D} (1 - v) \Phi(\check{s}) - \delta'(\beta_D) \\ &= 0. \end{aligned} \tag{C.46}$$

For ease of notation, we write this condition simply as $U = 0$. Applying the Implicit

Function Theorem to this condition, we obtain $\frac{dU}{dq} = 0$, or equivalently,

$$\begin{aligned}\frac{\partial U}{\partial q} + \frac{\partial U}{\partial s} \frac{d\check{s}}{dq} + \frac{\partial U}{\partial \beta_D} \frac{d\beta_{LD}}{dq} &= 0 \\ \frac{\partial U}{\partial q} + \frac{\partial U}{\partial s} \left(\frac{\partial \check{s}}{\partial q} + \frac{\partial \check{s}}{\partial \beta_D} \frac{d\beta_{LD}}{dq} \right) + \frac{\partial U}{\partial \beta_D} \frac{d\beta_{LD}}{dq} &= 0 \\ \frac{\partial U}{\partial q} + \frac{\partial U}{\partial s} \frac{\partial \check{s}}{\partial q} + \frac{dU}{d\beta_D} \frac{d\beta_{LD}}{dq} &= 0.\end{aligned}$$

Therefore, since $\frac{\partial U}{\partial q} = (\alpha(r-w)\gamma'(\check{s}) + p) \frac{R(\check{s})}{1-\beta_D} (1-v)\Phi(\check{s}) > 0$ and $\frac{dU}{d\beta_D} < 0$ at the optimal solution $\beta_D = \beta_{LD}$ (see the proof of Lemma 3), we can rearrange terms and obtain

$$\frac{d\beta_{LD}}{dq} > - \left(\frac{dU}{d\beta_D} \right)^{-1} \frac{\partial U}{\partial s} \frac{\partial \check{s}}{\partial q}. \quad (\text{C.47})$$

In addition, we also know that $\frac{dU}{d\beta_D} = \frac{\partial U}{\partial \beta_D} + \frac{\partial U}{\partial s} \frac{d\check{s}}{d\beta_D}$, and hence $\frac{\partial U}{\partial s} \frac{d\check{s}}{d\beta_D} = \frac{dU}{d\beta_D} - \frac{\partial U}{\partial \beta_D}$. Similarly, we know that $\frac{d\check{s}}{dq} = \frac{\partial \check{s}}{\partial q} + \frac{\partial \check{s}}{\partial \beta_D} \frac{d\beta_{LD}}{dq}$. Therefore, replacing terms, we get the following inequality:

$$\begin{aligned}\frac{d\check{s}}{dq} &> \frac{\partial \check{s}}{\partial q} - \frac{d\check{s}}{d\beta_D} \left(\frac{dU}{d\beta_D} \right)^{-1} \frac{\partial U}{\partial s} \frac{\partial \check{s}}{\partial q} \\ &> \frac{\partial \check{s}}{\partial q} \left(1 - \left(\frac{dU}{d\beta_D} \right)^{-1} \left(\frac{dU}{d\beta_D} - \frac{\partial U}{\partial \beta_D} \right) \right) \\ &> \frac{\partial \check{s}}{\partial q} \left(\frac{dU}{d\beta_D} \right)^{-1} \frac{\partial U}{\partial \beta_D}.\end{aligned} \quad (\text{C.48})$$

From Lemma 1, we know that $\frac{\partial \check{s}}{\partial q} > 0$. Similarly, we also know that $\frac{dU}{d\beta} < 0$ and, from the definition of the function U in Equation (C.46), $\frac{\partial U}{\partial \beta} = -\delta''(\beta_{LD}) < 0$, as $\delta(\beta_D)$ is strictly convex in β_D . Therefore, we conclude that $\frac{d\check{s}}{dq} > 0$.

The same logic can be applied to the cases where $\beta_D^* = \beta_{HD}$, defining U using the first order condition in Equation (C.13); or where $\beta_D^* = \hat{\beta}_D$, defining U using the first order condition in Equation (C.25). The only difference is that in the last case, the inequality in Equation (C.48) is replaced by an equality.

Thus, we have shown that if $\beta_D = \beta_{LD}$ or $\beta_D = \bar{\beta}_{HD} = \min\{\beta_{HD}, \hat{\beta}_D\}$, then in any disclosure setting, the supplier's final SR level is non-decreasing in q . Finally, by Proposition 9 we also know that τ_D is increasing in q , which means that for a given

signal, an increase in q never causes β_D^* to shift from $\bar{\beta}_{HD}$ to β_{LD} for any signal. \square

Proof of Proposition 7. This result follows directly from the following observations. First, from Corollary 4, the supplier's choice of SR depends only on β_D . In addition, Proposition 4 shows that for any \tilde{s} and v , $\beta_D > \beta$. Combining these two results and using (i) the structure of β^* and β_D^* in Theorems 2 and 6, respectively, and (ii) by Proposition 4, $\tau_D > \tau$, we know that for any level of visibility $v \in (0, 1)$, $\beta_D^* = \beta_{HD} > \beta^*$ for every $\tilde{s} \in [m, \tau_D]$ and that $\beta_D^* = \beta_{LD} \geq \beta^* = \beta_L$ for $\tilde{s} \in (\tau_D, M]$.

