The Dissertation Committee for Jonathan Richard Lhost
certifies that this is the approved version of the following dissertation:

Essays on Competition, Cooperation, and Market Structures

Committee:

David S. Sibley, Co-Supervisor

Dale O. Stahl, Co-Supervisor

Laurent Mathevet

Eugenio J. Miravete

Abraham L. Wickelgren

Essays on Competition, Cooperation, and Market Structures

by

Jonathan Richard Lhost, B.A.; M.S.Econ.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2014

To the family and friends whose endless support made this possible.

Acknowledgments

I would like to thank my advisors, David Sibley and Dale Stahl, for all of their help and support throughout this process. I would also like to thank the other members of my committee, Laurent Mathevet, Abe Wickelgren, and Eugenio Miravete. I am also grateful for the help and support of Padmanabhan Srinagesh, Brijesh Pinto, Justin Leroux, Thomas Wiseman, Dwight Steward, Sandra Black, Jason Abrevaya, Gerald Oettinger, and Vivian Goldman-Leffler.

Essays on Competition, Cooperation, and Market Structures

Jonathan Richard Lhost, Ph.D.
The University of Texas at Austin, 2014

Supervisors: David S. Sibley
Dale O. Stahl

My dissertation examines competition, cooperation, and efficiency in three market settings in which a population of economic agents interact, either directly with each other in pairwise matches, directly with firms, or with firms via a platform.

In one chapter I consider a population of customers who have different valuations for a good sold by competing merchants, as well as varying preferences over the merchant from which to purchase the good and the payment form with which to make the purchase, and examine what the effects might be if a merchant placed an additional surcharge on transactions completed with a payment form that is more costly for the merchant. The cost for the merchant can vary dramatically depending on the payment form used. For example, a credit card transaction is generally more expensive for the merchant than a debit card transaction, even if the transaction is completed using the same technology and is processed over the same network (e.g., a MasterCard signature debit transaction and a MasterCard credit card transaction). Historically, with limited exceptions, merchants have been prohibited, both by law and by the contract permitting the acceptance of that network's cards, from charging customers different prices for transactions completed using different payment cards,

despite the different costs these transactions impose on them. Recent concessions made by several major payment networks in response to legal challenges raises the possibility that this paradigm might change in the future. This chapter examines what the effects might be if merchants were permitted to charge customers different prices based on the payment form and whether these effects depend on differences between the merchants, such as differences in the marginal cost of providing the good.

In another chapter, I consider a population of individuals made up of more-patient and less-patient types who interact directly with each other in a repeated prisoner's dilemma embedded in a search model. A player is matched anonymously with another player to play a prisoner's dilemma game repeatedly until the match is ended, either exogenously or endogenously by one of the players, at which point each player may receive another random match. I first determine when it is feasible to achieve the best outcome in which all players cooperate. When it is not possible to achieve full cooperation, I examine how welfare can be improved over the outcome in which no players cooperate. When conditions are such that less-patient players choose not to cooperate, I first examine how separation by action within a single market can increase welfare for all players over the uncooperative equilibrium, with more-patient players choosing to cooperate in hopes of forming a cooperative relationship, despite the risk of being matched with a less-patient player who chooses not to cooperate. I then examine how full separation of the more- and less-patient players, made possible by introduction of a second market, can increase the welfare of the more-patient players without harming the less-patient players.

In a third chapter, customers choose to purchase a good from one of several competing firms in a setting in which network congestion and firms' investment in capacity plays an important role in firm costs and product quality, e.g., the wireless industry. Wireless carriers (e.g., Verizon) compete not only on the price of their service but also on its quality. The quality of a carrier's service is determined in part by the quantity of customers it serves and by investment in capacity with which to serve them. While the primary effect of a carrier increasing its capacity is an increase in that carrier's service quality, there are also externality effects on other wireless carriers. For example, if carrier A increases its capacity, thereby increasing its service quality, and causes some customers to leave a competing carrier B, the service quality experienced by customers who remain with carrier B will increase as a result of the decreased congestion in carrier B's network. This chapter examines the interplay between these effects alongside traditional price competition in this oligopoly setting.

Table of Contents

Acknowledgments **v**

Abstract **vi**

Chapter 1. Credit or Debit? How Surcharging Affects Customers, Merchants, and the Platform **1**

1.1 Introduction 1

1.2 Structure of Payment Networks 10

1.3 Model 13

1.3.1 Utility Maximization 14

1.3.2 Merchant Profit Maximization 21

1.3.2.1 Price Setting 22

1.3.2.2 Credit Acceptance 24

1.3.3 Credit Platform Profit Maximization 25

1.3.4 Equilibrium 27

1.4 The Effects of Surcharging 29

1.4.1 Estimation Technique for Numerical Simulations 30

1.4.2 Merchant Credit Acceptance 33

1.4.2.1 Without Cash-Back 33

1.4.3 With Cash-Back 41

1.4.4 Quantity of Credit Transactions 42

1.4.5 Quantity-Weighted Average Price 45

1.4.6 Merchant Profits 46

1.4.6.1 Symmetric Merchants 46

1.4.6.2 Asymmetric merchants 47

1.4.6.3 Effects of Platform's Optimal Choices on Merchant Profits 48

1.5 Conclusion 49

Chapter 2. Worth the Wait? Cooperation in a Repeated Prisoner's Dilemma with Search **51**
2.1 Introduction 51
2.2 Cooperation in the Prisoner's Dilemma 55
2.3 Environment 58
2.3.1 Strategy of player i type t_i 61
2.3.2 Equilibrium 67
2.4 One Market 68
2.5 Two Markets 72
2.5.1 Two Market Setting 73
2.5.2 Separating Equilibrium with Two Markets 74
2.5.3 Welfare Gains from Second Market 80
2.5.4 One versus Two Market Separating Equilibria 82
2.6 Endogenous versus Exogenous Breakup 84
2.6.1 First-Best Equilibrium with Endogenous Breakup 86
2.6.2 One-Market Separating Equilibrium with Endogenous Breakup 92
2.6.2.1 Distribution of types and states 93
2.6.2.2 Lowest matching probability: both types choose to stay matched 94
2.6.2.3 Middle matching probability: high type choose to end uncooperative matches 96
2.6.2.4 Highest matching probability: high and low type choose to end uncooperative matches 97
2.6.3 Hybrid Model: Between Fully Exogenous and Fully Endogenous Breakup 98
2.7 Model with Endogenous Breakup 99
2.7.1 Model with Endogenous Breakup 101
2.7.2 Equilibrium with Endogenous Breakup 106
2.7.2.1 Endogenous Breakup Decision of the Currently Matched 106
2.7.2.2 Stage Game Strategy Decision of the Newly Matched . 107
2.7.2.3 Market Decision of the Unmatched 108
2.7.2.4 Equilibrium with Endogenous Breakup 108
2.7.3 First Best with Endogenous Breakup 109
2.7.3.1 Endogenous Breakup Decision of the Currently Matched 109

2.7.3.2 Stage Game Strategy Decision of the Newly Matched . 116
2.7.3.3 Completing the Specification of the First Best Equilibrium with Endogenous Breakup 120
2.7.4 Separating Equilibrium with Endogenous Breakup 121
2.7.4.1 Separating Equilibrium with One Market 122
2.7.4.2 Separating Equilibrium with Two Markets 123
2.7.5 Optimality of Equilibria with Endogenous Breakup 124
2.7.6 Hybrid Model: Between Fully Exogenous and Fully Endogenous Breakup 125
2.8 Conclusion 125

Chapter 3. Effects of Spectrum Holdings on Equilibrium in the Wireless Market 128
3.1 Introduction 128
3.2 Model 133
3.2.1 Optimal Capacities 136
3.2.2 Two Firm Example 138
3.3 Calibration 142
3.3.1 Data 142
3.3.2 Model Calibration 146
3.4 Effects of Spectrum Holdings 148
3.4.1 Additional Spectrum for One Firm 148
3.4.2 Equalizing Spectrum Holdings 150
3.4.3 Transfer of Spectrum Between Two Firms 153
3.5 Robustness 155
3.6 Conclusion 157

References 168

Chapter 1

Credit or Debit? How Surcharging Affects Customers, Merchants, and the Platform

1.1 Introduction

"Please press credit or debit," the cashier instructs you. You press the button corresponding to the card you just pulled from your wallet, giving little thought to the complex, multi-party transaction you just initiated or to how your specific choice might affect the merchant. Suppose your purchase totaled $40. If the card you just used was a Visa credit card, then, on average, the merchant just paid around $0.80 to the bank handling the merchant's Visa transactions. What if instead of a credit card you pulled a debit card out of your wallet? If you signed, then it might have cost the merchant $0.60, and if you entered your PIN, it might have only cost the merchant $0.20.[1]

While these costs may seem insignificant, the cumulative effect is large. A significant fraction of economic activity is completed using payment cards. In 2012, 87 billion transactions were completed using payment cards in the United States, with over $2.2 trillion worth of goods and services purchased using credit cards and over $1.8 trillion using debit cards. The costs for merchants associated with these transactions are significant, with merchants paying over $66 billion in fees in 2012.[2]

[1]Examples are based on the average merchant discount rates for transactions completed using a Visa credit card, Visa signature debit card, and Visa PIN debit card.

[2]Source: Nilson Report.

The economics literature generally finds that merchants would have higher profits if they were not prohibited from what is known as "surcharging," that is, charging a customer a higher or lower price depending upon the cost to the merchant of the customer's payment method. Merchants have raised legal challenges against the prohibition on surcharging, claiming it is anti-competitive, increases their costs, and reduces their profits. The economics literature generally agrees, finding that the ability to surcharge increases merchant profits. I also find this to be the case when I make the common assumption that merchants are identical. However, while this assumption greatly simplifies the analysis, it conceals an important alternative result: small merchants are likely to be hurt by surcharging. When I relax the assumption that merchants have identical marginal costs, the merchant with lower costs, typically a larger retailer, benefits from surcharging, whereas the merchant without an ability to reduce costs, typically a smaller retailer, does not. Trade groups representing small merchants, such as local convenience stores and independent bookshops, have been among the parties suing for, among other things, the right to surcharge. This result calls into question what benefit, if any, these merchants would receive from surcharging.

If payment cards imposed costs without conferring benefits, of course, no merchants would accept them. Compared to checks, cards save time and do not run the same risk of non-payment. And unlike cash, theft from employees and customers is not a concern. In addition to transactional benefits, cards may increase sales. If a merchant does not accept payment via cards, some customers who wish to pay by card may go to a competing merchant who does. And, especially for larger purchases, using credit cards allows some customers to finance purchases they may not otherwise be able to make.

Payment networks (e.g., Visa, MasterCard, American Express, Discover) require a merchant to sign a contract specifying the terms of accepting that platform's cards. These contracts have historically prohibited merchants from treating a customer paying with one of that platform's cards differently from a customer paying with any other card. Several of these anti-discrimination rules have been challenged in court.[3] In 1996, Wal-Mart and several other retailers sued Visa and MasterCard over those networks' "Honor All Cards" rule. The Honor All Cards rule required any merchant who wanted to accept any of a network's cards to accept all of that network's cards. That is, if a merchant wanted to accept any payment card branded with a MasterCard logo, the merchant was required to accept all cards branded with a MasterCard logo. In 2003, Visa and MasterCard agreed to limit the scope of the Honor All Cards rule so that a merchant accepting that network's debit cards was no longer required to accept that network's credit cards, and vice versa. However, the Honor All Cards rule still applies within the class of credit or debit cards, so any merchant who accepts a MasterCard credit (debit) card must accept all MasterCard credit (debit) cards, and similarly for Visa-branded credit and debit cards.

The other anti-discrimination rule that has received much legal attention is what is know as the "No Surcharge Rule." Under the No Surcharge Rule imposed by all major payment card networks, merchants are typically allowed to give a discount to customers paying with cash, but are not allowed to set different prices for customers paying with different payment cards. For example, a merchant cannot set different prices for customers paying with a Visa credit card and a Visa debit card, or with an American Express credit card and a Discover credit card.

The costs of cash transactions vary wildly for different merchants in different industries, depending on factors such as the overall quantity of cash, the likelihood

[3]For a background on the history of payment cards, see Evans and Schmalensee (2005). For the history of the antitrust litigation involving payment networks, see Wildfang and Marth (2006).

of mistakes, and the ease of theft,[4] and thus, so too do the potential benefits for merchants of providing cash discounts. Some merchants may have low cash-handling costs, and thus, might benefit from providing cash discounts, while other merchants might prefer to avoid cash altogether. Regardless, what the No Surcharge Rule prohibits is what many merchants argue is the form of differential pricing merchants would benefit from most. For example, suppose a customer has two MasterCard credit cards in his wallet, one a standard consumer card and one a rewards card. The physical details of the transaction are identical for both cards. The card is swiped through the same equipment and processed over the same network, and the customer provides his signature in the same way. However, the rewards card might be double the cost for the merchant compared to the standard card. Merchants, they argue, should be allowed to place a surcharge on the more costly transaction, either to recoup some of the additional cost and/or to steer the customer towards using the less-costly option.

In response to recent litigation brought by groups of merchants and the United States Department of Justice, Visa and MasterCard, the two largest payment card networks, agreed to a settlement that would allow merchants to start surcharging in January 2013, subject to certain restrictions. Merchants may surcharge based on the "brand level," for example, setting one price for all Visa credit cards and a different price for all MasterCard credit cards. Merchants may also surcharge based on the "product level," for example, setting different prices for different Visa credit card types (e.g., Traditional, Traditional Rewards, Signature, Signature Preferred). In each case, the surcharge may not exceed the merchant's cost of accepting that brand or type of card, or 4%, whichever is lower.[5] Merchants who choose to surcharge

[4]See, for example, Garcia-Swartz et al. (2006).

[5]Additional details are available at `www.visa.com/merchantsurcharging`.

are required to disclose this practice to customers, both through signage and on the receipt. Despite this agreement, two important restrictions may limit merchants' ability to surcharge in practice. First, surcharging is still prohibited by laws in some states.[6] Second, merchants who accept payment cards from platforms other than Visa or MasterCard are still bound by the contractual agreements they have with those platforms. For example, American Express, who was sued regarding its No Surcharge Rule at the same time as Visa and Mastercard were, has not agreed to a settlement and continues to defend its No Surcharge Rule. The outcome of this litigation will affect the future landscape of surcharging. However, with surcharging now permitted by the two largest payment networks, pressure on the other payment networks to allow surcharging will increase, making it increasingly like that U.S. consumers will start to encounter merchants who surcharge. Thus, it is important to understand the potential effects of merchant surcharging.

The economics literature on surcharging is part of the more general literature on two-sided markets. Payment systems are a two-sided market because the payment card networks serve as a platform to facilitate interactions between two groups, merchants and customers. In general, to succeed, a platform must attract enough users from both groups. For example, a newspaper must attract enough advertisers and readers to be profitable; it can only attract enough advertisers if it also attracts enough readers to view the advertisements, but if it includes too many advertisers, it may lose readers. The same is true for payment card platforms. A customer has little use for an American Express card if he is not able to use it to make purchases, and it is only beneficial for a merchant to accept American Express for payment if doing so attracts enough customers. In each case, the platform must balance the payments

[6] Currently states that prohbit surcharging inclue California, Colorado, Connecticut, Florida, Kansas, Maine, Massachusetts, New York, Oklahoma, and Texas.

received by or paid to each side of the market. Some newspapers are provided for free at local businesses and receive payment only from the advertising side of the market, while other newspapers require readers to pay for a subscription. Similarly, there is a wide range of prices paid by merchants and by customers for using various payment cards. One strain of the economics literature seeks to explain the optimal way to balance the prices paid by each side of the two-sided market.[7]

Much of the literature on two-sided markets that is specifically about payment card systems focuses on the role of the interchange fee.[8] Some payment card systems consist of two separate groups of banks, ones that interact with the customer side of the market and others that interact with the merchant side of the market, each setting the price for the side of the market with which it interacts. The interchange fee is a monetary transfer between these two groups and is the main component of the total price paid to the platform by merchants and customers. A complete explanation of the structure of payment networks is provided in Section 1.2. After that section, there will be no further mention of the interchange fee because the credit platform considered in the model will be a three party system in which the platform is a single entity and the concept of interchange does not apply.

However, a few aspects of the interchange fee warrant mentioning here. First, because the interchange fee is a monetary transfer between banks, it falls under the purview of the Federal Reserve and its regulatory capacity. Second, the interchange fee is the main component of the total fee paid by merchants for debit card transactions and certain credit card transactions. Thus, higher interchange fees usually translate into higher costs for merchants. This has been given much popular attention

[7]For examples, see Evans (2003), Rochet and Tirole (2003b), Rochet and Tirole (2006), Armstrong (2006), Kaiser and Wright (2006), Weyl (2010), and Rysman (2009).

[8]Examples include Gans and King (2003), Rochet and Tirole (2003a), Wright (2003), Wright (2004), and Guthrie and Wright (2007).

recently in the aftermath of the financial crisis of the late 2000s and the Dodd-Frank Wall Street Reform and Consumer Protection Act, specifically in what is known as the "Durbin Amendment." The Durbin Amendment instructed the Federal Reserve to place a cap on debit card interchange fees at levels that are "reasonable and proportional." The Federal Reserve originally proposed a cap of around \$0.12 per debit card transaction, but implemented a cap of what amounts to \$0.23 on the average debit card transaction of \$38.00.[9] However, various consumer groups, including the National Retail Federation and the National Restaurant Association sued the Fed, claiming that the Fed considered factors in setting the cap that were outside the scope of what it was allowed to consider according to the language of the law, and consequently, set the cap not at a "reasonable" level, but rather, at an "unreasonable" level. A judge in federal district court agreed, instructing the Fed to formulate a new rule.[10]. The Federal Reserve has stated its intent to appeal the decision, and what the cap on debit card interchange fees may be in the future is uncertain. However, it seems clear that a cap will be set at some level. In the model presented here, the cost for merchants of transactions completed by debit card will be held fixed.

There are two main strands in the economics literature regarding the No Surcharge Rule. One follows from Rochet and Tirole (2002) and the other from Schwartz and Vincent (2006). In each, as will be the case in the model presented here, customers pay by one of two payment forms, one that is exogenous ("cash") and one that is controlled by a profit-maximizing platform ("card"). In Rochet and Tirole (2002), there are two symmetric merchants who compete à la Hotelling to attract

[9]Full text of the Federal Reserve's ruling is available here: `www.gpo.gov/fdsys/pkg/FR-2011-07-20/pdf/2011-16861.pdf`.

[10]The full memorandum opinion issued by United States District Judge Richard Leon, an interesting example of law and economics, is available here: `https://ecf.dcd.uscourts.gov/cgi-bin/show_public_doc?2011cv2075-38`.

the business of a unit mass of customers. Each merchant decides whether to accept cash only, or both cash and the payment card. Each customer has perfectly inelastic demand, always purchasing one unit of the good for sale by the two merchants, choosing which merchant to purchase it from based on the price set by each merchant and his preferences for the two merchants, specified as a "transportation cost" as in a standard "linear city" model. Customers differ in the benefit they receive from paying with the payment card instead of cash, each choosing to use the payment form that maximizes his utility. In Schwartz and Vincent (2006) there is a monopoly merchant that sells a single good and decides whether to accept both cash and credit, or cash only. Customers have elastic demand, but are exogenously divided into those who always pay by cash and those who always pay by card. So while each customer is able to choose the quantity he would like to purchase, he is unable to decide how he would like to pay for it. Both Rochet and Tirole (2002) and Schwartz and Vincent (2006) find similar effects of lifting the No Surcharge Rule and allowing merchants to charge customers different prices based on their payment form. When allowed to surcharge, merchants pass through to customers their cost of card transactions. Customers paying by card pay a higher price, while cash customers pay a lower price. Merchants have higher profits.

The model presented in Section 1.3 will combine aspects of both Rochet and Tirole (2002) and Schwartz and Vincent (2006). There are two merchants who compete on prices and card acceptance for customers who have preferences over merchants and payment forms, as in Rochet and Tirole (2002). However, customers have elastic demand, as in Schwartz and Vincent (2006), and thus are heterogeneous along all three dimensions. Another key feature of the model presented in Section 1.3 is that the merchants are not constrained to be symmetric. One consequence of allowing merchants to differ is that they do not always make the same equilibrium decision regarding card acceptance, with both accepting or both not accepting.[11] Thus, customers, who have preferences over both merchants and payment forms, face a richer set of trade-offs when maximizing utility. This, combined with the fact that I allow for customers who have elastic demand, leads to interesting new results on the effects of the No Surcharge Rule.

In general, the credit card platform has two potential sources of profits, the merchant side of the market and the customer side of the market. Under the No Surcharge Rule, the credit card platform finds it optimal to earn profits from the merchant side of the market while subsidizing credit card use by providing rewards (e.g., "cash-back") to customers who use the credit card. When merchants are allowed to surcharge, however, the platform finds it optimal to shift to making more profits from the customer side of the market, charging customers a fee for using the credit card while simultaneously lowering the price merchants must pay the platform when customers pay using the credit card. Merchants now have lower costs and choose to lower the prices they charge customers. However, the effect on merchant profits depends on the similarity of the merchants.

When I assume merchants have identical marginal costs, I find the same result as in the previous literature, that merchants have higher profits when allowed to surcharge. However, relaxing the assumption that merchants are identical reveals an alternative possibility that surcharging might not always benefit merchants. Specifically, if one merchant has a marginal cost advantage over the other, as might be the case with a large national retailer competing with a small local merchant, then only the merchant with lower marginal costs benefits from surcharging. As small busi-

[11] In the model presented in Section 1.3, the two payment forms will be labeled "debit" and "credit" instead of "cash" and "card," but serve similar roles, with "card" acceptance in Rochet and Tirole (2002) equivalent to "credit" acceptance here. In Schwartz and Vincent (2006), there is only one merchant, so the concept of "both" merchants accepting does not apply.

nesses have been among the most vocal about their desire to be allowed to surcharge, this finding that being allowed to surcharge might reduce their profits is of particular interest.

The remainder of this paper proceeds as follows. Section 1.2 describes the structure of a typical payment card network, focusing on the "three-party" structure that will be used in the following section, but also describing a "four-party" system. Section 1.3 describes the formal model. Results are presented in Section 1.4 and Section 1.5 concludes.

1.2 Structure of Payment Networks

This section describes the general structure of a three-party payment network, including who the three parties are and the interactions between them.[12] Consider a three-party payment network as shown in Figure 1.1, consisting of customers, merchants, and a payment platform. The payment platform provides a method of payment that a customer can use at a merchant that is part of the network. Typical examples of the payment method include credit cards and debit cards.

Suppose a customer makes a purchase from a merchant for price P. If the customer pays with cash, the customer hands the cash to the merchant in exchange for the good, and the transaction is complete. Suppose instead that the customer consummates the transaction using a card provided by the payment platform. In this case, the customer receives the good without directly transferring money to the merchant. Instead, the merchant is paid by the platform. The merchant receives the purchase price P from the platform, minus a fee. This fee is typically called

[12] The other common type of structure, a four-party payment network, will be discussed briefly at the end of this section.

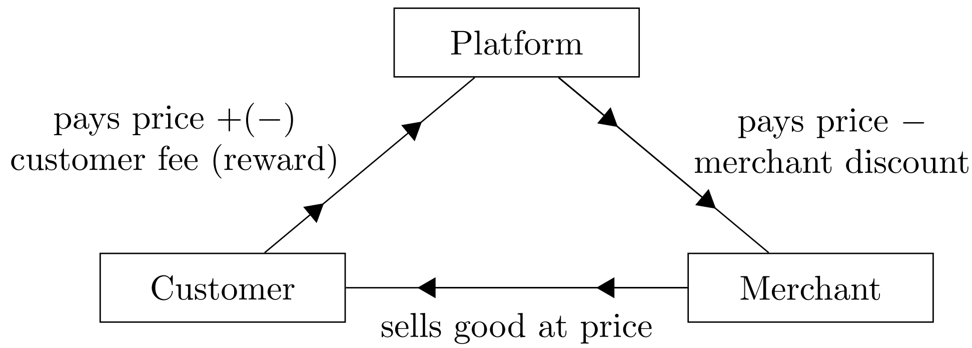

Figure 1.1: Three-party payment network

the "merchant discount" or "discount rate" because what the merchant receives in payment is less than the purchase price P.[13]

The customer pays the payment platform the purchase price P, plus any fees assessed and minus any rewards provided by the platform. The timing of the payment made by the customer to the platform depends on the type of card. For a credit card, the customer typically receives a monthly bill and must make a monthly payment. The option to pay only a portion of the total charges and borrow the remainder at a pre-agreed-upon interest rate schedule is an important aspect of credit cards, but not one we will focus on in this paper. For this paper, the cost of financing purchases will simply be included in the fee paid by the customer to the payment platform in addition to the purchase price P. For a debit card, payment from the customer to the payment platform is typically deducted automatically from the customer's account.[14]

[13]The merchant discount is also commonly called a "swipe fee" because it is a fee the merchant pays after swiping a customers card through a card reader.

[14]There are two main types of debit card transactions, PIN and signature. For PIN debit card transactions, the customer enters a Personal Identification Number into a keypad at the time of purchase, and the funds are removed from the customer's account immediately. Because of this direct connection to the customer's account, PIN debit card transactions are also called "online" debit card transactions. Signature debit card transactions (or "offline") are processed in the same manner as credit card transactions, with the customer signing a receipt authorizing the payment

Customers might also pay other fees to both credit and debit payment platforms, such as fees for late or incomplete payments, or an annual fee for use of that platform's payment card.

It is also common for customers to receive payments from the payment platform. Credit card platforms often give rewards to cardholders for using that platform's cards, including cash-back bonuses, airline miles, discounts on goods and services, etc. Debit card platforms sometimes provide similar rewards. In addition, the "reward" sometimes provided to users of a platform's debit card is free or discounted use of banking services, such as a "free checking" account.

In the model presented in Section 1.3, there will be a debit card and a credit card. The credit card platform will be an active player in the model, maximizing its profits by optimally setting the discount rates paid by merchants and by setting the reward paid to or fee received from customers paying with the credit card. Customers in the model will purchase at most a single good. Consequently, all fees and rewards will be combined into a single term. The debit card platform will be taken as exogenous, and thus, any fees or rewards associated with the debit card are subsumed into the customer's preferences over payment forms.

The model presented in Section 1.3 will be a three-party payment network as just discussed. The other common type of payment network is a four-party network. In this type of network, the payment platform consists of two separate groups, "acquires" that interact with the merchants and "issuers" that interact with cardholders. The issuers issue cards to cardholders, receive cardholder payments and fees, and pay cardhold rewards. The "acquires" acquire merchants to be part of the payment network and pay merchants for purchases, minus the merchant discount. In order to pay

---

and funds later removed from the customer's account.

the merchants, acquires receive the purchase price from issuers, minus what is called the "interchange fee." The interchange fee is effectively a portion of the purchase price that issuers retain to finance the issuing of cards and the payment of rewards to cardholders. The acquires, in turn, finance their payment of this interchange fee through the merchant discount. We are interested in examining the interaction between merchants, customers, and the payment platform, and not the inner workings of the financial institutions that make up payment network. Thus, for our purposes, it does not matter if we examine a three-party network in which one entity interacts with both merchants and customers or a four-party network with separate issuers and acquires, so we will focus on a three-party payment network.

In order to accept payments using a platform's cards, merchants typically must sign a contract dictating the terms of accepting that platform's cards. These contracts typically require merchants treat all cards equally. That is, the merchant cannot charge a customer one price if the customer uses card A and a different price if he uses card B. This requirement is often called the No Surcharge Rule because it prohibits merchants from placing a surcharge on purchases made by specific cards. In addition to the No Surcharge Rule imposed by payment platforms, legal restrictions in many locations also prohibit merchants from surcharging. Consequently, if a customer makes a purchase from a merchant, the customer will likely pay the same price regardless of the card they choose to use. The rest of this paper will examine the effects of removing this No Surcharge Rule.

1.3 Model

Two competing merchants sell an identical good to customers from a population of customers who have different valuations for the good and heterogeneous preferences over the two merchants and the payment method used, either credit or debit. Customers weight their utility of purchasing from each of the merchants and with the available payment forms and choose the option that maximizes their utility, or opt to make no purchase at all.

The two merchants compete both on price, as well as on whether they accept the credit card. Under the No Surcharge Rule, they must charge all customers the same price, regardless of how the customer chooses to pay. When surcharging is allowed, they can set a different price for credit and debit purchases.

The credit platform sets the price each merchant must pay when a customer makes a purchase using the credit card. The credit platform also interacts directly with customers by providing customers with a benefit for using the credit card, or by charging them a fee. The debit card platform is taken as exogenous.

The timing is as follows. First, the credit platform sets the price paid by merchants for each transaction completed by the credit card and the bonus paid to or fee paid by customers using the credit card. Second, the merchants decide whether to accept the credit card and then set the price they will charge customers for the good. Third, customers decide whether to buy the good, and if so, from which merchant and with which payment form. The players and these interactions are depicted in Figure 1.2.[15] The notation and details of each stage are presented in the following sections. We will find a subgame perfect Nash equilibrium solving by backwards induction, starting with the utility maximizing decisions of customers.

1.3.1 Utility Maximization

Let customer i denote a customer from a unit mass of potential customers. Customer i values the good at $v_i \in [\underline{v}, \overline{v}]$, distributed according to CDF V. Customer

[15]See also the game tree shown in Figure 1.4.

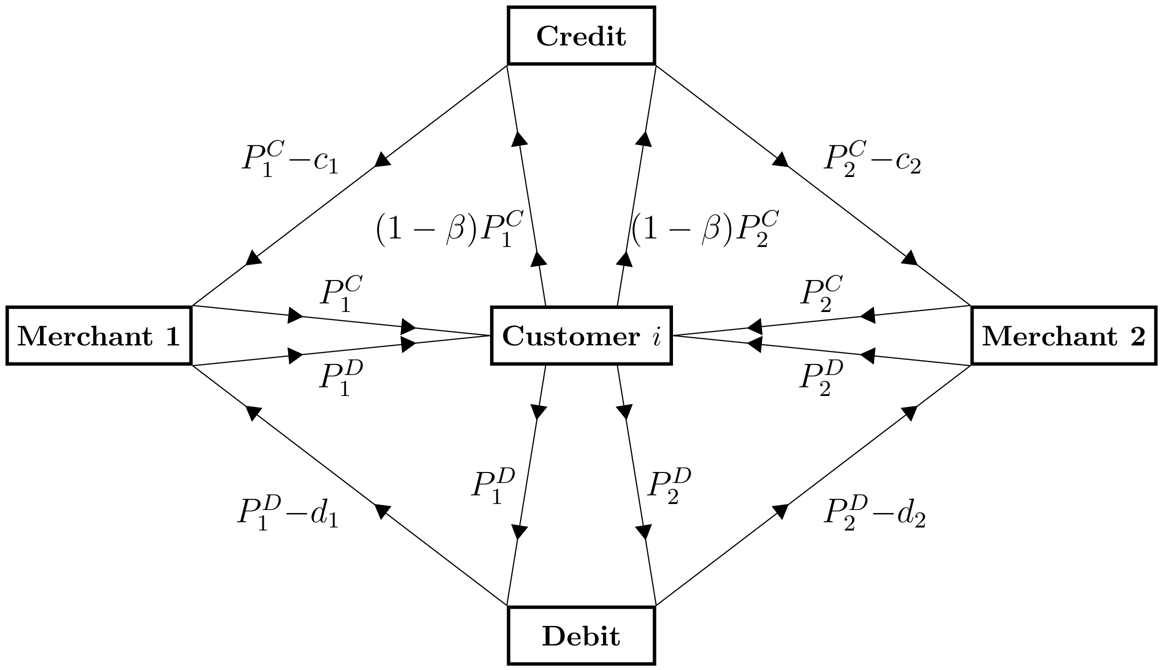

Figure 1.2: Possible choices for customer i and the corresponding interactions between customer, merchant, and platform

i has preferences over the two merchants, specified as disutilities similar to a "travel cost" often seen in Hotelling models, with customer i's disutility of purchasing from merchant $j \in \{1, 2\}$ denoted $t_i^j \in [\underline{t}^j, \bar{t}^j]$. Similarly, each customer i's preferences for paying with a credit card ($k = C$) or a debit card ($k = D$) are also specified as disutilities, with $t_i^k \in [\underline{t}^k, \bar{t}^k]$ denoting the disutility customer i receives from paying with payment form $k \in \{C, D\}$. Preferences over merchants are independently and identically distributed according to CDF T^M and preferences over payment forms, or "wallet preferences," are independently and identically distributed according to CDF T^W. If customer i purchases the good from merchant j using payment form k at price

P_j^k, then customer i's utility is

$$U_i^{kj} = v_i - P_j^k - t_i^j - t_i^k \tag{1.1}$$

Differences in preferences for the two merchants and the two payment forms are one reason why a customer might choose one merchant, payment-form combination over the others. The other reason is differences in the price paid by the customer. If the only payment form available is the debit card, then there are two possible prices, one for each merchant. However, with the credit card, there are two other potential sources of differences in price. Under the No Surcharge Rule, the price charged by each merchant must be the same whether a customer pays by credit or by debit: $P_j^C = P_j^D$. However, when surcharging is allowed, each merchant can set P_j^C different from P_j^D, bring the number of potentially different prices up to four.

Even under the No Surcharge Rule, the effective price paid by customers using credit might be different from those paying by debit because of rewards given or fees charged by card platforms. In this model, the debit card platform is taken as exogenous and any such fees or rewards for use of the debit card are normalized to zero. The credit platform, however, as part of its profit-maximizing behavior examined in Section 1.3.3, either pays a bonus to customers who use the credit card or charges them a fee. Specifically, if a customer makes a purchase for price P_j^C using the credit card, the net price paid by the customer is $(1-\beta)P_j^C$. For simplicity, we will refer to β as a "cash-back bonus" as the customer receives βP_j^C from the credit platform if $\beta > 0$. However, this "bonus" is actually a fee if β is negative, with the customer paying βP_j^C to the credit platform. Whether β is positive or negative, as well as its magnitude, is determined in equilibrium as part of the profit-maximizing decisions of the credit card platform. Note that the presence of β means that even under the

No Surcharge Rule, customers paying with credit may effectively pay a different price than customers paying with debit.[16]

If both merchants accept credit cards, there are four possible merchant, payment-form combinations. Customer i, given his valuation for the good (v_i), his preferences over merchants (t_i^j) and payment forms (t_i^k), the price at each merchant (P_j^k), and the cash-back bonus (β), chooses the merchant, payment-form combination from the available choices in order to maximize his utility. To make this choice, each customer i makes pairwise comparisons between the available choices. For example, if customer i were to go to merchant 2, he would choose to pay with the credit card if the utility from doing so is higher than that from paying with the debit card instead.

Specifically, customer i prefers credit at merchant 2 over debit at merchant 2 if

$$U_i^{C2} > U_i^{D2}$$
$$v_i - (1-\beta)P_2^C - t_i^2 - t_i^C > v_i - P_2^D - t_i^2 - t_i^D$$
$$w_i \equiv t_i^D - t_i^C > (1-\beta)P_2^C - P_2^D$$

where w_i is customer i's relative preference for credit over debit.[17] The other comparisons are made in an analogous fashion. Customer i prefers debit at merchant 2 over debit at merchant 1 if

$$U_i^{D2} > U_i^{D1}$$
$$v_i - P_2^D - t_i^2 - t_i^D > v_i - P_1^D - t_i^1 - t_i^D$$
$$P_1^D - P_2^D > t_i^2 - t_i^1 \equiv m_i$$

where m_i is customer i's relative preference for merchant 1 over merchant 2.

[16]As customers are making at most one purchase, we assume that their utility is affected by the effective purchase price and ignore any time delay between the time of purchase and the time the cash-back bonus is received (or paid, in the case of a fee).