Therefore, for any signal \tilde{s} , $s^*(s_0, \beta_D, s_D) \geq s^*(s_0, \beta)$. Finally, we can define $\tau_{SR} = \check{s}(\beta_L)$ such that for all $s_0 \leq \tau_{SR}$ and for any value of β or β_D that the manufacturer may consider (i.e., β_L , β_H , β_{LD} , or β_{HD}), the supplier's best response is equal to $s^* = \check{s}(\beta_D)$ (or $s^* = \check{s}(\beta)$ in the No Manufacturer Disclosure setting). This, as by Theorem 6, $\beta_H > \beta_L$; by Proposition 4, for any \tilde{s} and v , $\beta_D^* > \beta^*$; and by Theorem 5, the supplier's best response is equal to $\min\{s_0, \check{s}(\beta_D)\}$ (or $\min\{s_0, \check{s}(\beta)\}$ in the No Manufacturer Disclosure setting). As a result, since $\check{s}(\beta_D)$ is strictly increasing in β_D and $\beta_D^* > \beta^*$, it follows that for any supplier type $s_0 \leq \tau_{SR}$ and for any \tilde{s} and v , the supplier's final level of SR is strictly greater in the Manufacturer Disclosure setting than in the No Manufacturer Disclosure setting. \square

Proof of Lemma 11. First, we note that given s_0 , the final level of SR is equal to $s^*(s_0, \beta^*(v, \tilde{s} = s_0))$ if the signal is correct (which happens with probability v); and equal to the expected value of $s^*(s_0, \beta^*(v, \tilde{s}))$ over all possible values of \tilde{s} if the signal is not correct (which happens with probability $1 - v$). Thus, the expected SR level given v and s_0 can be written as

$$ES(v, s_0) \equiv v s^*(s_0, \beta^*(v, \tilde{s} = s_0)) + (1 - v) \int_m^M s^*(s_0, \beta^*(v, \tilde{s})) \phi(\tilde{s}) d\tilde{s}. \quad (\text{C.49})$$

Let $v_2 \leq 1$ be the level of visibility such that $\tau(v_2) = l_2$, where $l_2 \equiv \max_{v \in [0, 1]} \{\tau(v)\}$. From the definition of l_2 , $\tau(v_2) \geq \tau(v)$ for all v . By Corollary 5, we know that for any v , $\check{s}(\beta_L(v)) < \tau(v) < \check{s}(\bar{\beta}_H(v))$. Also, Proposition 1 shows that β_L and $\bar{\beta}_H$ are

decreasing and non-decreasing in v , respectively. Consequently, since $\check{s}(\beta)$ is increasing in β , $\check{s}(\beta_L)$ and $\check{s}(\bar{\beta}_H)$ are also decreasing and non-decreasing in v . Combining these results, we obtain that $\check{s}(\beta_L(v)) < l_2 < \check{s}(\bar{\beta}_H(v)) \forall v \geq v_2$. Since both inequalities are strict, $\check{s}(\beta)$ is continuous in β , and $\beta_L(v)$ and $\bar{\beta}_H(v)$ are continuous in v , there must also exist $\epsilon > 0$ such that $\check{s}(\beta_L(v)) < l_2 < \check{s}(\bar{\beta}_H(v)) \forall v \geq \hat{v} \equiv v_2 - \epsilon$, with $\hat{v} < 1$. As a result, for any $v \geq \hat{v}$ and $s_0 \in [l_2, \check{s}(\bar{\beta}_H(v))]$, the expected level of SR in Equation (C.49) can be rewritten as:

$$\text{ES}(v, s_0) = v s_0 + (1 - v) \left(\Phi(\tau(v)) \check{s}(\bar{\beta}_H(v)) + (1 - \Phi(\tau(v))) s_0 \right), \quad (\text{C.50})$$

since (i) $s^*(s_0, \beta) = \max \{s_0, \min \{\hat{s}, \check{s}(\beta)\}\}$, or equivalently, $s^*(s_0, \beta) = \max \{s_0, \check{s}(\beta)\}$ if $\beta \leq \hat{\beta}$; and (ii) the manufacturer selects $\beta^* = \bar{\beta}_H(v)$ for any signal $\tilde{s} \leq \tau(v)$ and $\beta^* = \beta_L(v)$ for any signal $\tilde{s} > \tau(v)$.

For ease of notation, let us write $\tau(v)$, $\beta_L(v)$, and $\beta_H(v)$ simply as τ , β_L , and β_H ; and let us define $s_L = \check{s}(\beta_L)$ and $s_H = \check{s}(\bar{\beta}_H)$. Using this notation and the result in Equation (C.50), for any $v \geq \hat{v}$ and $s_0 \in [l_2, \check{s}(\bar{\beta}_H(v))]$, the partial derivative of the expected level of SR with respect to v can be written as

$$\begin{aligned} \frac{d\text{ES}(v, s_0)}{dv} &= s_0 - \left(\Phi(\tau) s_H + (1 - \Phi(\tau)) s_0 \right) \\ &\quad + (1 - v) \left(\phi(\tau) \frac{d\tau}{dv} s_H + \Phi(\tau) \frac{ds_H}{dv} - \phi(\tau) \frac{d\tau}{dv} s_0 \right) \\ &= \Phi(\tau) (s_0 - s_H) + (1 - v) \left(\phi(\tau) \frac{d\tau}{dv} s_H + \Phi(\tau) \frac{ds_H}{dv} - \phi(\tau) \frac{d\tau}{dv} s_0 \right). \end{aligned} \quad (\text{C.51})$$