Customer i prefers credit at merchant 2 over debit at merchant 1 if

$$
\begin{aligned}
U_i^{C2} &> U_i^{D1} \\
v_i - (1-\beta)P_2^C - t_i^2 - t_i^C &> v_i - P_1^D - t_i^1 - t_i^D \\
t_i^D - t_i^C &> (1-\beta)P_2^C - P_1^D + (t_i^2 - t_i^1) \\
w_i &> (1-\beta)P_2^C - P_1^D + m_i
\end{aligned}
$$

Customer i prefers credit at merchant 2 over credit at merchant 1 if

$$
\begin{aligned}
U_i^{C2} &> U_i^{C1} \\
v_i - (1-\beta)P_2^C - t_i^2 - t_i^C &> v_i - (1-\beta)P_1^C - t_i^1 - t_i^C \\
(1-\beta)(P_1^D - P_2^D) &> t_i^2 - t_i^1 \equiv m_i
\end{aligned}
$$

Customer i prefers credit at merchant 1 over debit at merchant 1 if

$$
\begin{aligned}
U_i^{C1} &> U_i^{D1} \\
v_i - (1-\beta)P_1^C - t_i^1 - t_i^C &> v_i - P_1^D - t_i^1 - t_i^D \\
w_i \equiv t_i^D - t_i^C &> (1-\beta)P_1^C - P_1^D
\end{aligned}
$$

Customer i prefers credit at merchant 1 over debit at merchant 2 if

$$
\begin{aligned}
U_i^{C1} &> U_i^{D2} \\
v_i - (1-\beta)P_1^C - t_i^1 - t_i^C &> v_i - P_2^D - t_i^2 - t_i^D \\
t_i^D - t_i^C &> (1-\beta)P_1^C - P_2^D - (t_i^2 - t_i^1) \\
w_i &> (1-\beta)P_1^C - P_2^D - m_i
\end{aligned}
$$

By examining the pairwise comparisons made by each customer i, we obtain self-selection constraints that describe how the unit mass of customers segments itself into regions corresponding with each of the merchant, payment-form combinations.

[17] w_i, as well as the m_i term found in the next equation, will be discussed in more detail shortly. This representation of customer preferences is an extension of Lhost, Srinagesh, and Sibley (2012).

We can write these self-selection constraints in terms of the parameters characterizing customer valuations and preferences, the prices charged by each merchant, and the bonus (or fee) set by the credit platform.

Each of the self-selection constraints that involves a comparison between the two merchants involves a term that is the difference in disutility obtained by going to each merchant, $t_i^2 - t_i^1$. It is convenient to replace this difference in disutilities with a third parameter, $m_i \equiv t_i^2 - t_i^1$, which represents customer i's preference for merchant 1 relative to merchant 2. Customer i's "merchant preference," m_i, is distributed according to CDF M between lower bound $\underline{m} \equiv \underline{t}_i^2 - \overline{t}_i^1$ and upper bound $\overline{m} \equiv \overline{t}_i^2 - \underline{t}_i^1$, where CDF M is derived directly from the underlying distribution of $t_i^j \sim T^M$.[18]

Similarly to how we just defined customer i's merchant preference, it is also convenient to define customer i's "wallet preferences" as customer i's relative preference for paying by credit relative to debit. In each self-selection constraint that involves a comparison between the two payment forms, it is possible to replace the term that is customer i's difference in disutility of paying by the two payment forms with customer i's wallet preferences, $w_i \equiv t_i^D - t_i^C$, where w_i is distributed according to CDF W between lower bound $\underline{w} \equiv \underline{t}_i^D - \overline{t}_i^C$ and upper bound $\overline{w} \equiv \overline{t}_i^D - \underline{t}_i^C$, with CDF W derived directly from the underlying distribution of $t_i^k \sim T^W$.

By replacing the four disutility parameters, t_i^j for $j \in \{1, 2\}$ and t_i^k for $k \in \{C, D\}$, with two parameters representing customer i's merchant and wallet preferences, m_i and w_i, we are able to depict the self-selection constraints, and thus the segmentation of customers into the four merchant, payment-form combinations, as show in Figure 1.3. The specific way in which customers are segmented shown in Figure 1.3 is only one possibility. For example, it is drawn in the case where $P_1^D > P_2^D$,

[18] For example, if T^M is the standard uniform distribution, then $m_i \equiv t_i^2 - t_i^1$ has the triangular distribution between -1 and 1.

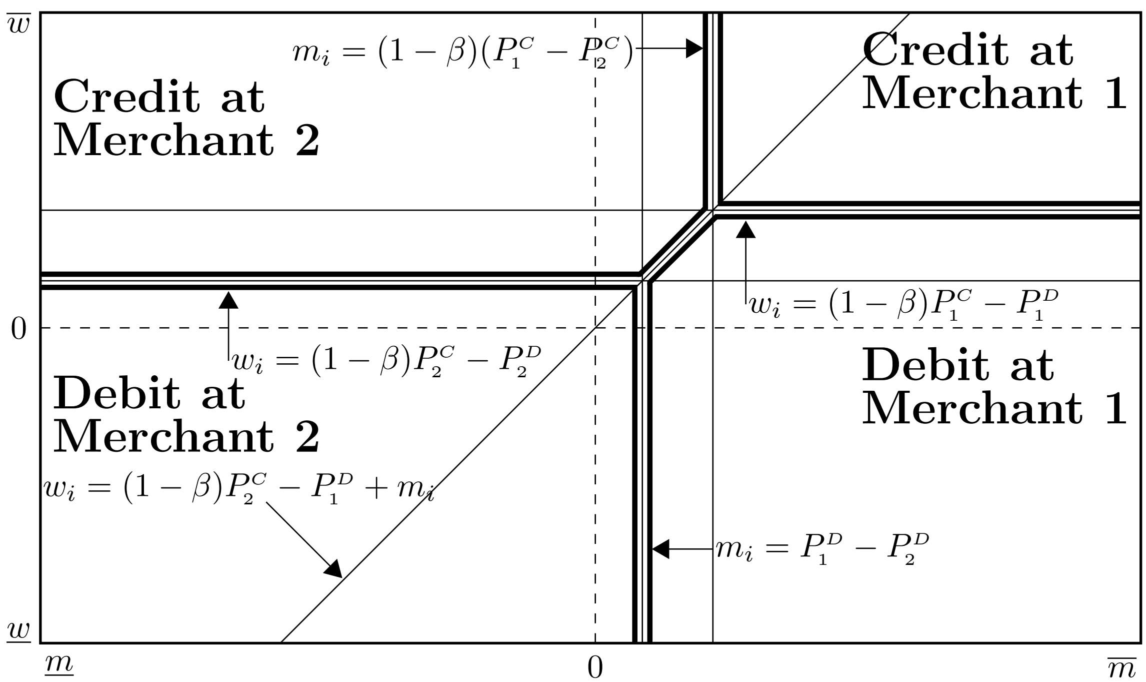

Figure 1.3: One case for segmentation of customers by self-selection constraints in terms of merchant preferences, $m_i \equiv t_i^2 - t_i^1$, and wallet preferences, $w_i \equiv t_i^D - t_i^C$

and thus the vertical line $m_i = P_1^D - P_2^D$ is to the right of $m_i = 0$. Intuitively, a customer paying by debit card who is exactly indifferent between the two merchants (a customer with $m_i = 0$) will choose merchant 2 if the price is lower at merchant 2 ($P_1^D > P_2^D$). If instead it was the case that $P_2^D > P_1^D$, then the line $m_i = P_1^D - P_2^D$ would be to the left of $m_i = 0$. Additional examples are shown in the online appendix, which includes an interactive version of Figure 1.3.

We have just seen how each customer i chooses which of the available merchant, payment-form combinations he would choose if he were to purchase the good. By integrating over the fraction of customers with preferences in each region we can obtain the fraction of customers who would choose each of the available merchant, payment-form combinations. If we assumed all customers purchase the good no matter the price, as is commonly assumed in models similar to the model presented here, the discussion of how customers maximize utility would be complete. However, customers in this model are assumed to have heterogeneous valuations for the good, and only customers with sufficiently-high valuations choose to purchase. To obtain the quantity of customers who actually purchase the good we must also integrate over the distribution of customer valuations, with only those for whom the utility of purchasing is positive choosing to purchase the good.

For example, in the case depicted in Figure 1.3, the customers who choose to make a purchase with a debit card from merchant 2 are those customers with preferences for merchants between the lower bound on the left side of Figure 1.3 ($\underline{m} \equiv \underline{t}_i^{\,2} - \bar{t}_i^{\,1}$) and $m_i \equiv t_i^2 - t_i^1 = P_1^D - P_2^D$, with preferences for payment forms between the lower bound on the bottom of Figure 1.3 ($\underline{w} \equiv \underline{t}_i^{\,D} - \bar{t}_i^{\,C}$) and $w_i \equiv t_i^D - t_i^C = (1 - \beta)P_2^C - P_2^D$, and with valuations such that $v_i > P_2^D + t_i^2 + t_i^D$, or equivalently, such that $U_i^{D2} > 0$. By integrating over customers with these preferences and valuations, an expression for the quantity of purchases made with a debit card at merchant 2, $\boldsymbol{Q}_2^D$, is obtained. The remaining quantity functions are formulated in the same manner.

The expressions $\boldsymbol{Q}_j^k$ are functions of the parameters describing customer preferences and valuations, the prices set by the merchants, and the cash-back bonus set by the credit platform. Derivatives of these functions will appear in the first order conditions of the merchants and the credit platform found in the following sections. A bold font will be used to distinguish expressions that refer to functions. However, to save on notation, they will be written without listing the arguments.

1.3.2 Merchant Profit Maximization

The two merchants compete both on price, as well as on whether they accept the credit card. Merchant profit maximization occurs in two stages, each deciding simultaneously whether to accept the credit card in the first stage and then simultaneously setting prices given these publicly observable credit acceptance choices. We will start by examining the second, price setting stage, determining each merchant's best response price function for given first stage credit acceptance decisions. We will then examine the optimal credit acceptance decisions of each merchant given optimal price setting behavior.

1.3.2.1 Price Setting

Merchant j has marginal cost μ_j of providing the good. In addition, each merchant also has a cost that is specific to the payment form used. This cost is the "merchant discount" discussed in Section 1.2. The merchant discount for merchant j is d_j if a customer uses the debit card and c_j if the customer uses the credit card.[19]

It is generally the case that if a purchase is made using a credit card it is more costly for the merchant than if the same purchase were made with a debit card. For ease of exposition, the two payment forms are labeled credit and debit, but this applies more generally to a lower cost and higher cost payment form. In this model, the credit card is meant to represent the payment form that is more costly for the merchant (i.e., in equilibrium it will generally be true that $c_j > d_j$). Consequently, merchants only find it optimal to accept the credit card if the benefit of attracting more customers outweighs the additional costs. Let $A_j = 1$ if merchant j accepts the credit card and $A_j = 0$ if he only accepts the debit card.

[19]The cost of transactions completed by the credit card, c_j, is set by the credit platform as part of its profit maximization, as described in Section 1.3.3.

If merchant j's credit acceptance decision is A_j and that of the competing merchant is A_{-j}, the profit of merchant j is

$$\boldsymbol{\pi}_j(A_j, A_{-j}) = \left(P_j^D - \mu_j - d_j\right) \boldsymbol{Q}_j^D + A_j \left(P_j^C - \mu_j - c_j\right) \boldsymbol{Q}_j^C \tag{1.2}$$

where under the No Surcharge Rule, the restriction is imposed that $P_j^C = P_j^D$. When merchant j accepts the credit card ($A_j = 1$) and the No Surcharge Rule is not in effect so that P_j^C can be different from P_j^D, merchant j has two first order conditions:

$$\frac{\partial \boldsymbol{\pi}_j(1, A_{-j})}{\partial P_j^D} = \boldsymbol{Q}_j^D + \left(P_j^D - \mu_j - d_j\right) \frac{\partial \boldsymbol{Q}_j^D}{\partial P_j^D} + \left(P_j^C - \mu_j - c_j\right) \frac{\partial \boldsymbol{Q}_j^C}{\partial P_j^D} = 0 \tag{1.3}$$

$$\frac{\partial \boldsymbol{\pi}_j(1, A_{-j})}{\partial P_j^C} = \boldsymbol{Q}_j^C + \left(P_j^D - \mu_j - d_j\right) \frac{\partial \boldsymbol{Q}_j^D}{\partial P_j^C} + \left(P_j^C - \mu_j - c_j\right) \frac{\partial \boldsymbol{Q}_j^C}{\partial P_j^C} = 0 \tag{1.4}$$

If merchant j does not accept the credit card, then the first order condition with respect to P_j^C, equation (1.4), does not apply. Thus, in the case of $A_j = 0$, merchant j's first order condition is

$$\frac{\partial \boldsymbol{\pi}_j(0, A_{-j})}{\partial P_j^D} = \boldsymbol{Q}_j^D + \left(P_j^D - \mu_j - d_j\right) \frac{\partial \boldsymbol{Q}_j^D}{\partial P_j^D} = 0 \tag{1.5}$$

When surcharging is allowed, each merchant has either one or two first order conditions. Thus, with surcharging, the two, three, or four best response price functions are determined by the solution of the system of two, three, or four first order conditions, respectively, two if neither merchant accepts credit, three if only one merchant accepts credit, and four if both merchants accept credit.

Under the No Surcharge Rule, each merchant is only allowed to set one price that applies to both debit customers and, if the merchant accepts credit and there are credit customers, to credit customers. Thus, under the No Surcharge Rule, the two best response price functions are determined by the solution of the system of two first order conditions, one for each merchant:

$$\frac{\partial \boldsymbol{\pi}_j(A_j, A_{-j})}{\partial P_j} = \boldsymbol{Q}_j^D + (P_j - \mu_j - d_j)\frac{\partial \boldsymbol{Q}_j^D}{\partial P_j} + A_j\left[\boldsymbol{Q}_j^C + (P_j - \mu_j - c_j)\frac{\partial \boldsymbol{Q}_j^C}{\partial P_j}\right] = 0 \quad (1.6)$$

where $P_j = P_j^D = P_j^C$.

For each setting, under the No Surcharge Rule and when surcharging is allowed, there are four possible systems of first order conditions corresponding with the four possible combinations of credit acceptance decisions, $(A_1, A_2) \in \{0,1\} \times \{0,1\}$. Consider one of these systems, say, when surcharging is allowed and both merchants accept credit. In this case there are four first order conditions, given by equations (1.3) and (1.4) for $j \in \{0,1\}$ (and with $A_{-j} = 1$). This system of four first order conditions can be solved for the optimal prices of each merchant. These prices are functions of the parameters describing customer preferences and valuations, the merchant discounts, and the cash-back bonus. We will denote these best response price functions as $\widehat{\boldsymbol{P}}_j^k(A_j, A_{-j})$. We will denote a quantity function evaluated at optimal prices determined by these best response price functions as $\widehat{\boldsymbol{Q}}_j^k(A_j, A_{-j})$, and (1.2) evaluated with $\widehat{\boldsymbol{P}}_j^k(A_j, A_{-j})$ and $\widehat{\boldsymbol{Q}}_j^k(A_j, A_{-j})$ as $\widehat{\boldsymbol{\pi}}_j(A_j, A_{-j})$.

1.3.2.2 Credit Acceptance

We have just seen how merchants determine how to set prices optimally for a given combination of credit acceptance decisions, (A_1, A_2). We now turn our attention to the determination of whether to accept credit.

By determining how each merchant will optimally set prices in each case, as well as what the resulting utility-maximizing behavior of customers will be and the quantities that will result, each merchant can determine the profit each will receive in each case, $\widehat{\boldsymbol{\pi}}_j(A_j, A_{-j})$ for $j \in \{0,1\}$ and $(A_j, A_{-j}) \in \{0,1\} \times \{0,1\}$. The result is the 2×2 game shown in Table 1.1. A Nash equilibrium of this game, denoted

	$A_2 = 1$	$A_2 = 0$
$A_1 = 1$	$\widehat{\boldsymbol{\pi}}_1(1,1), \widehat{\boldsymbol{\pi}}_2(1,1)$	$\widehat{\boldsymbol{\pi}}_1(1,0), \widehat{\boldsymbol{\pi}}_2(1,0)$
$A_1 = 0$	$\widehat{\boldsymbol{\pi}}_1(0,1), \widehat{\boldsymbol{\pi}}_2(0,1)$	$\widehat{\boldsymbol{\pi}}_1(0,0), \widehat{\boldsymbol{\pi}}_2(0,0)$

Table 1.1: Merchant credit acceptance subgame

$(\widehat{A}_1, \widehat{A}_2)$, yields the equilibrium credit acceptance decisions of the two merchants, and equilibrium prices are determined by the best response price functions corresponding with that combination of credit acceptance decisions (denoted $\widehat{\boldsymbol{P}}_j^k(\widehat{A}_j, \widehat{A}_{-j})$).

The 2×2 credit acceptance game could have multiple Nash equilibria. However, how to select which equilibrium will occur is not an issue in this framework because there is another player, the credit platform, who moves first and can thereby select the equilibrium which is best for it by its choice of merchant discounts and the cash back bonus. We will now turn our attention to how the credit platform maximizes profits.

1.3.3 Credit Platform Profit Maximization

When examining the profit-maximization problem of the merchants, we discussed the payment-form specific costs faced by merchant j, d_j when a customer pays with the debit card and c_j when a customer pays with the credit card. While the merchant discount c_j is a cost for merchant j, it is revenue for the credit platform. The credit platform can also receive revenue from customers who purchase using the credit card by setting $\beta < 0$. Alternatively, the platform can choose to set $\beta > 0$ and pay customers who purchase using the credit card, thereby providing customers with extra incentive to use the card but making β a cost for the platform. The platform has a marginal cost for transactions completed at merchant j, denoted μ_c^j. The profit of the credit platform is given by

$$\boldsymbol{\pi}_C = \left(c_1 - \mu_c^1 - \beta \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)\right) \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2) + \left(c_2 - \mu_c^2 - \beta \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)\right) \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1) \tag{1.7}$$

where the "hats" over the quantity functions, prices, and credit acceptance decisions indicate that both merchants are maximizing profits and customers are maximizing utility. Note that if merchant j's optimal decision is to not accept credit, then $\widehat{\boldsymbol{Q}}_j^C(0, \widehat{A}_{-j}) = 0$. Thus, it is unnecessary to multiply the first term of (1.7) by $\widehat{A}_1$ or the second by $\widehat{A}_2$.

Both $\widehat{\boldsymbol{P}}_1^C$ and $\widehat{\boldsymbol{P}}_2^C$ are affected by both c_1 and c_2. c_1 is a cost for merchant 1, and thus directly affects its optimal choice of $\widehat{\boldsymbol{P}}_1^C$. The same is true of c_2's effect on $\widehat{\boldsymbol{P}}_2^C$, and since merchant 1's choice of $\widehat{\boldsymbol{P}}_1^C$ is affected by $\widehat{\boldsymbol{P}}_2^C$, it is also affected by c_2. In addition, both $\widehat{\boldsymbol{P}}_1^C$ and $\widehat{\boldsymbol{P}}_2^C$ affect the optimal choices of customers, and thus the quantity expressions obtained by integrating over the optimal choices of customers are also both affected by both c_1 and c_2. While β does not directly affect the merchants, it does affect the decisions of customers, which affect the quantities and in turn the merchants' best response price functions. Consequently, all the expressions in (1.7) with a "hat" must be differentiated in each of the platform's first order conditions. The credit platform's optimal choice of c_1, c_2, and β is characterized by the solution to the following three first order conditions:

$$
\begin{aligned}
\frac{\partial \boldsymbol{\pi}_C}{\partial c_1} = 0 = &\left(1 - \beta \frac{\partial \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial c_1}\right) \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2) \\
&+ \left(c_1 - \mu_c^1 - \beta \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)\right) \frac{\partial \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial c_1} \\
&+ \left(1 - \beta \frac{\partial \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial c_1}\right) \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1) \\
&+ \left(c_2 - \mu_c^2 - \beta \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)\right) \frac{\partial \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial c_1}
\end{aligned} \tag{1.8}
$$

$$
\begin{aligned}
\frac{\partial \boldsymbol{\pi}_C}{\partial c_2} = 0 = &\left(1 - \beta \frac{\partial \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial c_2}\right) \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2) \\
&+ \left(c_1 - \mu_c^1 - \beta \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)\right) \frac{\partial \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial c_2} \\
&+ \left(1 - \beta \frac{\partial \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial c_2}\right) \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1) \\
&+ \left(c_2 - \mu_c^2 - \beta \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)\right) \frac{\partial \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial c_2}
\end{aligned} \tag{1.9}
$$

$$
\begin{aligned}
\frac{\partial \boldsymbol{\pi}_C}{\partial \beta} = 0 = &-\left(\widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2) + \beta \frac{\partial \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial \beta}\right) \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2) \\
&+ \left(c_1 - \mu_c^1 - \beta \widehat{\boldsymbol{P}}_1^C(\widehat{A}_1, \widehat{A}_2)\right) \frac{\partial \widehat{\boldsymbol{Q}}_1^C(\widehat{A}_1, \widehat{A}_2)}{\partial \beta} \\
&- \left(\widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1) + \beta \frac{\partial \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial \beta}\right) \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1) \\
&+ \left(c_2 - \mu_c^2 - \beta \widehat{\boldsymbol{P}}_2^C(\widehat{A}_2, \widehat{A}_1)\right) \frac{\partial \widehat{\boldsymbol{Q}}_2^C(\widehat{A}_2, \widehat{A}_1)}{\partial \beta}
\end{aligned} \tag{1.10}
$$

Recall that while we are examining the profit maximization problem of the credit platform last, that is because the credit platform selects the merchant discounts and the cash-back bonus first and we are solving by backwards induction. Thus, the credit platform's solution to the three first order conditions above will dictate the equilibrium outcome. As such, we will denote the solution to the platform's system of three first order conditions as c_1^*, c_2^*, β^*. The credit platform's equilibrium profits are analogously denoted $\boldsymbol{\pi}_C^*$, while that of the merchants is denoted $\widehat{\boldsymbol{\pi}}_j^*$.

1.3.4 Equilibrium

We have examined how customers maximize utility (Section 1.3.1), how merchants maximize profits by optimally setting prices (Section 1.3.2.1) and deciding whether to accept the credit card (Section 1.3.2.2), and how the credit platform optimally sets the merchant discounts and the cash-back bonus (Section 1.3.3). To find a subgame perfect Nash equilibrium, we need only to follow the steps discussed in each of these sections in the order presented.

First, we determine the optimal choice of all customers for all possible prices and combinations of credit acceptance decisions made by the merchants and all possible choices of the cash-back bonus set by the credit platform. Second, we determine the best response price functions for each merchant for all possible combinations of credit acceptance decisions and for all possible choices of merchant discounts and the cash-back bonus set by the credit platform, given the optimal choices made by all customers in each case. Third, we determine the optimal credit acceptance decisions of the merchants, given their optimal price-setting behavior, for all possible choices of merchant discounts and the cash-back bonus set by the credit platform. Fourth, we determine the optimal merchant discounts and cash-back bonus for the credit platform, given the profit-maximizing choices of the merchants and the utility-maximizing choices of the customers. This process is followed twice, once under the No Surcharge Rule and once when surcharging is allowed, to find the equilibrium in each setting.

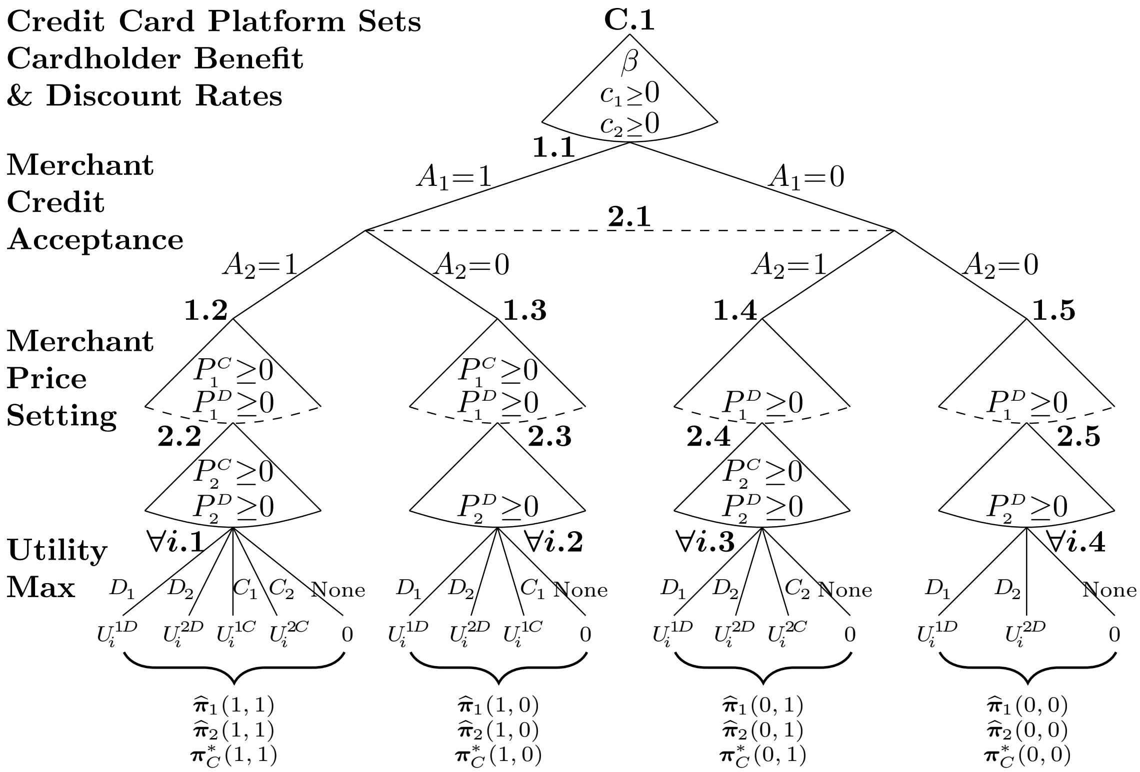

Figure 1.4: Note that the decision nodes found in the final level, those labeled $\forall i.1$ through $\forall i.4$, represent the choices available to each customer i from the unit mass of customers. The payoffs shown at the bottom of the tree include the utility received by customer i below the corresponding choices at these decision nodes, either U_i^{kj} if choosing one of the available merchant, payment-form combinations, or 0 if choosing not to purchase. The payoffs for the two merchants and the credit platform are centered below the set of choices for customer i because these profits involve the quantities obtained by integrating over the entire unit mass of customers.

A game tree corresponding with this process is shown in Figure 1.4.

1.4 The Effects of Surcharging

We are interested in examining the effects of lifting the No Surcharge Rule and allowing merchants to surcharge. To do this, we must determine equilibrium outcomes in both settings and compare them. For example, suppose we are interested in determining how surcharging affects the profits of merchant j. One approach would be to solve for a closed-form expression for $\widehat{\boldsymbol{\pi}}_j^*$ under the No Surcharge Rule and again with surcharging, and then examine properties of the function resulting from taking the difference between the two, such as determining when it is positive and performing comparative static analysis. Instead, we will use numerical simulations. It is true that numerical simulations require us to make assumptions about the distribution of customer preferences and valuations, and about values of certain parameters. However, finding closed-form solutions for equilibrium outcomes of interest also requires distributional assumptions. And when it is even possible to find closed-form solutions, the resulting equations of interest are sufficiently long to be of little use in achieving our goal. The resulting equations tend to be highly nonlinear because each comes from solutions to multiple systems of nonlinear equations. Thus determining whether an expression of interest is positive or negative requires choosing specific parameter values. It is possible to explore a much wider range of possibilities by using numerical simulations, and thus, this is the approach we will pursue.

1.4.1 Estimation Technique for Numerical Simulations

To examine the effects of surcharging, an equilibrium was found in both settings, under the No Surcharge Rule and with surcharging, and the equilibrium outcomes compared. The steps taken to find an equilibrium were as follows:

1. Choose the equilibrium to be found, including the setting, either under the No Surcharge Rule or with surcharging, the distributions for customer preferences

and valuations, values for the debit discount rates, d_1 and d_2, values of the marginal costs of the merchants, μ_1 and μ_2, and values of the marginal costs of the platform, μ_c^1 and μ_c^2.

2. Solve the customer utility maximization problem to obtain closed-form solutions for the quantity expressions, $\boldsymbol{Q}_j^k$, in all possible cases.

3. Using the quantity expressions found in step 2, formulate closed-form expressions for merchant profits, $\boldsymbol{\pi}_j(A_j, A_{-j})$, in all possible cases.

4. Using the profit equations found in step 3, find the first order conditions with respect to the prices, P_j^k, and formulate the system of first order conditions that must be solved for each possible case.

5. For each possible system of merchant first order conditions found in step 4, formulate the corresponding second order conditions that must be satisfied for a solution to the system to be a solution to the merchants' profit maximization problem.

6. For a given numerical choice of merchant discounts and the cash-back bonus set by the platform, (c_1, c_2, β), solve the merchants' profit maximization problem.

   6.1 For all four possible combinations of merchant credit acceptance decisions, $(A_1, A_2) \in \{0, 1\} \times \{0, 1\}$, solve the merchant price-setting profit maximization problem by finding a solution to the system of first order conditions found in step 4, subject to the constraint that the second order conditions found in step 5 are strictly satisfied. Repeat this step multiple times from multiple starting values to ensure that the solution found is the solution to the merchants' price-setting profit maximization problem.

6.2 Formulate the profit of each merchant at the optimal prices found in step 6.1, and determine the Nash equilibrium or equilibria of the merchant credit acceptance subgame.[20]

7. Repeat step 6 over a grid of possible values of the merchant discounts and cash-back bonus, (c_1, c_2, β).

8. At each point evaluated in step 7, calculate the credit platform's profit.

9. Determine the credit platform's optimal choice of merchant discounts and the cash-back bonus, (c_1^*, c_2^*, β^*), by selecting the choice that results in the highest profit in step 8. If (c_1^*, c_2^*, β^*) is close to the edge of the grid explored in step 7, return to step 7 and expand the grid.

By following these steps, we are able to obtain equilibrium outcomes for the chosen setting and parameters. By doing so for both settings, and then comparing these outcomes, we are able to examine the effects of ending the No Surcharge Rule and allowing merchants to surcharge.

For the numerical results presented in the remainder of Section 1.4, merchant and wallet preferences were independently and identically distributed according to a uniform distribution with a width of 5, customer valuations were distributed uniformly between 35 and 43, and the merchant discount for debit transactions was $d_1 = d_2 =$ \$0.23, a value chosen to be approximately equal to the merchant discount for debit card transactions. The marginal cost of merchant 1 was held constant at $\mu_1 =$ \$38, a value approximately equal to the average card purchase. When considering symmetric

[20]If there are multiple equilibria, the credit platform is able to select alternative merchant discounts to select the equilibrium it finds optimal. An example of this will be discussed in Section 1.4.2.1.

merchants, $\mu_1 = \mu_2 = 38$, and when considering asymmetric merchants, $\mu_1 = 38 > \mu_2 = 36$.

Equilibria were estimated over a grid spanning $(\$0.00, \$3.00) \times (\$0.00, \$3.00)$ for merchant discount rates (c_1, c_2), and $(-0.05, 0.05)$ for the cash-back bonus β.[21] For each grid point, merchants' optimal prices, $\widehat{\boldsymbol{P}}_j^k$, were found by minimizing the sum of squares of the first order conditions of the merchant problem, subject to the constraints that each individual first order condition was zero and the second order conditions were strictly satisfied.[22] At many grid points, finding a solution required attempts from multiple starting values. Starting values were obtained by following a three-dimensional recursive algorithm, using successful solutions from neighboring points on the grid as starting values for new attempts and returning to unsuccessful points on the grid after a neighbor's success. All estimation was implemented in C.

1.4.2 Merchant Credit Acceptance

The first equilibrium outcome we will examine is the merchants' decision of whether to accept credit. Before doing so in the context of the full model, it is useful to first consider a simplified environment. In this spirit, we will first examine a setting in which the credit platform is constrained from paying a cash-back bonus to customers who pay with the credit card, or from charging them a fee. Even though $P_j^D = P_j^C$ under the No Surcharge Rule, the effective price paid by a credit customer is still generally different from the price paid by a debit customer when $\beta \neq 0$ because $P_j^D \neq (1-\beta)P_j^C$. In this setting with $\beta = 0$, however, credit and debit customers pay

[21]Recall that a customer purchasing with the credit card for price P_j^C receives an amount βP_j^C from the credit platform if $\beta > 0$ and pays an amount βP_j^C to the credit platform if $\beta < 0$.

[22]Alternative implementations using fewer constraints were tested. Each attempt required less time using these alternative approaches, but was less likely to find a solution. As a result, a greater number of attempts was required, increasing total estimation time.

the same effective price under the No Surcharge Rule. So while holding β constant at zero is unlikely to be the optimal choice for the platform, examining equilibrium outcomes under this restriction makes it simple to see how different choices by the platform of merchant discounts translate into different credit acceptance decisions for the merchants, which in turn makes it simple to see a few key aspects about the platform's optimal choice of merchant discounts as well.

1.4.2.1 Without Cash-Back

In this section, we impose the restriction that $\beta = 0$. In general, we are interested in outcomes that are a subgame perfect Nash equilibrium of the full model presented in Section 1.3 in which the credit platform selects β in addition to the merchant discounts. However, outcomes described in this section as a subgame perfect Nash equilibrium, while they are a subgame perfect Nash equilibrium in this restricted setting with $\beta = 0$, are only a subgame perfect Nash equilibrium in the general setting if the credit platform's optimal choice is $\beta^* = 0$. Credit acceptance decisions of the merchants that are a Nash equilibrium of the merchant credit acceptance subgame are a Nash equilibrium in both settings. With this in mind, we will turn our attention to Figure 1.5.

Figure 1.5 shows four graphs. The top row is under the No Surcharge Rule (labeled NSR), while the bottom row is with surcharging (labeled SUR). The right column shows the subgame perfect Nash equilibrium (labeled SPNE) outcome of the merchant credit acceptance subgame for given marginal costs of the credit platform (μ_c^1,μ_c^2). The left column shows the Nash equilibrium of the merchant credit acceptance subgame (labeled CANE for credit acceptance Nash equilibrium) for given merchant discounts (c_1, c_2), both those that the credit platform finds optimal and chooses as part of a SPNE as well as those that it does not.