Thus, $\frac{d\text{ES}(v, s_0)}{dv} \rightarrow \Phi(\tau) (s_0 - s_H) < 0$ as $v \rightarrow 1$. There must therefore exist $\bar{v} < 1$ such that for all $v \geq \bar{v}$, if $s_0 \in [l_2, \check{s}(\bar{\beta}_H(v))]$, then $\frac{d\text{ES}(v, s_0)}{dv} < 0$. \square

Bibliography

- Agrawal, V., D. Lee. 2016. The effect of sourcing policies on a supplier's sustainable practices. Working paper, Georgetown University, Washington, DC.
- Akerlof, G. 1982. Labor contracts as partial gift exchange. *Quart. J. Econom.* **97**(4) 543–569.
- Alexander, R. D. 1987. *The Biology of Moral Systems*. Transaction Publishers, New Brunswick, New Jersey.
- Anctil, R., J. Dickhaut, C. Kanodia, B. Shapiro. 2004. Information transparency and coordination failure: Theory and experiment. *Journal of Accounting Research* **42**(2) 159–195.
- Anderson, Lisa R, Jennifer M Mellor, Jeffrey Milyo. 2004. Social capital and contributions in a public-goods experiment. *The American Economic Review* **94**(2) 373–376.
- Andersson, O., H.J. Holm, J.R. Tyran, E. Wengstrom. 2016. Deciding for others reduces loss aversion. *Management Science* **62**(1) 29–36.
- Andreoni, J., J. Miller. 2002. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica* **70**(2) 737–753.
- Angrist, Joshua D, Jörn-Steffen Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Arora, N., T. Henderson. 2007. Embedded premium promotion: Why it works and how to make it more effective. *Marketing Science* **26**(4) 514–531.
- Au, W., J. Kwong. 2004. Measurements and effects of social-value orientation in social dilemmas. *Contemporary Psychological Research on Social Dilemmas*. Cambridge University Press New York, NY, 71–98.
- Babcock, L., G. Loewenstein. 1997. Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives* **11**(1) 109–126.
- Babcock, L., X. Wang, G. Loewenstein. 1996. Choosing the wrong pond: Social comparisons in negotiations that reflect a self-serving bias. *Quarterly Journal of Economics* **111**(1) 1–19.
- Babich, V. 2010. Independence of capacity ordering and financial subsidies to risky suppliers. *Manufacturing & Service Operations Management* **12**(4) 583–607.
- Babich, V., C.S. Tang. 2012. Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing & Service Operations Management* **14**(2) 301–314.
- Bagnoli, M., S.G. Watts. 2003. Selling to socially responsible consumers: Competition and the private provision of public goods. *Journal of Economics & Management Strategy* **12**(3) 419–445.
- Bartling, B., R. Weber, L. Yao. 2015. Do markets erode social responsibility? *Quarterly Journal of Economics* **130**(1) 219–266.
- Becker, G. M., M. H. DeGroot, J. Marschak. 1964. Measuring utility by a single-response sequential

- method. *Behav. Sci.* **9**(3) 226–232.
- Bénabou, Roland, Jean Tirole. 2006. Incentives and prosocial behavior. *The American economic review* **96**(5) 1652–1678.
- Bloomfield, Robert J, T Jeffrey Wilks. 2000. Disclosure effects in the laboratory: Liquidity, depth, and the cost of capital. *The Accounting Review* **75**(1) 13–41.
- Blount, S. 1995. When social outcomes aren't fair: The effect of causal attributions on preferences. *Organ. Behav. Human Decis. Process.* **63**(2) 131–144.
- Bohm, Peter, Johan Lindén, Joakim Sonnegård. 1997. Eliciting reservation prices: Becker–degroot–marschak mechanisms vs. markets. *The Economic Journal* **107**(443) 1079–1089.
- Bolton, G., B. Greiner, A. Ockenfels. 2013. Engineering trust: Reciprocity in the production of reputation information. *Management Science* **59**(2) 265–285.
- Bolton, G., E. Katok. 1998. An experimental test of the crowding out hypothesis: The nature of beneficent behavior. *Journal of Economic Behavior & Organization* **37**(3) 315–331.
- Bolton, G. E., A. Ockenfels. 2000. ERC: A theory of equity, reciprocity, and competition. *Amer. Econom. Rev.* **90**(1) 166–193.
- Bolton, L.E., J.W. Alba. 2006. Price fairness: Good and service differences and the role of vendor costs. *Journal of Consumer Research* **33**(2) 258–265.
- Borzaga, C., L. Becchetti. 2010. *The Economics of Social Responsibility: The World of Social Enterprises*. Routledge.
- Brennan, G., L. G. González, W. Güth, M. V. Levati. 2008. Attitudes toward private and collective risk in individual and strategic choice situations. *J. Econom. Behav. Organ.* **67**(1) 253–262.
- Brock, J. M., A. Lange, E. Y. Ozbay. 2013. Dictating the risk: Experimental evidence on giving in risky environments. *Amer. Econom. Rev.* **103**(1) 415–437.
- Buell, R.W., T. Kim, C. Tsay. 2016. Creating reciprocal value through operational transparency. *Management Science*, forthcoming.
- Buell, R.W., M.I. Norton. 2011. The labor illusion: How operational transparency increases perceived value. *Management Science* **57**(9) 1564–1579.
- Chakravarty, S., G.W. Harrison, E.E. Haruvy, E.E. Rulstrom. 2011. Are you risk averse over other people's money? *Southern Economic Journal* **77**(4) 901–913.
- Charness, G. 2000. Responsibility and effort in an experimental labor market. *J. Econom. Behav. Organ.* **42**(3) 375–384.
- Charness, G. 2004. Attribution and reciprocity in an experimental labor market. *J. Labor Econom.* **22**(3) 665–688.
- Charness, G., E. Haruvy. 2002. Altruism, equity, and reciprocity in a gift-exchange experiment: An encompassing approach. *Games & Economic Behavior* **40**(2) 203–231.
- Charness, G., M. Rabin. 2002. Understanding social preferences with simple tests. *Quart. J. Econom.* **117**(3) 817–869.
- Charness, Gary, Martin Dufwenberg. 2006. Promises and partnership. *Econometrica* **74**(6) 1579–1601.
- Chen, L., H.L. Lee. 2016. Sourcing under supplier responsibility risk: The effects of certification, audit, and contingency payment. *Management Science*. Forthcoming.
- Chen, S., Q. Zhang, Y. Zhou. 2015. Impact of supply chain transparency on sustainability under NGO scrutiny. Working paper, University of Washington, Seattle, WA.

- Cho, S.H., X. Fang, S. Tayur, Y. Xu. 2016. Combating child labor: Incentives and information transparency in global supply chains. Working paper, Carnegie Mellon University, Pittsburgh, PA.
- Cohen, Taya R, Brian C Gunia, Sun Young Kim-Jun, J Keith Murnighan. 2009. Do groups lie more than individuals? honesty and deception as a function of strategic self-interest. *Journal of Experimental Social Psychology* **45**(6) 1321–1324.
- Cone Communications/Echo. 2013. Global CSR Study. Available at <http://www.conecomm.com/global-csr-study>.
- Corbett, C.J., G.A. DeCroix. 2001. Shared-savings contracts for indirect materials in supply chains: Channel profits and environmental impacts. *Management Science* **47**(7) 881–893.
- Cox, J. C. 2004. How to identify trust and reciprocity. *Games Econom. Behav.* **46**(2) 260–281.
- Croson, Rachel, Karen Donohue. 2003. Impact of pos data sharing on supply chain management: An experimental study. *Production and Operations management* **12**(1) 1–11.
- Croson, Rachel, Karen Donohue. 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management science* **52**(3) 323–336.
- Dahlsrud, A. 2008. How corporate social responsibility is defined: An analysis of 37 definitions. *Corp. Soc. Responsib. Environ. Management* **15**(1) 1–13.
- Dana, J., R. A. Weber, J. X. Kuang. 2007. Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Econ. Theory* **33**(1) 67–80.
- de Kwaadsteniet, E. W., E. van Dijk, A. Wit, D. de Cremer. 2006. Social dilemmas as strong versus weak situations: Social value orientations and tacit coordination under resource size uncertainty. *J. Exp. Soc. Psych.* **42**(4) 509–516.
- De Pelsmacker, P., L. Driesen, G. Rayp. 2005. Do consumers care about ethics? Willingness to pay for fair-trade coffee. *Journal of Consumer Affairs* **39**(2) 363–385.
- Delucchi, K. L., A. Bostrom. 2004. Methods for analysis of skewed data distributions in psychiatric clinical studies: Working with many zero values. *Am. J. Psychiat.* **161**(7) 1159–1168.
- Devinney, T. M., P. Auger, G. M. Eckhardt. 2010. *The myth of the ethical consumer*. Cambridge University Press, Cambridge, UK.
- Doorey, D. J. 2011. The transparent supply chain: From resistance to implementation at Nike and Levi-Strauss. *J. Bus. Ethics* **103**(4) 587–603.
- Dreber, Anna, Magnus Johannesson. 2008. Gender differences in deception. *Economics Letters* **99**(1) 197–199.
- Egels-Zandén, Niklas, Niklas Hansson. 2016. Supply chain transparency as a consumer or corporate tool: The case of nudie jeans co. *Journal of Consumer Policy* **39**(4) 377–395.
- Engel, C. 2011. Dictator games: A meta study. *Experimental Economics* **14**(4) 583–610.
- Engelmann, D., U. Fischbacher. 2009. Indirect reciprocity and strategic reputation building in an experimental helping game. *Games Econom. Behav.* **67**(2) 399–407.
- Evans, Anthony M, William Revelle. 2008. Survey and behavioral measurements of interpersonal trust. *Journal of Research in Personality* **42**(6) 1585–1593.
- EY, UN Global Compact. 2016. The state of sustainable supply chains: Building responsible and resilient supply chains. Available at <https://www.unglobalcompact.org/library/4451>.
- Falk, A., E. Fehr, U. Fischbacher. 2008. Testing theories of fairness – Intentions matter. *Games Econom. Behav.* **62**(1) 287–303.