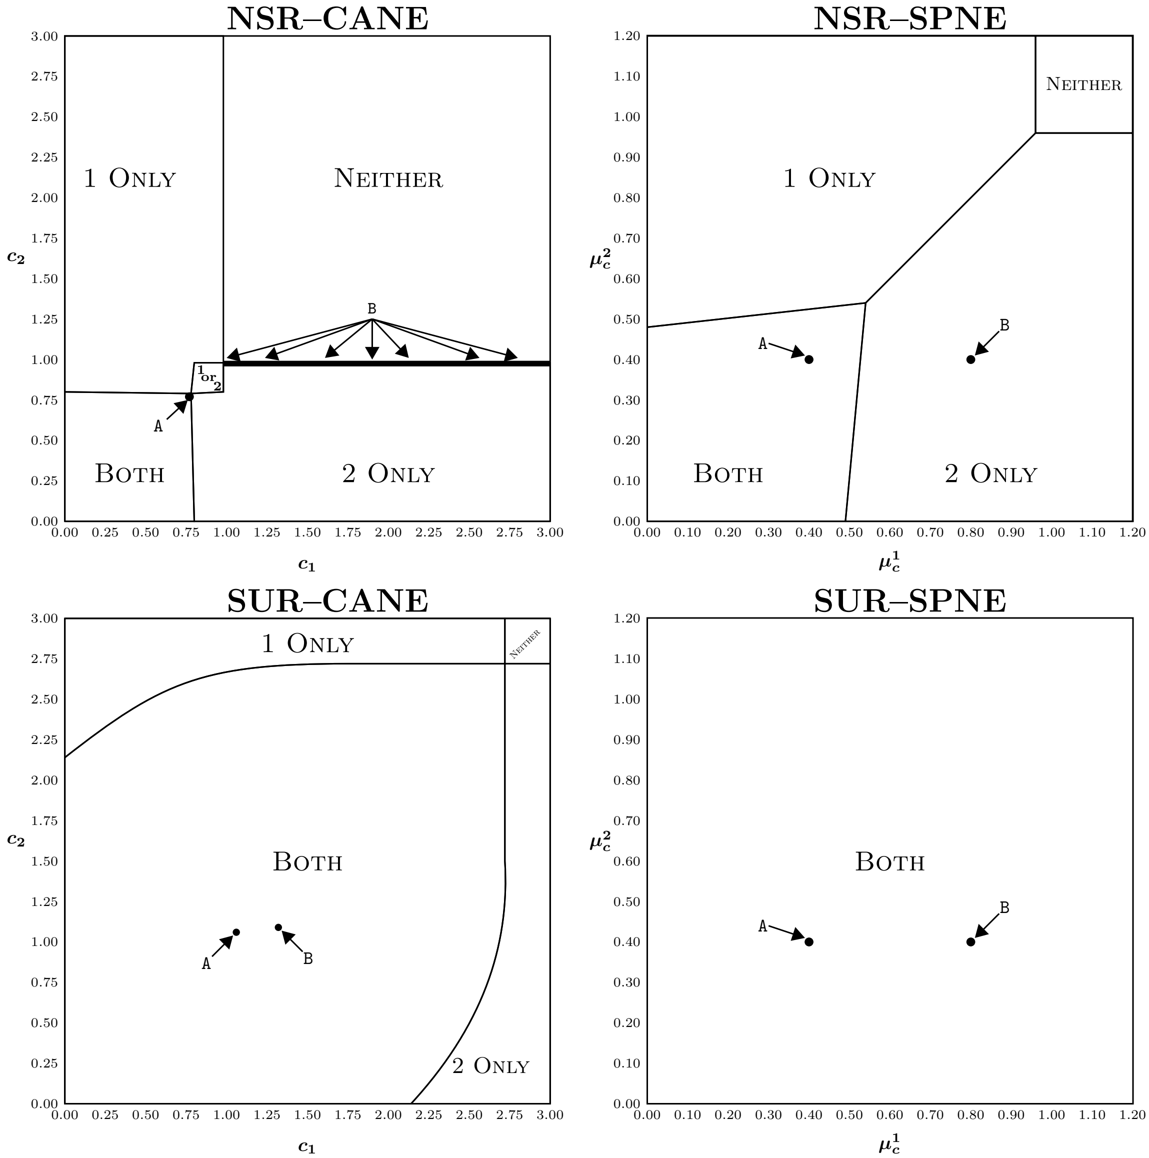

Figure 1.5: Merchant credit acceptance decisions with $\beta = 0$. SPNE=subgame perfect Nash equilibrium. CANE=Nash equilibrium of the merchant credit acceptance subgame. NSR=under the No Surcharge Rule. SUR=when surcharging is allowed. A full description is provided in Section 1.4.2.1.

In each of the four graphs in Figure 1.5, there are between one and five regions labeled with the Nash equilibrium of the merchant credit acceptance subgame. The label BOTH corresponds with the $(1, 1)$ Nash equilibrium in which both merchants accept credit, while NEITHER corresponds with the $(0, 0)$ Nash equilibrium in which neither merchant accepts credit. The 1 ONLY and 2 ONLY labels correspond with the $(1, 0)$ and $(0, 1)$ Nash equilibria, respectively. The small region labeled 1 OR 2 in the NSR–CANE graph is a region of merchant discounts for which both the $(1, 0)$ and $(0, 1)$ Nash equilibria of the merchant credit acceptance subgame exist. However, it is always possible for the credit platform to select merchant discounts outside of this 1 OR 2 region, as we shall see shortly. We will consider first the top row of graphs that are under the No Surcharge Rule, and then examine the bottom row where surcharging is allowed.

Consider the point labeled A in the NSR–SPNE graph. At this particular point, the credit platform's marginal costs are $(\mu_c^1, \mu_c^2) = (0.4, 0.4)$. This point lies in the region labeled BOTH. When the platform's marginal costs are 0.4 for transactions completed at each merchant, the solution to the platform's profit maximization problem is to select merchant discounts equal to 0.78 for each merchant. That is, the credit platform chooses $c_1 = c_2 = 0.78$. Given this, the merchants each determine the price they would set in each of the four possible combinations of merchant credit acceptance decisions, evaluate the profit they would get in each case, and then choose to accept credit or not so that their choices are a Nash equilibrium of the merchant credit acceptance game shown in Table 1.1. In this case, when $c_1 = c_2 = 0.78$, the outcome of this process is that both merchants find it optimal to accept credit. And based on this outcome, the profit-maximizing price set by each merchant, and the utility-maximizing decisions of customers, the platform determines that the profit it achieves by setting $c_1 = c_2 = 0.78$ is the highest profit achievable when it's marginal costs are

$\mu_c^1 = \mu_c^2 = 0.4$. The point labeled `A` in the NSR–CANE graph corresponds with this optimal choice of merchant discounts for the credit platform when its marginal costs are given by point `A` in the NSR–SPNE graph.

It warrants mentioning that point `A` is the upper- and right-most point in the BOTH region. This is generally the case under the No Surcharge Rule because causing merchants to stop accepting credit is the binding constraint on the credit platform rather than traditional quantity effects. Ignore merchant 2 for a moment and consider the effects of a small increase in c_1. Recall that c_1 is effectively the price received by the credit platform, but is a cost for merchant 1. If the credit platform increases c_1 by a small amount, it faces effects analogous to traditional price and quantity effects; in response to the increase in its costs (c_1), merchant 1 increases its price slightly, causing its quantity to decrease slightly. The credit platform has increased revenue from receiving a higher price (c_1), but decreased revenue from the decrease in quantity. If merchant 1 always accepted credit, balancing these two effects would be how the credit platform would optimally set c_1. However, in this setting, there is another potential effect of increasing c_1, and that is that merchant 1 will stop accepting credit. The effect of this decrease in quantity (to 0) clearly outweighs the slight decrease in quantity from the traditional quantity effect. The same is true when setting c_2 for merchant two. Together, the result is that the credit platform, when it determines that it can achieve the highest profit when both merchants accept credit, sets the merchant discounts as high as possible without causing either merchant to stop accepting credit. We will see a similar result when the credit platform finds it optimal to have only one merchant accept credit.

Now consider the point labeled `B` in the 2 ONLY region of the NSR–SPNE graph. At this point, the platform's marginal costs are $(\mu_c^1, \mu_c^2) = (0.8, 0.4)$. Compared to point `A` just discussed, the platform's marginal cost of transactions completed at merchant 1 has doubled, while the marginal cost of transactions completed at merchant 2 has remained the same. As μ_c^1 increases while μ_c^2 remains constant, a point is reached at which transactions completed at merchant 1 become sufficiently costly for the credit platform that it becomes optimal to set merchant discounts such that only merchant 2 accepts credit. The credit platform maximizes profits by setting $c_2 = 0.97$ and $c_1 \geq 0.98$. Because transactions completed at merchant 1 are sufficiently costly for the platform, it wants to set the merchant discount for merchant 1 high enough that merchant 1 does not want to accept credit.

If the credit platform were to set c_1 between 0.8 and 0.97 (with $c_2 = 0.97$), then (c_1, c_2) would lie in the 1 OR 2 region where multiple Nash equilibria of the merchant credit acceptance subgame are possible. The credit platform has higher profit when only merchant 2 accepts credit and thus wishes to avoid the equilibrium in which only merchant 1 accepts credit. Any uncertainty from having both outcomes possible, however, is easy for the credit platform to avoid by simply setting c_1 sufficiently high that merchant 1 does not find it optimal to accept credit. Any value $c_1 \geq 0.98$ achieves this goal, and thus there are a continuum of merchant discounts consistent with SPNE.

The choice of $c_2 = 0.97$, however, is unique. If the credit platform were to set c_2 any higher, merchant 2 would no longer find it optimal to accept credit and the platform's profit would be 0 (because the quantity of credit transactions would be 0, as neither merchant would accept credit). The platform could set c_2 lower and remain in the 2 ONLY region. However, just as the optimal choice of discount rates when the platform wants both merchants to accept credit was the upper- and right-most point in the BOTH region, the optimal choice when the platform wants only merchant 2 to accept credit is the upper boundary of the 2 ONLY region. The binding constraint for the platform on increasing the merchant discount in each case is the complete loss of revenue resulting from the merchants no longer accepting credit, rather than the slight decrease in revenue from a traditional quantity effect.

No points are labeled in Figure 1.5 in the 1 Only region, but the outcome and intuition are analogous to that for points in the 2 Only region. When the platform's marginal costs are such that it is optimal for the platform if only merchant 1 accepts credit, the credit platform sets the merchant discount for merchant 1 as high as possible without causing merchant 1 to stop accepting credit and sets the merchant discount for merchant 2 sufficiently high that merchant 2 will not find it optimal to accept credit. In the NSR–CANE graph, this is the right border of the 1 Only region with c_2 above the 1 or 2 region.

If the credit platform's marginal costs are sufficiently high that any merchant discount high enough to result in positive margin is so high that the merchant does not accept credit, then it is faced with the choice of earning negative profit by lowering the merchant discounts to the point where one or both merchants accept credit, or earning zero profit by setting the merchant discounts sufficiently high that neither merchant accepts credit. Thus, if the platform's marginal costs lie in the Neither region of the NSR–SPNE graph, it will set merchant discounts anywhere in the Neither region of the NSR–CANE graph, neither merchant will accept credit, and the platform's profit will be zero.

In the preceding discussion under the No Surcharge Rule of how the credit platform optimally sets merchant discounts, given its marginal costs, and the SPNE credit acceptance outcome that results, we have characterized the SPNE credit acceptance outcome for any possible value of the platform's marginal costs. This is because under the No Surcharge Rule, the binding constraint for the platform on increasing the merchant discount is the complete loss of revenue resulting from the merchant no longer accepting credit rather than the slight loss of revenue from a slight decrease in quantity as occurs with a traditional quantity effect. Consequently, the optimal merchant discounts lie on the borders of the Both, 1 Only, or 2 Only regions.[23] This occurs because under the No Surcharge Rule, merchants are constrained in how they are able to optimize in response to an increase in the merchant discount. Since the price must be the same for all customers, any increase in price for credit customers must also be an increase in price for debit customers, even though the cost for customers paying by debit has not changed. The only response available to merchants that is directed at only customers paying by credit is the all or nothing action of credit acceptance.

When the No Surcharge Rule is lifted, merchants are no longer restricted in how they can respond to an increase in the merchant discount. If the cost of transactions completed by customers paying by credit increases, the merchant has the option to increase the price paid by those customers, and only those customers, if he finds it optimal to do so. The merchant could also choose to spread the cost increase over all transactions by increasing the price paid by debit customers. And the option to stop accepting credit altogether always remains. However, when the subtler, continuous response of a slight increase in price is available, merchants generally find it preferable to the all or nothing response of no longer accepting credit. This can be seen in the bottom row of graphs in Figure 1.5.

For easy comparison between settings, the points labeled `A` and `B` in the SUR–SPNE graph are for the same platform marginal costs as in the NSR–SPNE graph. Point `A` is in the region labeled Both, as it was under the No Surcharge Rule, while point `B` has switched from 2 Only to Both. In fact, all points over the range of marginal costs shown are now in the Both region. This is because the merchant

[23]Or, if the platform's marginal costs are sufficiently high that zero profit is the best it can do, then the optimal merchant discounts lie anywhere in the Neither region.

discount at which the merchants no longer find it optimal to accept credit is much higher when they are able to respond to a higher discount rate by increasing the price paid by credit customers. Under the No Surcharge Rule, at the point when its costs were those given by point `A`, it found it optimal to set the merchant discounts both equal to 0.78, the highest possible value for which both merchants accept credit. When merchants are allowed to surcharge, the merchant discount at which they stop accepting credit is much higher. Consequently, the merchants no longer accepting credit is no longer the binding constraint on the optimal choice of merchant discounts for the credit platform. Rather, the platform's choice is guided by standard price and quantity effects, setting merchant discounts to balance the two. For point `A`, the optimal choice for the credit platform is $(c_1, c_2) = (1.06, 1.06)$. This is higher than under the No Surcharge Rule, but well away from the point at which the merchants would stop accepting credit. If the platform's marginal cost for transactions completed at merchant 1 increases to 0.8 (point `B`), then the optimal choice of discount rates increases to $(c_1, c_2) = (1.32, 1.09)$. This is a small increase in c_2 because the marginal cost of transactions completed at merchant 2 has not changed. The increase in c_1 is much larger. However, the value chosen by the credit platform is still well below the value that would cause merchant 1 to stop accepting credit.

1.4.3 With Cash-Back

When the restriction that $\beta = 0$ is removed and the platform is able to choose β optimally in addition to the merchant discounts, the intuition for the platform's optimal choices remains similar, as it does for the credit acceptance decisions of the merchants. Graphs for $\beta \neq 0$ of the merchant credit acceptance Nash equilibrium as a function of the merchant discounts look very similar to those shown in Figure 1.5 for $\beta = 0$. If we were to look at the graph corresponding to each value of β

chosen in SPNE by the credit platform, what we would notice is that the acceptance regions would look very similar to those in Figure 1.5, with the borders between them shifting further out as β increases. As β increases, the overall desire of customers to pay by credit increases as well, and thus, so too does the threshold merchant discount at which credit acceptance ceases to be optimal. Under the No Surcharge Rule, it remains optimal for the platform to choose merchant discounts on the border between acceptance regions, just as it did when $\beta = 0$. And with surcharging, the optimal choice remains far away from the borders.

When $\beta = 0$, we saw that for the entire range of platform marginal costs, it is always optimal for the credit platform to induce both merchants to accept credit. This remains the case when the platform chooses β, meaning that the SUR–SPNE graph of Figure 1.5 is representative of all SUR–SPNE graphs we could draw with surcharging. As a result, Figure 1.6 includes only SPNE graphs under the No Surcharge Rule. Consider first the graph on the left, which shows the SPNE in the case when merchants are symmetric, as was the case Figure 1.5. If the platform's marginal cost is much higher for transactions completed at one merchant or the other, then it remains optimal to induce that merchant to not accept credit. However, now that the platform is able to choose β, the range of marginal costs for which the platform finds it optimal for both merchants to accept credit is greatly expanded. When the merchants have similar marginal costs, it is likely that the platform's marginal cost of transactions completed at each merchant will be similar as well. Thus, outcomes in the BOTH region seem most likely. When merchants are asymmetric (see the graph on the right side of Figure 1.6), the region for which inducing only the high cost merchant, merchant 1, to accept credit, is much smaller. Only if transactions completed with the higher cost merchant are significantly cheaper for the platform does the platform wish to induce acceptance by only that merchant, a situation that seems highly unlikely.

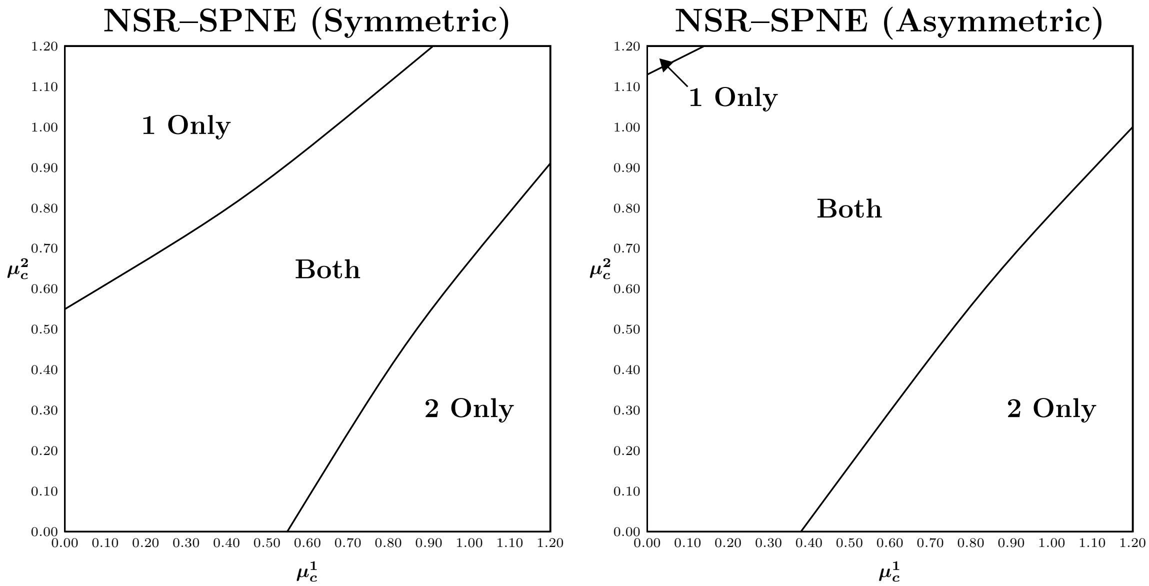

Figure 1.6: Subgame perfect Nash equilibrium merchant credit acceptance decisions with $\beta = \beta^*$ when merchants are symmetric ($\mu_1 = \mu_2 = 38$) and asymmetric ($\mu_1 = 38 > \mu_2 = 36$)

1.4.4 Quantity of Credit Transactions

Let us now turn our attention to three main equilibrium outcomes of interest. First, consider the total quantity of credit transactions. Most models on payment card systems that investigate the effects of lifting the No Surcharge Rule find that the total quantity of credit transactions is lower when merchants are allowed to surcharge. The same results holds true here. This is shown for the case of symmetric merchants in Figure 1.7. The left graph shows the total quantity of credit under the No Surcharge Rule, the middle when surcharging is allowed, and the right shows the difference between the two (i.e., $\widehat{\boldsymbol{Q}}_1^C + \widehat{\boldsymbol{Q}}_2^C$ with surcharging minus $\widehat{\boldsymbol{Q}}_1^C + \widehat{\boldsymbol{Q}}_2^C$ under the No Surcharge Rule). The x- and y-axes are the platform's marginal costs, just as in the SPNE graphs discussed in the previous section. Recall that there is a unit mass of customers, each of whom purchases either one or zero of the good for sale. Thus, the magnitude in the z direction, $\widehat{\boldsymbol{Q}}_1^C + \widehat{\boldsymbol{Q}}_2^C$, is the faction of all customers who both purchase a good and do so using credit.

The distinct regions clearly seen in the No Surcharge Rule graph on the left correspond with the different merchant acceptance regions shown in the left graph of Figure 1.6. The parallel lines on the surface of the graph run perpendicular to the axis corresponding with the platform's marginal cost at that merchant, making the Both region appear as if it is covered with a grid. As we saw in Figure 1.6, the platform finds it optimal to induce both merchants to accept credit when surcharging is allowed. Consequently, there is only one region in the graph depicting the situation with surcharging. This enables us to use the same convention as in the No Surcharge Rule graph to indicate the merchant acceptance regions under the No Surcharge Rule in the final graph that shows the change between the situation with surcharging and under the No Surcharge Rule. The entire third graph is negative values, indicating that the total quantity of credit transactions decreases when surcharging is allowed.

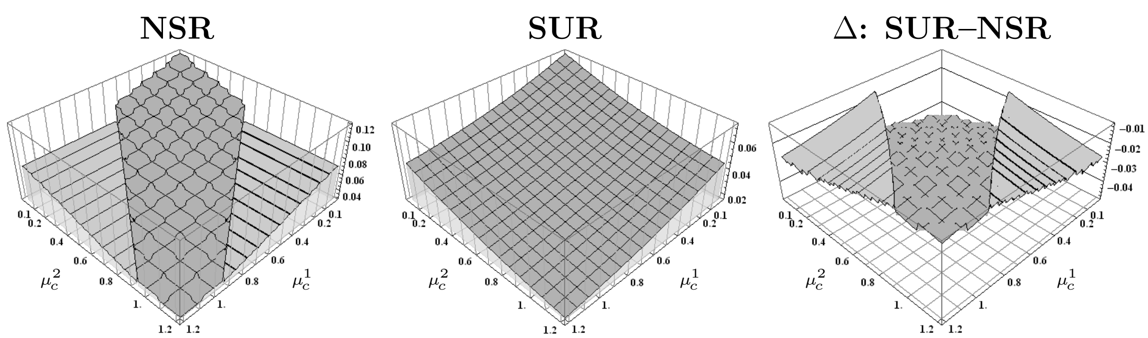

Figure 1.7: Quantity of credit transactions (symmetric merchants)

While the result that the total quantity of credit transactions decreases when surcharging is allowed is similar to what is found in other models examining the effects of lifting the No Surcharge Rule, the intuition for this result is different. The decrease in the quantity of credit transactions when surcharging is allowed is often explained by differences in the prices set by merchants; credit transactions are more expensive for merchants under the No Surcharge Rule, so when the No Surcharge Rule is lifted, merchants raise the price they charge customers who pay using the credit card, thus decreasing the quantity of credit card transactions. However, in this case the decrease in the quantity of credit card transactions is driven by changes in the optimal fee structure set by the credit card platform. Under the No Surcharge Rule the platform finds it optimal to earn profits from the merchant side of the market via higher merchant discounts, increasing the volume of transactions by paying customers a reward for using the credit card (e.g., "cash-back"). However, when surcharging is allowed, the platform finds it optimal to stop paying customers rewards and to instead charge customers a fee for using the credit card. This causes many customers to stop using the credit card. To increase the volume of credit transactions, the platform lowers the merchant discounts. This causes merchants to lower the price they charge customers paying by credit, which partially offsets the decrease in credit transactions caused by the platform charging customers a fee. However, the decrease in credit transactions brought about by the cardholder fee outweighs the increase in credit transactions brought about by merchants lowering the price, leading to an overall decrease in the quantity of credit transactions.

1.4.5 Quantity-Weighted Average Price

Another main finding in the literature examining the effects of lifting the No Surcharge Rule and allowing merchants to surcharge is that there is an overall decrease in the prices paid by consumers. This occurs here as well, but again, with a slightly different explanation. The standard expectation is that when the No Surcharge Rule is lifted, merchants raise the price for credit card customers and lower the price for debit card customers. This tends to decrease the quantity-weighted average price by lowering the price paid by customers who use a debit card and, in addition, by lowering the fraction of customers who use the more-costly-for-the-merchant credit card. In this case, however, the platform lowers the merchant discounts when surcharging is allowed, prompting merchants to lower the price for customers paying by credit card. In addition, merchants also lower the price for customers paying by debit card. As merchants lower the price paid by both credit and debit customers, the quantity-weighted average price is higher under the No Surcharge Rule than when surcharging is allowed. This can be clearly seen in Figure 1.8.

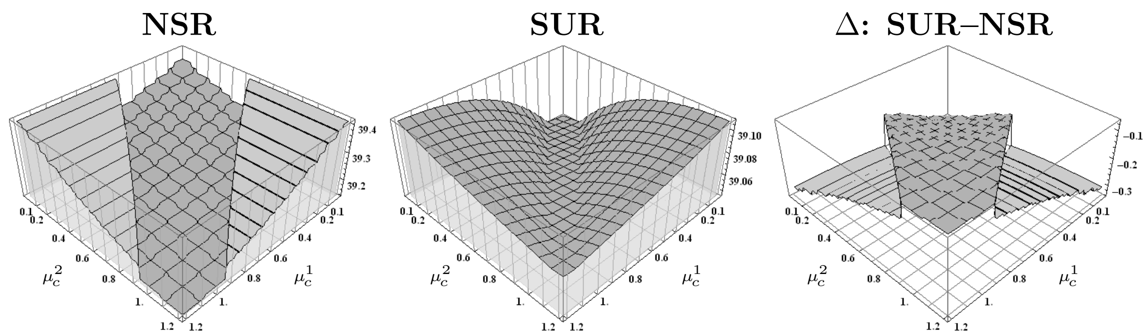

Figure 1.8: Quantity-weighted average price (symmetric merchants)

1.4.6 Merchant Profits

So far, the outcomes of interest occurring in equilibrium in this model have corresponded closely with what has been found previously in the literature. When we examine the effect of surcharging on merchant profits when merchants are symmetric, this will continue to be the case.

1.4.6.1 Symmetric Merchants

Consider first the case when merchants have the same marginal cost, shown in Figure 1.9. The columns are the same as in the previous sections. The top row

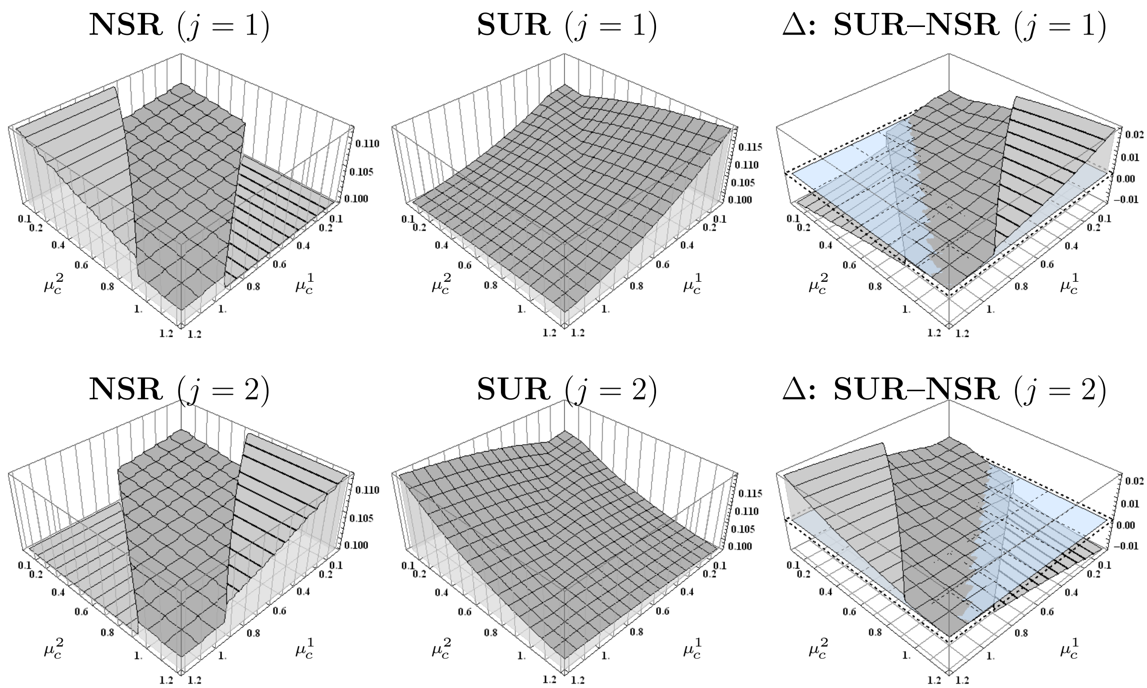

Figure 1.9: Merchant profits (symmetric merchants)

shows the profit of merchant 1, while the second row shows the profit of merchant 2. In the third column, a horizontal plane is included in the graphs at z=0. Thus, wherever the graph is above the plane, profits increase with surcharging, and wherever it is below the plane, profits decrease. At first glance, the results appear ambiguous, with part of the graph above the plane and part below. However, not all points in the graph are equally likely to occur in reality. Recall that the x- and y-axes are the platform's marginal costs of transactions completed at merchant 1 and 2, respectively. The region of the top-right graph that is below the plane is the region of platform marginal costs where transactions completed with merchant 1 are significantly more costly than transactions completed at merchant 2. Taken out of context, this could seem plausible or not. However, recall that the two merchants are identical. If they both have the same marginal cost of providing the good ($\mu_1 = \mu_2$), it seems likely that

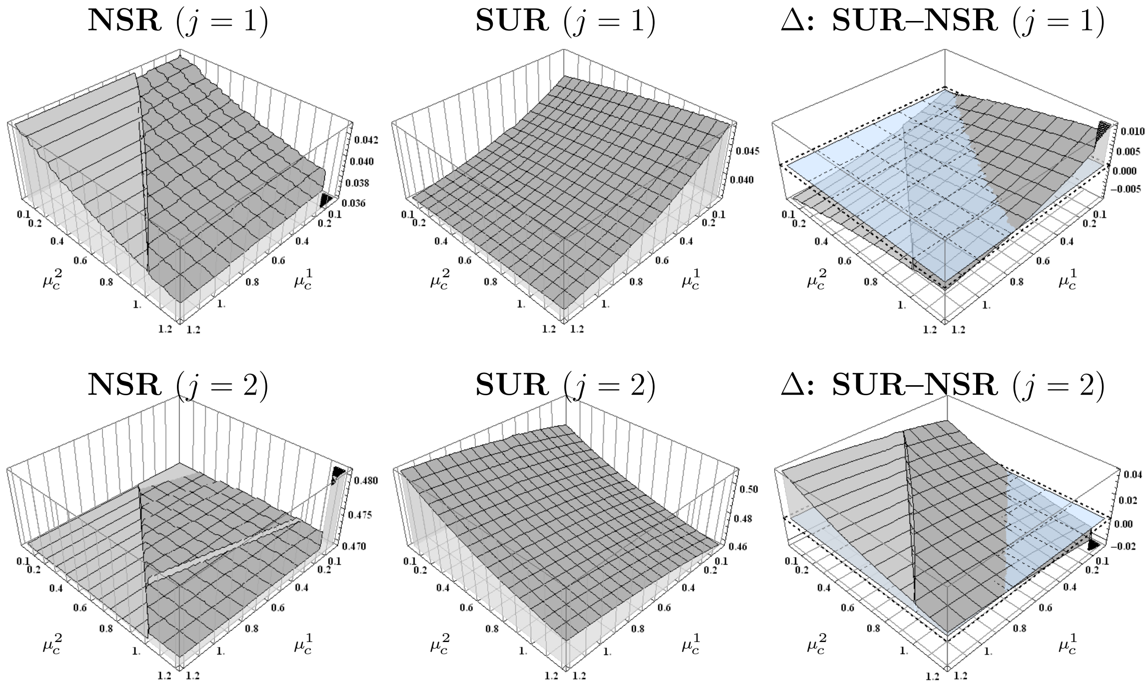

Figure 1.10: Merchant profits (assymmetric merchants)

their transactional costs of credit card transactions would be similar as well. Thus, the points along the 45 degree line through the center of the graph seem most likely to occur in the real world. Near this 45 degree line, the graph is above the plane, and thus, it seems most likely that both merchants have higher profits when allowed to surcharge.

1.4.6.2 Asymmetric merchants

Now consider the case when the two merchants are asymmetric in terms of their marginal costs. This is shown in Figure1.10. It seems reasonable to believe that the platform's marginal costs would be higher for transactions completed at merchant 1, the merchant with higher overall marginal costs, rather than at merchant 2, the merchant with lower overall marginal costs. At most, we might believe that the marginal costs for the platform should be roughly equal. It does not seem reasonable to believe that it would be more costly to complete transactions with the lower-cost merchant. Thus, points to the left, or, perhaps, on the 45 degree line, seem most likely to occur in the real world. For these points, merchant 2's graph is above the plane, indicating higher profits from surcharging. However, the graph for merchant 1 is below the plane, indicating that merchant 1, the smaller merchant, has lower profits when allowed to surcharge.

1.4.6.3 Effects of Platform's Optimal Choices on Merchant Profits

As discussed in Section 1.4.5, the removal of the No Surcharge Rule causes the platform to switch from earning profits mostly from the merchant side of the market to the customer side of the market. The platform stops giving customers rewards for using the credit card and instead starts charging them a fee. This causes many customers to stop using the credit card. In order to increase the volume of credit card transactions, the platform lowers the merchant discounts. This causes merchants to lower prices, thereby increasing volume. When the merchants are symmetric, the platform lowers the merchant discounts symmetrically. In response to this reduction in their costs, both merchants lower their prices, sell a higher quantity of goods, and, in equilibrium, end up with higher profits.

When merchants are asymmetric, the platform no longer lowers merchant discounts symmetrically. The goal of lowering the merchant discount is to increase the volume of credit transactions. In order to increase the volume of transactions, the platform finds it optimal to decrease the merchant discount by more for the higher-volume, lower-cost merchant. While this benefits the lower-cost merchant, the higher-cost merchant was already at a competitive disadvantage under the No Surcharge Rule, which is made worse by the platform lowering costs by more for the lower-cost merchant. Consequently, the lower-cost merchant has higher profits when surcharging is allowed, but the higher-cost merchant has lower profits.

1.5 Conclusion

This paper examined the effects of merchants being allowed to place a surcharge on transactions completed using payment forms that are more costly to the merchants. The model combined aspects of several strands of the economics literature on the No Surcharge Rule, while relaxing a few key limiting assumptions made in each. Customers are heterogeneous along three key dimensions, having varying preferences for both merchants and payment forms as well as different valuations for the good. Merchants compete not only on price, but also on the acceptance of the credit card. In addition, allowing the merchants to be different means that the same choices are not always optimal for both in equilibrium, adding to the richness of the interactions between the parties.

In this more general framework, I find results similar to those found elsewhere in the literature for most outcomes of interest, with one important difference. Surcharging of credit card transactions reduces the quantity of credit card transactions and the quantity-weighted average price. Merchants, in most cases, have higher profits when allowed to surcharge. Because merchants have been quite vocal over many years about their desire to be allowed to surcharge, finding they have higher profits when allowed to do so is to be expected.

However, by relaxing the common assumption that merchants are identical, an important alternative outcome was revealed. When merchants differ in their marginal costs, they are affected differently by surcharging. The merchant who is able to reduce marginal costs, as might be the case with a large retailer, benefits from surcharging.

However, a merchant competing with this larger retailer that does not have the same ability to reduce marginal costs, typically a smaller business, often has lower profits when allowed to surcharge.

Chapter 2

Worth the Wait? Cooperation in a Repeated Prisoner's Dilemma with Search

2.1 Introduction

When can sorting tip the balance in favor of maximizing the mutual gains from collaboration in spite of conflicting personal incentives? When the first-best outcome in which all individuals in a population cooperate is not feasible because some individuals in the population choose not to cooperate, is it possible that cooperation by only some individuals can improve the welfare of those who do choose to cooperate, or perhaps, improve the welfare of all, despite a fraction of the population choosing not to cooperate? The literature examining cooperation in a repeated prisoner's dilemma situation has largely focused on sustaining cooperation between all individuals in a population. Different ways of sustaining cooperation under different informational structures have been found, each striking a different balance between complexity and robustness to real world challenges. When the actions of all are observable by all, sustaining cooperation is relatively simple (e.g., standard folk theorem results). When observability is limited, cooperation can still be sustained despite the imperfect information, but often at the expense of strategic simplicity and with a greater need for a common understanding of and desire to implement coordinated punishments (e.g., Kandori (1992), Ellison (1994)). Largely unexamined is the question: when the first-best outcome, in which all individuals cooperate, is not feasible, can welfare be improved over the worst possible outcome, in which no individuals cooperate?

To examine this question I consider a population of individuals with different rates of time preference interacting in a repeated prisoner's dilemma. The appeal of the prisoner's dilemma as a tool for modeling human behavior in strategic situations lies in the tension it exemplifies between rational individuals acting to maximize their own utility and the inefficient outcome that results in spite of this optimizing behavior. A repeated prisoner's dilemma is embedded in a search market model. Individuals fully observe the actions taken in a current match but nothing else. Players follow uncomplicated strategies and there is no informational structure to facilitate labeling of individuals or public devices to coordinate punishments. The fact that cooperation is sustained despite the simplicity of strategies followed by players and limited reliance on complicated information structures or common understanding or even desire to implement coordinated punishments is attractive when considering facilitating cooperation in real-world situations.

When all players find it optimal to cooperate, there is no need to examine ways to increase welfare because welfare is already maximized. The first-best outcome in which all players cooperate is feasible if the least patient player finds cooperation optimal. After establishing this baseline of comparison, I examine how partial and full separation can improve welfare when the first-best outcome is not feasible.

In this paper I establish when cooperation by only some individuals in the population can improve the welfare of all, both cooperators and non-cooperators. When the first-best outcome is not feasible, separation by action within a market can improve the welfare of all individuals over the fully uncooperative outcome. Players who do not choose to cooperate have the opportunity to take advantage of those who do. And despite sometimes finding themselves matched with individuals who do not cooperate, cooperators find themselves better off as well, as long as a sufficiently large fraction of the population find cooperation to be optimal. No matter how impatient are the low type, cooperation can still be sustained among the more-patient players as long as enough of the population are the more-patient type, providing higher expected utility for all. This allows welfare to be improved for all players, even when some players are very impatient or even completely myopic.