- Falk, A., U. Fischbacher. 2006. A theory of reciprocity. *Games & Economic Behavior* **54**(2) 293–315.
- Fashion Revolution. 2016. Fashion Transparency Index. Available at http://fashionrevolution.org/wp-content/uploads/2016/04/FR_FashionTransparencyIndex.pdf.
- Fehr, E., U. Fischbacher. 2004. Third-party punishment and social norms. *Evol. Human Behav.* **25**(2) 63–87.
- Fehr, E., K. M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quart. J. Econom.* 817–868.
- Fehr, E., K. M. Schmidt. 2006. The economics of fairness, reciprocity and altruism – Experimental evidence and new theories. *Handbook of the Economics of Giving, Altruism and Reciprocity*, vol. 1. Elsevier Science, 615–691.
- Feinberg, F., T. Kinnear, J. Taylor. 2012. *Modern Marketing Research: Concepts, Methods, and Cases*. 2nd ed. Cengage Learning, Boston, MA.
- Fischbacher, U. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experiment. Econom.* **10**(2) 171–178.
- Forsythe, R., J. L. Horowitz, N. E. Savin, M. Sefton. 1994. Fairness in simple bargaining experiments. *Games Econom. Behav.* **6**(3) 347–369.
- Gächter, Simon, Elke Renner. 2010. The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics* **13**(3) 364–377.
- García-Gallego, A., N. Georgantzís. 2011. Good and bad increases in ecological awareness: Environmental differentiation revisited. *Strateg. Behav. Environ.* **1**(1) 71–88.
- Gillai, B., A.H. Porteous, S.V. Rammohan, H.L. Lee. 2013. The relationship between responsible supply chain practices and performance. White paper, Stanford University, Stanford, CA.
- Gino, F., M. Norton, R. Weber. 2016. Motivated Bayesians: Feeling moral while acting egoistically. *Journal of Economic Perspectives* **30**(3) 189–212.
- Glaeser, Edward L, David I Laibson, Jose A Scheinkman, Christine L Soutter. 2000. Measuring trust. *The Quarterly Journal of Economics* **115**(3) 811–846.
- Gneezy, A., U. Gneezy, L. Nelson, A. Brown. 2010. Shared social responsibility: A field experiment in pay-what-you-want pricing and charitable giving. *Science* **329**(5989) 325–327.
- Gneezy, Uri. 2005. Deception: The role of consequences. *The American Economic Review* **95**(1) 384–394.
- Goldstein, W., E. Weber. 1995. Content and discontent: Indications and implications of domain specificity in preferential decision making. *Psychology of Learning & Motivation* **32** 83–136.
- Gummerum, Michaela, Yaniv Hanoach, Monika Keller, Katie Parsons, Alegra Hummel. 2010. Preschoolers’s allocations in the dictator game: The role of moral emotions. *Journal of Economic Psychology* **31**(1) 25–34.
- Guo, R., H.L. Lee, R. Swinney. 2016. Responsible sourcing in supply chains. *Management Science* **62**(9) 2722–2744.
- Güth, W., S. Huck, P. Ockenfels. 1996. Two-level ultimatum bargaining with incomplete information: An experimental study. *Econom. J.* **106**(436) 593–604.
- Hainmueller, J., M. J. Hiscox, S. Sequeira. 2015. Consumer demand for the Fair Trade label: Evidence from a multi-store field experiment. *Review of Economics & Statistics* **97**(2) 242–256.
- Haisley, E. C., R. A. Weber. 2010. Self-serving interpretations of ambiguity in other-regarding behavior. *Games Econom. Behav.* **68**(2) 614–625.