If enough of the population do not find cooperation to be optimal in the presence of defectors, the only one-market equilibrium consistent with individual rationality is the fully uncooperative equilibrium. In these situations, I find that some individuals can be made better off, without making the rest of the population worse off, by using time, a resource possessed by all, to separate players with different levels of patience. Specifically, a second, slower market is introduced, allowing the more-patient types to separate themselves from the less-patient types. More-patient players find it worth the wait to enter the slower market to wait for a cooperative match, while less-patient players do not, preferring to receive an uncooperative match more quickly. This ability to fully separate themselves from the less-patient players improves the welfare of the more-patient players without making the less-patient players worse off. This Pareto improvement opportunity is not sensitive to the fraction of the population that is each type and can allow for cooperation to be sustained among more-patient individuals, even if they make up a small fraction of the population and even if the risks associated with cooperation are significantly higher than the maximum level sustainable with one market.

The second market unambiguously improves welfare when no cooperation is sustainable with only one market. However, when conditions are such that partially-cooperative equilibria with separation-by-action within one market and full separation across two markets exist simultaneously, the less-patient players are always better off with only one market. This is because when all players must meet in a single market, the less-patient players are able to take advantage of the more patient players and receive the temptation payment whenever they are matched with a player who cooperates, and when they are matched with another less-patient player who defects, they receive exactly what they would with full separation across two markets.

The other side of the less-patient types' gain from taking advantage of the more-patient types when all players co-exist in one market is of course losses experienced by the more-patient types when matched with a player who defects. This risk can be avoided entirely through the full separation facilitated by the second market. However, the more-patient players sometimes prefer one market to two despite this risk, depending on how much longer they must wait to receive a match in the second, slower market. The welfare gains from introducing the second market when no cooperation is sustainable with only one market come from the second market being sufficiently slow that less-patient players are not tempted by it but not so slow that more-patient players do not find it to be worth the wait as a means to achieve separation from the less-patient players. If the second market does not have to be too slow, then the more-patient players prefer the separation it affords them while the less-patient players receive exactly what they would receive with repetition of the stage game Nash equilibrium, no longer receiving any gains from taking advantage of cooperating players. However, if the second market is too slow, even though full separation may be possible, it may not be desirable. If the slow market is sufficiently slow, the more-patient players prefer to co-exist with the less-patient players in a single market, despite the risks of being matched with a player who does not cooperate, because the loss of utility associated with waiting for a cooperative match in the slower market is too great. In this case, both types of players prefer an equilibrium in one market in which cooperation and defection exist simultaneously.

This paper proceeds as follows. Section 2.2 discusses implementations of cooperation in the prisoner's dilemma found in the existing literature. Section 2.3 presents the basic model, while Section 2.4 explores possible equilibria with one market. The model is extended to include two markets in Section 2.5 and possible welfare implications of this addition are analyzed. Because the notation is much simpler and much of the intuition for the model and results can be obtained in the simpler environment, the model presented in Sections 2.3 through 2.5 does not allow for players to end matches endogenously. However, the qualitative results remain unchanged when players are allowed to end matches endogenously. The similarities and few key changes to incentives when endogenous breakup is allowed are discussed in Section 2.6 and presented formally in Section 2.7. Section 2.8 concludes.

2.2 Cooperation in the Prisoner's Dilemma

It is well known that not cooperating is the dominant action in a one shot prisoner's dilemma. The simplest and earliest approach to facilitating cooperation is to consider two players who play a repeated prisoner's dilemma, forever. By standard folk theorem results (e.g., Fudenberg and Maskin (1989)), cooperation can be sustained if players are sufficiently patient. Kandori (1992) extends standard folk theorem results to the case of a population that plays forever but in which individuals are randomly matched. For example, suppose there is a market in which members of a large population are randomly matched each day to play a prisoner's dilemma stage game. Players cooperate against other players who have always cooperated and defect against players who have not. As long as there is an information structure to facilitate the labeling and thus punishment of defectors, cooperation can be sustained.

Kandori (1992) also considers random anonymous matching in a setting where labeling is not possible. He finds that despite the inability for direct punishment, cooperation is still sustainable by community enforcement. Extending the example to this setting, since it is no longer possible to only defect against defectors since they cannot be identified, now players cooperate until defected against, at which point they switch to defecting against all future matches. Once one player defects one time, defection spreads through the population and cooperation is lost forever, providing a harsh but effective punishment that sustains cooperation.

One objection to the contagion equilibrium of Kandori is the lack of forgiveness and lack of robustness to mistakes. On the equilibrium path efficiency is obtained, but if one player fails to cooperate even once, even if by mistake, cooperation is lost forever. Ellison (1994) considers a similar framework with anonymous random matching but introduces a public random variable on which players can coordinate. Continuing the example, players cooperate until defected against, and then defect against all future matches until some day it is raining when they show up to be matched. Once it rains, the punishment phase ends and cooperation is restored. If this does not provide strong enough incentives for cooperation, the punishment phase could be extended to ending only after two occurrences of rain, and so on. Thus the public signal allows for coordinated punishments and can be adjusted to provide optimal punishment intensity, strong enough to facilitate cooperation but not so strong that players do not follow through with punishment when it is their equilibrium strategy to do so. This public coordination device allows for near efficiency even with minimal noise and trembles. It does, however, increase the complexity of the strategies followed by the players and requires common understanding of the public signal and what actions are required in response to it.

We are interested in considering a framework similar to Kandori and Ellison with random matching in a population. However, we would like to limit the complexity required in terms of strategies followed by the players and the requirements for informational structures. One approach of relaxing the informational requirements is to consider imperfect monitoring. Fudenberg et al. (1994) extend folk theorem results to the case of imperfect public monitoring, as do Abreu et al. (1990) in a dynamic programming context, while Ely and Välimäki (2002) do so for private (but almost perfect) monitoring. We will instead assume players observe perfectly the actions taken by each player in a current match, but nothing else. The strategies followed will limit the cognitive burden placed on players, thus eliminating the requirement that players have a common understanding of public signals and the complex punishments coordinated on them. The result is facilitation of cooperation in a setting with limited informational requirements and appealingly simple strategic requirements. Players need only cooperate if they wish, and punish a current match if they defect.

The desire to simplify the strategic requirements placed on players is not unique to this work. Recent economic events have increased the attention given to behavioral economics in general, with one focus being on modeling the way individuals actually behave together with constraints that may exist on their cognitive abilities. The model presented here does not formally include any behavioral modifications, such as non-standard preferences or cognitive limitations.[1] However, in line with this research, the results derived here come from an environment that attempts to limit the memory recall required by players to actions taken in the current match, the need for complex common understanding of signals or other events, and the complexity of strategies needed to sustain cooperation, all of which serve the dual role of decreasing the cognitive requirements on players while at the same time reducing informational requirements.

The setting with one market is similar to Ghosh and Ray (1996) and Kranton (1996). However, in the model presented here, each period, players are only able to

[1] For an overview, see the book *Advances in Behavioral Economics* edited by Camerer et al. (2002), or the more recent article "Advancing Beyond *Advances in Behavioral Economics*" by Fudenberg (2006).

cooperate or not, rather than having the option to increase the level of cooperation, and thus payoffs from being in a cooperative match, over time. An alternative examination of sustaining cooperation in a repeated prisoner's dilemma with a similar focus on separation is an experimental paper by Bohnet and Kübler (2005). These authors conducted an experiment in which a group of individuals played a repeated prisoner's dilemma. Prior to playing the repeated stage game, subjects were given the opportunity to bid in an n-price auction. The n highest bidders then played a modified stage game that was more conducive to cooperation, basically providing insurance for cooperators against the losses of being defected against, while the remainder played a non-insured version of the repeated stage game. They found that if the number of spots in the modified game did not exceed the fraction of the population inclined to cooperate, then cooperation could be sustained by separating cooperators from non-cooperators. This paper will have a similar feel in terms of sustaining cooperation by separating types, but without requiring the resources necessary to essentially pay players to cooperate. Here stage game payoffs will not be altered, but rather cooperation will be facilitated by separation alone.

2.3 Environment

Consider a population consisting of two types of players who differ only by their discount rates. Player i discounts the future at rate δ_{t_i}. The high type of players, denoted $t_i = h$, are more patient players who discount the future at rate $\delta_{t_i} = \delta_h$. The low type of players, denoted $t_i = \ell$, are less patient and discount the future at rate $\delta_{t_i} = \delta_\ell < \delta_h$. Players are otherwise identical. The fraction of players who are the high type is given by π and is commonly known.

Each discrete period, a player is either matched or unmatched. The utility received by unmatched players is normalized to 0. Unmatched players can choose to stay unmatched or to enter the market. When an unmatched player enters the market he is matched randomly with another unmatched player with probability μ_m. Once matched, players play a repeated prisoner's dilemma stage game until the match dissolves exogenously, with probability β_m, or endogenously, when endogenous breakup is allowed. Matched players play the stage game, receive their utility payoffs for the period, and then learn if the match will dissolve that period or not. If the match does breakup, players start the next period in the unmatched state, choosing to either enter the market again or to stay unmatched. The timing is such that each period can be thought of as having three subperiods, with matching occurring in the first subperiod, stage game play and utility received in the second, and resolution of breakup uncertainty in the third. Matched players know their own action and payoff, as well as that of their match, for the duration of the match. However, once a match dissolves, players are unable to identify players with whom they have previously been matched.

Later, in Section 2.5, we will explore an environment with two markets. The two markets differ only by the search probabilities, μ_m and β_m, where $m \in \{s, f\}$ denotes the two markets. The slow market, $m = s$, is slower than the fast market, $m = f$. A player entering the slow market expects to wait longer before receiving a match ($\mu_s < \mu_f$). Once matched, the breakup probability in the new slower market is no higher than the breakup probability in the fast market ($\beta_s \leq \beta_f$). Before considering what happens with two markets, in Section 2.4, we will consider equilibrium when there is only one market. Because the two market environment nests the one market environment, to avoid duplication, the full environment with two markets will be presented now. To reduce the two-market setting to a setting with only one market, let $\mu_s = 0$ so that a player who enters the slow market never receives a match. Most of the model will be developed in terms of general market m, but all one market equilibria occur in the fast market, $m = f$.

Both the probability of breakup, β_m, and the matching probability, μ_m, are exogenously determined and are the same for all players in each period. Three versions of the model are explored, one in which all breakup is exogenous, one in which endogenous breakup is allowed each period, and a hybrid between the two.[2] Much of the intuition can be gained from exploring the model without endogenous breakup, and it requires less notation than and produces similar results to the model with endogenous breakup. Thus, the model without endogenous breakup will be developed first here and explored in Sections 2.4 and 2.5, while the similarities, as well as the few key differences, between the models will be discussed in Section 2.6. The formal model with endogenous breakup is presented in Section 2.7. The hybrid model (e.g., endogenous breakup is allowed with some probability each period) is the most flexible, nesting the other two models. However, because it is simply a weighted average of the models with and without endogenous breakup and requires even more notation, it will be discussed only briefly in Section 2.7.

When player i is matched with another player, they play a prisoner's dilemma game with possible actions *Cooperate* and *Defect*. Stage game utility payoffs are given in Table 2.1. Payoffs are finite and satisfy the standard prisoner's dilemma incentive structure with $(c + \tau) > c > d > (d - \lambda)$ so that D is the dominant action.[3] We will

[2]Note that there is still an exogenous chance of breakup, β_m, in all three models.

[3]The analysis presented here extends to the case where payoffs differ by type or differ for the row and column players as long as the prisoner's dilemma incentive structure is maintained. The payoffs given in Table 2.1 could be written as $(c^j_{t_i} + \tau^j_{t_i}) > c^j_{t_i} > d^j_{t_i} > (d^j_{t_i} - \lambda^j_{t_i})\ \forall t_i, j$, where j refers to row or column player to allow for two-sided markets. For $c^j_{t_i}, d^j_{t_i}, \tau^j_{t_i}, \lambda^j_{t_i} > 0\ \forall t_i, j$, the qualitative results are unchanged and thus for simplicity we will focus on the completely symmetric case shown in Table 2.1. The model could also be presented with several payoffs normalized, e.g., with $c = 2, d = 1$. However, the analysis is more transparent when it is possible to discuss $c - d$ directly as the value of cooperation relative to non-cooperation, so this normalization has not been made.

assume that $c, d, \tau, \lambda > 0$, but allow for $(d - \lambda)$ to be positive or negative. We will often consider the case where the temptation to defect, τ, is equal to the loss from cooperating against a player who defects, λ, but this assumption is not required.

	C	D
C	c, c	$d - \lambda, c + \tau$
D	$c + \tau, d - \lambda$	d, d

Table 2.1: Stage Game Payoffs

2.3.1 Strategy of player i type t_i

We are interested in comparing the static stage game Nash equilibrium strategy of always Defect with a strategy that involves cooperation. To reduce informational requirements and the cognitive burden placed on players, we are interested in having players follow strategies that are as simple as possible. Thus players who choose to cooperate instead of defect will follow the grim trigger strategy.

The main drawback to considering only these two strategies is the lack of forgiveness of the grim trigger strategy. It is a straightforward extension to allow for both a probability of accidentally defecting as well as a probability of forgiving an accidental defector. Parameters to capture these features would mostly appear as multipliers on the existing p_m term introduced shortly (e.g., a player's beliefs could be $(1 - \varepsilon)p_m$ instead of p_m). However, for simplicity we will assume that players do not make mistakes, making forgiveness of mistakes unnecessary.[4]

[4]Numerous other strategies, such as those that involve players cycling between cooperation and defection, are also possible. While we could certainly construct an equilibrium in which players cycle between cooperating for 17 periods and defecting for 11, it would require greater strategic complexity from the players and to little end. Furthermore, behavior such as this is not something we would expect to observe. We are also not considering strategies that require an additional information

The two strategies followed by the players are simple, which is appealing in its own due to the decreased cognitive burden it places on players as well as the limited informational requirements it places on the environment. Despite the simplicity of the strategies followed by players, however, these strategies allow for an interesting range of behavior. We will, of course, examine the standard cases in which all players cooperate or all players defect. But we will also examine situations in which cooperation exists simultaneously with defection; of interest is the fact that these situations can make players of both types better off. Thus we will start our analysis assuming that players follow one of these two strategies.

Formally, players who do not find cooperation optimal will follow the stage game strategy of always playing D. Players who choose to cooperate will follow a grim trigger strategy, σ_{gt}, defined by (2.1).

$$\sigma^{gt} = \begin{cases} C & \text{if neither player has played } D \text{ before in this match} \\ D & \text{if } D \text{ has been played before during this match} \end{cases} \tag{2.1}$$

Let $p_m \in [0,1]$ be the probability with which player i believes his match will follow the cooperative grim trigger strategy given by (2.1) in market m, and $1 - p_m$ be the probability with which he believes his match will play D.

An unmatched player chooses to enter market $m \in \{s, f\}$, with hopes of receiving a match, or to remain unmatched. The expected value of the decision

structure, such as coordination on public devices or labeling of defectors. Contagion equlibria, in addition to potentially imposing harsher punishments than the grim trigger, are ruled out by the assumption of a continuum of players, as the probability of being matched with the same player again, or with a player who has been matched with that player, is zero.

problem facing an unmatched player is[5]

$$V_{t_i} = \max\{0, V_{t_i}^s(p_s), V_{t_i}^f(p_f)\} \tag{2.2}$$

where 0 is the utility received from remaining unmatched and $V_{t_i}^m(p_m)$ is the utility the player expects to receive from entering market $m \in \{s, f\}$. When we examine equilibria with only one market, we will set the matching probability in the slow market to zero ($\mu_s = 0$) so that the value of entering the slow market is $V_{t_i}^s(p_s) = 0$, the same as the utility of remaining unmatched. In general, the utility the player expects to receive from entering market m is

$$V_{t_i}^m(p_m) = \overbrace{\mu_m W_{t_i}^m(p_m)}^{\text{receive match}} + \overbrace{(1-\mu_m)(0 + \delta_{t_i} V_{t_i})}^{\text{stay unmatched}} \tag{2.3}$$

$W_{t_i}^m(p_m)$ is the value player i type $t_i \in \{h, \ell\}$ expects to receive from a new match once matched, which occurs with probability μ_m. This expectation depends on the probability the player assigns to the likelihood of being matched with a player who chooses to cooperate.

The value of a match, $W_{t_i}^m(\sigma^{gt}|p_m)$, for player i type t_i choosing to follow the grim trigger strategy with the expectations p_m that his match will also follow the grim trigger strategy and $1 - p_m$ that he will not, is

$$\begin{aligned} W_{t_i}^m(\sigma^{gt}|p_m) = {} & \overbrace{p_m \left(c + \delta_{t_i}[(1-\beta_m) W_{t_i}^m(\sigma^{gt}|1) + \beta_m V_{t_i}]\right)}^{\text{matched with cooperator}} \\ & + \overbrace{(1-p_m)\left(d - \lambda + \delta_{t_i}[(1-\beta_m) W_{t_i}^m(\sigma^{gt}|0) + \beta_m V_{t_i}]\right)}^{\text{matched with non-cooperator}} \end{aligned} \tag{2.4}$$

After the first period of a match, players' actions reveal their strategies, and beliefs are updated accordingly to $p_m = 1$ if both players cooperated and to $p_m = 0$

[5] Note that V_{t_i}, the left hand side of (2.2), also depends on p_m. However, for notational simplicity and to avoid confusion later on, the p_m argument of V_{t_i} expressions will be omitted.

otherwise. This is reflected in the value expected in the second and subsequent periods of a match, $W_{t_i}^m(\sigma^{gt}|1)$ and $W_{t_i}^m(\sigma^{gt}|0)$. The second period of the match is reached only if the match does not dissolve, which occurs with probability $1-\beta_m$. If the match dissolves, which occurs with probability β_m, then the player re-enters the unmatched state.[6]

Similarly, the expected value of a match for a player with beliefs p_m who chooses to play D is

$$W_{t_i}^m(D|p_m) = \overbrace{p_m \left(c + \tau + \delta_{t_i}[(1-\beta_m)W_{t_i}^m(D|0) + \beta_m V_{t_i}]\right)}^{\text{matched with cooperator}} + \overbrace{(1-p_m)\left(d + \delta_{t_i}[(1-\beta_m)W_{t_i}^m(D|0) + \beta_m V_{t_i}]\right)}^{\text{matched with non-cooperator}} \quad (2.5)$$

where $W_{t_i}^m(D|0)$ reflects the fact that because player i plays D, his match will play D for the duration of the match, regardless of the match's strategy.