- Halevy, Y. 2007. Ellsberg revisited: An experimental study. *Econometrica* **75**(2) 503–536.
- Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwina Gwako, Natalie Henrich, et al. 2006. Costly punishment across human societies. *Science* **312**(5781) 1767–1770.
- Holt, C. A., S. K. Laury. 2002. Risk aversion and incentive effects. *Amer. Econom. Rev.* **92**(5) 1644–1655.
- HP. 2016. Supply chain responsibility. Available at <http://www8.hp.com/us/en/hp-information/global-citizenship/society/supplychain.html>.
- Huang, L., J.S. Song, R. Swinney. 2016. Managing social responsibility in multitier supply chains. Working paper, Duke University, Durham, NC.
- Inderfurth, Karl, Abdolkarim Sadrieh, Guido Voigt. 2013. The impact of information sharing on supply chain performance under asymmetric information. *Production and Operations Management* **22**(2) 410–425.
- Inditex. 2015. Annual report. Available at https://www.inditex.com/investors/investors_relations/annual_report.
- Joshi, Yatish, Zillur Rahman. 2015. Factors affecting green purchase behaviour and future research directions. *International Strategic Management Review* **3**(1) 128–143.
- Kagel, J., A. Roth. 1995. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.
- Kahneman, D., J. Knetsch, R. Thaler. 1986. Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* **76**(4) 728–741.
- Kalkanci, B., E. Ang, E.L. Plambeck. 2016. Strategic disclosure of social and environmental impacts in a supply chain. A. Atasu, ed., *Environmentally Responsible Supply Chains*. Springer, New York, NY, 223–240.
- Kalkanci, B., E.L. Plambeck. 2015. Measuring and improving supplier responsibility under voluntary vs. mandatory disclosure. Working paper, Georgia Institute of Technology, Atlanta, GA.
- Kanagaretnam, K., S. Mestelman, S.M.K. Nainar, M. Shehata. 2010. Trust and reciprocity with transparency and repeated interactions. *Journal of Business Research* **63**(3) 241–247.
- Kanagaretnam, Kiridaran, Stuart Mestelman, Khalid Nainar, Mohamed Shehata. 2009. The impact of social value orientation and risk attitudes on trust and reciprocity. *Journal of Economic Psychology* **30**(3) 368–380.
- Kang, Jiyun, Gwendolyn Hustvedt. 2014. Building trust between consumers and corporations: The role of consumer perceptions of transparency and social responsibility. *Journal of Business Ethics* **125**(2) 253–265.
- Karaer, Ö., T. Kraft, J. Khawam. 2017. Buyer and nonprofit levers to improve supplier environmental performance. *Production and Operations Management* **26**(6) 1163–1190.
- Katok, E. 2011. Using laboratory experiments to build better operations management models. *Foundations and Trends in Technology, Information and Operations Management* **5**(1) 1–86.
- Kim, B. 2000. Coordinating an innovation in supply chain management. *European Journal of Operations Research* **123**(3) 568–584.
- Kim, S.H., S. Netessine. 2013. Collaborative cost reduction and component procurement under information asymmetry. *Management Science* **59**(1) 189–206.
- Kimeldorf, H., R. Meyer, M. Prasad, I. Robinson. 2006. Consumers with a conscience: Will they pay more? *Contexts* **5**(1) 24–29.

- Klos, A., E. Weber, M. Weber. 2005. Investment decisions and time horizon: Risk perception and risk behaviour in repeated gambles. *Management Sci.* **51**(12) 1777–1790.
- Kosfeld, M., A. Okada, A. Riedl. 2009. Institution formation in public goods games. *American Economic Review* **99**(4) 1335–1355.
- KPMG and Global Reporting Initiative. 2014. Spotlight on corporate transparency: Insights from GLOBE 2014.
- Kraft, T., L. Valdés, Y. Zheng. 2017. Supply chain visibility and social responsibility: Investigating consumers' behaviors and motives. Working paper, Charlottesville, VA.
- Krawczyk, M., F. Le Lec. 2010. "Give me a chance!" An experiment in social decision under risk. *Experiment. Econom.* **13**(4) 500–511.
- Lachenbruch, P. A. 2002. Analysis of data with excess zeros. *Stat. Methods Med. Res.* **11**(4) 297–302.
- Lang, Mark, Mark Maffett. 2011. Transparency and liquidity uncertainty in crisis periods. *Journal of Accounting and Economics* **52**(2) 101–125.
- Lee, HL, Kevin O'Leary, Geraint John. 2012. The chief supply chain officer report 2012. *SCM World*.
- Letizia, P., G. Hendrikse. 2016. Supply chain structure incentives for corporate social responsibility: An incomplete contracting analysis. *Production and Operations Management* **25**(11) 1919–1941.
- Levine, D. K. 1998. Modeling altruism and spitefulness in experiments. *Rev. Econom. Dynam.* **1**(3) 593–622.
- Lewis, T.R., F. Liu, J.S. Song. 2016. A dynamic mechanism for achieving sustainable quality supply. Working paper, Duke University, Durham, NC.
- Li, C., L.G. Debo. 2009. Strategic dynamic sourcing from competing suppliers with transferable capacity investment. *Naval Research Logistics* **56**(6) 540–562.
- List, J. 2006. The behavioralist meets the market: Measuring social preferences and reputation effects in actual transactions. *Journal of Political Economy* **114**(1) 1–37.
- Liu, S., K.C. So, F. Zhang. 2010. Effect of supply reliability in a retail setting with joint marketing and inventory decisions. *Manufacturing & Service Operations Management* **12**(1) 19–32.
- Locke, R.M., F. Qin, A. Brause. 2007. Does monitoring improve labor standards? lessons from Nike. *Industrial and Labor Relations Review* **6**(1) 3–31.
- Luhan, W., M. Kocher, M. Sutter. 2009. Group polarization in the team dictator game reconsidered. *Experimental Economics* **12**(1) 26–41.
- Lundquist, Tobias, Tore Ellingsen, Erik Gribbe, Magnus Johannesson. 2009. The aversion to lying. *Journal of Economic Behavior & Organization* **70**(1) 81–92.
- MacLean, Richard, Kathee Rebernak. 2007. Closing the credibility gap: The challenges of corporate responsibility reporting. *Environmental Quality Management* **16**(4) 1–6.
- Marks & Spencer. 2015a. M&S Plan A Report. Available at <http://planareport.marksandspencer.com/>.
- Marks & Spencer. 2015b. Plan A Report 2015. Available at <http://corporate.marksandspencer.com/media/6e633a181b124309bab60137c8171017>.
- Marks & Spencer and Globescan. 2015. The M&S Transparency Study: Research Report.
- Marshall, D., L. McCarthy, P. McGrath, F. Harrigan. 2016. What's your strategy for supply chain disclosure? *MIT Sloan Management Review* **57**(2) 37–45.
- Mayer, Roger C, James H Davis. 1999. The effect of the performance appraisal system on trust for

- management: A field quasi-experiment. *Journal of applied psychology* **84**(1) 123.
- Mendoza, A.J., R.T. Clemen. 2013. Outsourcing sustainability: A game-theoretic approach. *Environmental Systems and Decisions* **33**(2) 224–236.
- Mintel Group. 2015. The ethical consumer – U.S. – July 2015. Report overview available at <http://store.mintel.com/the-ethical-consumer-us-july-2015>.
- Moorthy, K.S. 1988. Product and price competition in a duopoly. *Marketing Science* **7**(2) 141–168.
- Morris, S., H.S. Shin. 2004. Coordination risk and the price of debt. *European Economic Review* **48**(1) 133–153.
- Muthulingam, S., C. J. Corbett, S. Benartzi, B. Oppenheim. 2013. Energy efficiency in small and medium-sized manufacturing firms: Order effects and the adoption of process improvement recommendations. *Manufacturing Service Oper. Management* **15**(4) 596–615.
- New, S., D. Brown. 2011. The four challenges of supply chain transparency. *The European Business Review*, July 20, available at <http://www.europeanbusinessreview.com/challenges-supply-chain-transparency/>.
- Nielsen. 2013. Consumers who care – and say they’ll reward companies with their wallets. Available at <http://www.nielsen.com/us/en/insights/reports/2013/consumers-who-care.html>.
- Nielsen. 2015. The sustainability imperative. Available at <https://www.nielsen.com/content/dam/niensenglobal/dk/docs/global-sustainability-report-oct-2015.pdf>.
- Norman, W., C. MacDonald. 2004. Getting to the bottom of “triple bottom line”. *Business Ethics Quarterly* **14**(2) 243–262.
- Nowak, M. A., K. Sigmund. 1998. Evolution of indirect reciprocity by image scoring. *Nature* **393**(6685) 573–577.
- OECD/WTO. 2013. *Aid for Trade at a Glance 2013: Connecting to Value Chains*. OECD Publishing.
- Oosterbeek, H., R. Sloof, G. van de Kuilen. 2004. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics* **7**(2) 171–188.
- Orsdemir, A., B. Hu, V. Deshpande. 2016. Ensuring corporate social and environmental responsibility through vertical integration and horizontal sourcing. Working paper, UC Riverside, Riverside, CA.
- Ortmann, A., G. Gigerenzer. 2000. Reasoning in economics and psychology: Why social context matters. M. Streit, U. Mummert, D. Kiwit, eds., *Cognition, Rationality, and Institutions*. Springer, New York, NY, 131–145.
- Ostrom, E., J. Walker, R. Gardner. 1992. Covenants with and without a sword: Self-governance is possible. *American Political Science Review* **86**(2) 404–417.
- Özer, Özalp, Yanchong Zheng, Kay-Yut Chen. 2011. Trust in forecast information sharing. *Management Science* **57**(6) 1111–1137.
- Özer, Özalp, Yanchong Zheng, Yufei Ren. 2014. Trust, trustworthiness, and information sharing in supply chains bridging china and the united states. *Management Science* **60**(10) 2435–2460.
- Paharia, N., K. Vohs, R. Deshpandé. 2013. Sweatshop labor is wrong unless the shoes are cute: Cognition can both help and hurt moral motivated reasoning. *Organizational Behavior & Human Decision Processes* **121**(1) 81–88.
- Patagonia. 2014. The Footprint Chronicles. Retrieved August 7, 2014, <http://www.patagonia.com/us/footprint>.
- Pigors, M., B. Rockenbach. 2016a. Consumer social responsibility. *Management Science* **62**(11)

3123–3137.

- Pigors, Mark, Bettina Rockenbach. 2016b. The competitive advantage of honesty. *European Economic Review* **89** 407–424.
- Plambeck, E.L., T.A. Taylor. 2016. Supplier evasion of a buyer's audit: Implications for motivating supplier social and environmental responsibility. *Manufacturing & Service Operations Management* **18**(2) 184–197.
- Porteous, A.H., S.V. Rammohan. 2013. Integration, incentives, and innovation: Nike's strategy to improve social and environmental conditions in its global supply chain. Case study, Stanford Institute for the Study of Supply Chain Responsibility, Stanford, CA.
- Porteous, A.H., S.V. Rammohan, H.L. Lee. 2015. Carrots or sticks? Improving social and environmental compliance at suppliers through incentives and penalties. *Production and Operations Management* **24**(9) 1402–1413.
- Rangan, V.K., L. Chase, S. Karim. 2015. The truth about csr. *Harvard Business Review* **93**(1/2) 40–49.
- Roch, S., C. Samuelson. 1997. Effects of environmental uncertainty and social value orientation in resource dilemmas. *Organizational Behavior & Human Decision Processes* **70**(3) 221–235.
- Rode, J., R. Hogarth, M. Le Menestrel. 2008. Ethical differentiation and market behavior: An experimental approach. *Journal of Economic Behavior & Organization* **66**(2) 265–280.
- Rode, Julian. 2010. Truth and trust in communication: Experiments on the effect of a competitive context. *Games and Economic Behavior* **68**(1) 325–338.
- Roels, G., X. Su. 2013. Optimal design of social comparison effects: Setting reference groups and reference points. *Management Sci.* **60**(3) 606–627.
- Rooderkerk, R.P., H.J. Van Heerde, T.H.A. Bijmolt. 2011. Incorporating context effects into a choice model. *Journal of Marketing Research* **48**(4) 767–780.
- Rousseau, Denise M, Sim B Sitkin, Ronald S Burt, Colin Camerer. 1998. Not so different after all: A cross-discipline view of trust. *Academy of management review* **23**(3) 393–404.
- Rubin, J., R. Sheremeta. 2016. Principal–agent settings with random shocks. *Management Science* **62**(4) 985–999.
- Seinen, I., A. Schram. 2006. Social status and group norms: Indirect reciprocity in a repeated helping experiment. *Eur. Econom. Rev.* **50**(3) 581–602.
- Sen, S., C. Bhattacharya. 2001. Does doing good always lead to doing better? Consumer reactions to corporate social responsibility. *Journal of Marketing Research* **38**(2) 225–243.
- Smith, Tom W, Peter Marsden, Michael Hout, Jibum Kim. 2011. General social surveys, 1972–2010 [machine-readable data file]. *Sponsored by National Science Foundation. NORC, Chicago: National Opinion Research Center [producer]* .
- Stanca, L. 2009. Measuring indirect reciprocity: Whose back do we scratch? *J. Econom. Psych.* **30**(2) 190–202.
- Starbucks. 2014. Starbucks Global Responsibility Report 2014. Available at <http://globalassets.starbucks.com/assets/ea2441eb7cf647bb8ce8bb40f75e267e.pdf>.
- Starbucks. 2017. Ethical Sourcing: Coffee. Retrieved June 1, 2017, <https://www.starbucks.com/responsibility/sourcing/coffee>.
- Starbucks Corporation. 2015. Starbucks Global Responsibility Report 2015. Available at <https://www.starbucks.com/responsibility/global-report>.
- Talluri, S., R. Narasimhan, W. Chung. 2010. Manufacturer cooperation in supplier development

- under risk. *European Journal of Operational Research* **207**(1) 165–173.
- The Guardian. 2016. The expensive “Italian” shoes made for a pittance in east European sweatshops. Retrieved June 2, 2017, <https://www.theguardian.com/fashion/2016/aug/20/shoes-uk-high-street-made-for-a-pittance-eastern-europe-sweatshop>.
- The Sustainability Consortium. 2016. Greening Global Supply Chains - 2016 Impact Report. Available at <https://www.sustainabilityconsortium.org/wp-content/themes/enfold-child/assets/pdf/2016-impact-report.pdf>.
- Trudel, R., J. Cotte. 2009. Does it pay to be good? *MIT Sloan Management Review* **50**(2) 61–68.
- Unilever. 2014. Sustainable palm oil progress report. Available at https://www.unilever.com/Images/uslp-unilever-palm-oil-report-nov14_tcm244-424235_en.pdf.
- Vermeir, I., W. Verbeke. 2006. Sustainable food consumption: Exploring the consumer “attitude-behavioral intention” gap. *J. Agric. Environ. Ethics* **19**(2) 169–194.
- Wang, S., P. Sun, F. Véricourt. 2016. Inducing environmental disclosures: A dynamic mechanism design approach. *Operations Research* **64**(2) 371–389.
- Wang, Y., W. Gilland, B. Tomlin. 2010. Mitigating supply risk: Dual sourcing or process improvement? *Manufacturing & Service Operations Management* **12**(3) 489–510.
- Wells, W. 1975. Psychographics: A critical review. *Journal of Marketing Research* **12**(2) 196–213.
- Wooldridge, Jeffrey M. 2002. *Econometric analysis of cross section and panel data*. MIT press.
- Yamagishi, Toshio, Satoshi Akutsu, Kisuk Cho, Yumi Inoue, Yang Li, Yoshie Matsumoto. 2015. Two-component model of general trust: predicting behavioral trust from attitudinal trust. *Social Cognition* **33**(5) 436–458.
- Yamagishi, Toshio, Midori Yamagishi. 1994. Trust and commitment in the united states and japan. *Motivation and emotion* **18**(2) 129–166.
- Zhang, H., G. Aydin, H.S. Heese. 2017. Curbing the use of conflict minerals: A supply network perspective. Working paper, Indiana University, Bloomington, IN.
- Zhu, K., R.Q. Zhang, F. Tsung. 2007. Pushing quality improvement along supply chains. *Management Science* **53**(3) 421–436.