After the first period of a match, players know what their opponent will play each period for the duration of the match, except if there is a deviation, which we will consider later. Thus we can formulate the expressions explicitly for the value players expect to receive in the second and subsequent periods of a match, $W_{t_i}^m(\sigma^{gt}|1)$, $W_{t_i}^m(\sigma^{gt}|0)$, and $W_{t_i}^m(D|0)$, as follows

[6]Recall that in the version of the model presented in this section, there is no endogenous breakup. The analogous equations when endogenous breakup is allowed is given by equation (2.39).

$$W_{t_i}^m(\sigma^{gt}|1) = c + \delta_{t_i}\left[(1-\beta_m)W_{t_i}^m(\sigma^{gt}|1) + \beta_m V_{t_i}\right] \;\Rightarrow\; W_{t_i}^m(\sigma^{gt}|1) = \frac{c + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.6}$$

$$W_{t_i}^m(\sigma^{gt}|0) = d + \delta_{t_i}\left[(1-\beta_m)W_{t_i}^m(\sigma^{gt}|0) + \beta_m V_{t_i}\right] \;\Rightarrow\; W_{t_i}^m(\sigma^{gt}|0) = \frac{d + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.7}$$

$$W_{t_i}^m(D|0) = d + \delta_{t_i}\left[(1-\beta_m)W_{t_i}^m(D|0) + \beta_m V_{t_i}\right] \;\Rightarrow\; W_{t_i}^m(D|0) = \frac{d + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.8}$$

Substituting (2.6) and (2.7) into (2.4), we can solve for $W_{t_i}^m(\sigma^{gt}|p_m)$ in terms of parameters and V_{t_i}. Similarly, we can substitute (2.8) into (2.5) to solve for $W_{t_i}^m(D|p_m)$. Doing so yields

$$W_{t_i}^m(\sigma^{gt}|p_m) = p_m\left(\frac{c}{1-\delta_{t_i}(1-\beta_m)}\right) + (1-p_m)\left(d - \lambda + \frac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right) + \left(\frac{\delta_{t_i}\beta_m}{1-\delta_{t_i}(1-\beta_m)}\right)V_{t_i} \tag{2.9}$$

$$W_{t_i}^m(D|p_m) = p_m\left(c + \tau + \frac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right) + (1-p_m)\left(\frac{d}{1-\delta_{t_i}(1-\beta_m)}\right) + \left(\frac{\delta_{t_i}\beta_m}{1-\delta_{t_i}(1-\beta_m)}\right)V_{t_i} \tag{2.10}$$

The first two terms of each reflect what a player expects to receive from a new match given his beliefs about the probability his match will cooperate, while the last term is the value expected following breakup.

As long as $d > 0$, players can always obtain positive value by playing D. Thus $V_{t_i}^m(p_m) > 0$ and all unmatched players will choose to enter the market rather than to stay unmatched. If we assume that a player of a given type will always choose to follow the same strategy, which we will verify shorty as individually rational equilibrium behavior, then the decision problem of the unmatched player given in (2.2) becomes $V_{t_i} = V^m_{t_i}(p_m) > 0$. Equation (2.3) can then be solved for $V^m_{t_i}$, yielding the ex ante expected value of an unmatched player

$$V^m_{t_i}(p_m) = \frac{\mu_m W^m_{t_i}(p_m)}{1-\delta_{t_i}(1-\mu_m)} \tag{2.11}$$

Using (2.11), we can solve (2.9) for $W^m_{t_i}(\sigma^{gt}|p_m)$ and (2.10) for $W^m_{t_i}(D|p_m)$ in terms of parameters only, yielding

$$\begin{aligned} W^m_{t_i}(\sigma^{gt}|p_m) = &\left(\tfrac{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))}{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))-\delta_{t_i}\mu_m\beta_m}\right) \\ &\cdot \left(p_m \frac{c}{1-\delta_{t_i}(1-\beta_m)} + (1-p_m)\left(d - \lambda + \frac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right)\right) \end{aligned} \tag{2.12}$$

$$\begin{aligned} W^m_{t_i}(D|p_m) = &\left(\tfrac{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))}{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))-\delta_{t_i}\mu_m\beta_m}\right) \\ &\cdot \left(p_m(c+\tau) + (1-p_m)d + \frac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right) \end{aligned} \tag{2.13}$$

These expressions provide the lifetime value a player expects to receive immediately after receiving a new match but before playing the stage game for the first time with this new match, and are in terms of parameters only. The first term is a multiplier reflecting the agent's discount rate as well as the probabilities of matching and breakup, while the second terms reflect what a player expects from a match as in (2.9) and (2.10). Together, the terms involving δ_{t_i}, μ_m, and β_m are analogous to the familiar $\frac{1}{1-\delta}$ and $\frac{\delta}{1-\delta}$ terms common in repeated games without search.[7] With these expressions, we are now ready to consider equilibrium in this environment.

[7] A standard repeated game without search is nested in the present model if matching is certain and breakup never occurs. With $\mu_m = 1$ and $\beta_m = 0$, (2.12) and (2.13) reduce to $p_m \frac{c}{1-\delta_{t_i}} + (1-p_m)\left(d - \lambda + \frac{\delta_{t_i} d}{1-\delta_{t_i}}\right)$ and $p_m(c+\tau) + (1-p_m)d + \frac{\delta_{t_i} d}{1-\delta_{t_i}}$.

2.3.2 Equilibrium

In Section 2.3.1 we found the value player i type t_i expects to receive from following either the cooperative grim trigger strategy or from not cooperating and always playing D. Individually rational players will choose the strategy that maximizes their individual expected utility given beliefs about what the match will do. Accordingly, the value expected from a new match by a player with beliefs p_m is

$$W_{t_i}^m(p_m) = \max\{W_{t_i}^m(\sigma^{gt}|p_m), W_{t_i}^m(D|p_m)\} \tag{2.14}$$

A player with beliefs p_m will find it optimal to follow σ_{gt} and cooperate in the first period of a new match if and only if $W_{t_i}^m(\sigma^{gt}|p_m) \geq W_{t_i}^m(D|p_m)$. Equations (2.12) and (2.13), respectively, give these expected values of a new match in terms of parameters, allowing us to solve for a condition on the discount rate necessary for cooperation to be optimal. A newly matched player will choose the optimal strategy as follows:

$$\arg\max W_{t_i}^m(\cdot|p_m) = \begin{cases} \sigma^{gt} & \text{if } W_{t_i}^m(\sigma^{gt}|p_m) \geq W_{t_i}^m(D|p_m) \iff \delta_{t_i} \geq \underline{\delta}_{p_m}^m \\ D & \text{if } W_{t_i}^m(D|p_m) > W_{t_i}^m(\sigma^{gt}|p_m) \iff \delta_{t_i} < \underline{\delta}_{p_m}^m \end{cases} \tag{2.15}$$

$$\text{where} \quad \underline{\delta}_{p_m}^m \equiv \frac{p_m\tau + (1-p_m)\lambda}{(1-\beta_m)(p_m(c-d+\tau) + (1-p_m)\lambda)} \tag{2.16}$$

An increase in the potential loss from cooperating against a defector and the temptation to defect yourself, weighted by the probability of being faced with either situation, increases the discount rate required for cooperation to be optimal. The remaining term in the denominator of (2.16) reflects the benefit received each period from a cooperative relative to a non-cooperative match. The higher is $(c-d)$, the lower the discount rate required to induce cooperation. If $(c-d)$ is sufficiently low or $(p_m\tau + (1-p_m)\lambda)$ is sufficiently high, $\underline{\delta}_{p_m}^m \geq 1$, reflecting the fact that players will never cooperate if cooperation is not sufficiently attractive.

We are now ready to define equilibrium, defining equilibrium with one market first as a special case of equilibrium with two markets, and then providing the complete definition of equilibrium with two markets.

Definition 1 (Equilibrium with One Market)**.** *An equilibrium with one market is an equilibrium with two markets, as given by Definition 2 with conditions i through iv, in the special case of $\mu_s = 0$ so that $V_{t_i}^s(p_s) = 0$ according to (2.11), making players indifferent between entering market $m = s$ and staying unmatched as part of condition ii.*

Definition 2 (Equilibrium with Two Markets)**.** *An equilibrium with two markets is, for $m \in \{s, f\}$, a list of values, $(V_{t_i}, V_{t_i}^m, W_{t_i}^m)$, and beliefs, p_m, such that, given market probabilities of matching and breakup, (μ_m, β_m)*

i. (Matched) Given beliefs p_m, player i chooses strategy σ^{gt} or D according to (2.15) *and* (2.16) *that maximizes $W_{t_i}^m(p_m)$ as given by* (2.14)*, $\forall i$*

ii. (Unmatched) Given values $W_{t_i}^m(p_m)$, $m \in \{s, f\}$, expected once matched, player i forms expectations $V_{t_i}^m(p_m)$ about the value of entering each market according to (2.11) *and chooses to enter market $m \in \{s, f\}$ or to stay unmatched to maximize V_{t_i} according to* (2.2)*, $\forall i$*

iii. (Individual Rationality) Given $V_{t_i}^m(p_m)$ and $W_{t_i}^m(p_m)$, $m \in \{s, f\}$, player i does not have an incentive to deviate from his strategy determined by conditions i and ii, $\forall i$

iv. (Consistency) $\forall i$, player i's beliefs, p_m, are consistent with the strategies followed by all players $j \neq i$

2.4 One Market

In this section we will consider equilibrium with one market, as given by Definition 1. By setting $\mu_s = 0$, we are effectively shutting down the second market and considering only what happens when all players interact within one market, the fast market ($m = f$). After examining what is possible with one market, in the next section we will explore what further opportunities may be available when there are two markets. With the slower market shut down, all one-market equilibria occur in the fast market.

Proposition 1 (Existence with One Market)**.** *An Equilibrium with One Market always exists.*

One possible equilibrium that always exists is repetition of the stage game Nash equilibrium, which in this search framework is a pooling equilibrium in which all players defect. Given expectations that all other players will defect, the probability a match cooperates is $p_f = 0$.[8] Given $p_f = 0$, (2.16) yields $\underline{\delta}_0^f = \frac{1}{1-\beta_f} > 1, \forall \beta_f \in (0, 1)$,

[8]In this one market setting, the equations that have m sub and superscripts in Section 2.3 will have f sub and superscripts in this section because all equilibria occur in the fast market, $m = f$.

so by (2.15), each player maximizes (2.14) by playing D. Given $W^f_{t_i}(0) = W^f_{t_i}(D|0)$, $V^f_{t_i}(0) > 0$ is each player's solution to (2.2). Given $V^f_{t_i}(0)$ and $W^f_{t_i}(0)$, no player has an incentive to deviate from the strategies determined by conditions *i* and *ii*, satisfying condition *iii*. Since all players are playing D, $p_f = 0$, verifying condition *iv*. Thus an equilibrium with one market, as defined in Definition 1, always exists.

Proposition 2 (First Best)**.** *The first-best equilibrium in which all players cooperate exists if and only if* $\delta_\ell \geq \underline{\delta}^f_1$*, where* $\underline{\delta}^f_1$ *is given by* (2.16)*.*

Repetition of the stage game Nash equilibrium, while always a possible equilibrium, is inefficient. The best possible outcome is an equilibrium in which all players cooperate. Pooling on cooperation is an equilibrium if and only if the least-patient player is sufficiently patient. If all players expect all other players to cooperate, then $p_f = 1$. Given beliefs $p_f = 1$, if $\min\{\delta_h, \delta_\ell\} = \delta_\ell \geq \underline{\delta}^f_1$, where $\underline{\delta}^f_1 \equiv \frac{\tau}{(1-\beta_f)(c-d+\tau)}$ by (2.16), then each player, choosing the optimal strategy according to (2.15), maximizes (2.14) by following the grim trigger strategy. Thus $\forall i, t_i \in \{h, \ell\}$, conditions *i* and *ii* yield $W^f_{t_i}(1) = W^f_{t_i}(\sigma^{gt}|1)$ and $V^f_{t_i}(1) > 0$. If a player found it optimal to cooperate in the first period of a match, he will not find it optimal to deviate later on in the match,[9] satisfying condition *iii*. Since all players are cooperating, $p_f = 1$, and the consistency condition *iv* is satisfied. Thus an equilibrium with one market in which all players cooperate exists if $\delta_\ell \geq \underline{\delta}^f_1$, and the first-best outcome is achievable.

The first-best outcome is feasible under a fairly wide range of parameters, including at relatively high breakup probabilities. The condition $\delta_\ell \geq \underline{\delta}^f_1$ can be solved for the maximum breakup probability for which cooperation can be sustained for given discount rates, yielding

$$\beta_f \leq 1 - \frac{\frac{\tau}{c-d}}{\delta_\ell \left(1 + \frac{\tau}{c-d}\right)} \equiv \bar{\beta}_1\left(\tfrac{\tau}{c-d}\right) \tag{2.17}$$

where $\frac{\tau}{c-d}$ is the ratio of the temptation to defect to the value of sustaining a cooperative relationship. Figure 2.1 shows the relationship between the breakup probability and the temptation to defect for different values of δ_ℓ. If both types are fairly patient, with $\delta_\ell = 0.97$ and $\delta_h = 0.99$, cooperation can be sustained at breakup probabilities: $\bar{\beta}_1(\frac{\tau}{c-d} = 1) = 0.48$, $\bar{\beta}_1(5) = 0.14$, $\bar{\beta}_1(10) = 0.06$, $\bar{\beta}_1(25) = 0.01$. Surprisingly, players find it optimal to cooperate even if the likelihood of staying matched for another period

[9]Formally, it remains true that $\delta_{t_i} \geq \underline{\delta}^f_1$, so cooperation remains the best response for all players.

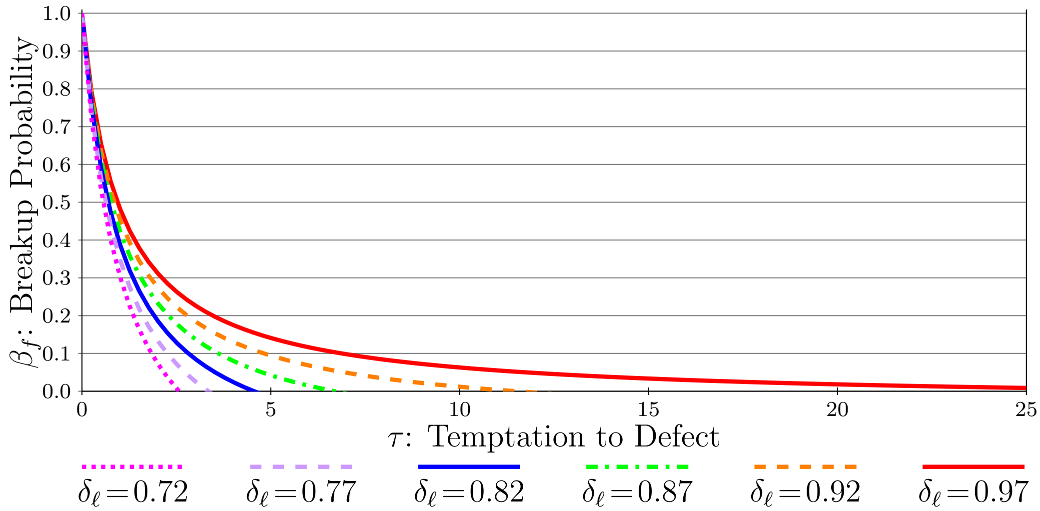

Figure 2.1: Maximum breakup probability (β_f) for existence of first-best equilibrium vs. temptation to defect (τ) for different values of δ_ℓ with $c - d = 1$. For (τ, β_f) pairs below and left of a line, the first-best equilibrium exists for that value of δ_ℓ.

is close to a flip of a coin if the temptation to defect and the value of a cooperative relationship are equal. Even more surprisingly, cooperation can be sustained if the temptation to defect is 25 times higher than the value of a cooperative relationship with a one percent probability of breakup, which is a small probability, but far from requiring infinite repetition.

If $\delta_\ell < \underline{\delta}^f_1$, then the first-best outcome is not achievable. However, if the high type are sufficiently patient, an outcome that provides higher ex ante expected value for all players than repetition of the stage game Nash equilibrium might still be possible.

Proposition 3 (Existence of Separating Equilibrium with One Market)**.** *Given* δ_ℓ *and* δ_h*, if there exists a* π *such that* $\delta_\ell < \underline{\delta}^f_\pi \leq \delta_h$*, where* $\underline{\delta}^f_\pi$ *is defined by* (2.16)*, then there exists an Equilibrium with One Market in which all high types cooperate and all low types defect.*

Consider an equilibrium in which all players who are the high type follow the grim trigger strategy and all players who are the low type defect. In such an equilibrium, the probability of being matched with a player who cooperates is the same as the fraction of the population who are the high type, so $p_f = \pi$.[10] With all players choosing which strategy to follow according to (2.15), a necessary condition for a separating equilibrium with one market to exists is $\delta_\ell < \underline{\delta}^f_\pi \leq \delta_h$ where $\underline{\delta}^f_\pi$ is defined by (2.16).

For both pooling equilibria, the conditions on the discount rates were also sufficient for existence, which remains the case here. If $\delta_\ell < \underline{\delta}^f_\pi$, then $W^f_\ell(\pi) = W^f_\ell(D|\pi)$ is the solution to (2.14) for all low types, and thus $V^f_\ell(\pi) > 0$. If $\delta_h \geq \underline{\delta}^f_\pi$, then $W^f_h(\pi) = W^f_h(\sigma^{gt}|\pi)$ is the solution to (2.14) for all high types types. When a high type is matched with a low type, if $\lambda > 1$, it is possible that the value realized from the match will be negative. However, $\delta_h \geq \underline{\delta}^f_\pi \Leftrightarrow W^f_h(\sigma^{gt}|\pi) \geq W^f_h(D|\pi) > 0$, and thus $V^f_h(\pi) > 0$ is the solution to (2.2) for all high types.

We next need to verify that no player has a profitable deviation from the strategies determined by conditions *i* and *ii*. Any player with which a low type is matched will play D in the second and subsequent periods of a match, regardless of the match's type, so deviation from D is not profitable for the low type. If a high type is matched with a low type, his strategy is to play D for the duration of the match, and it is not profitable for him to deviate from the grim trigger. If a high type is matched with another high type, his belief that the match will cooperate in the second and subsequent periods of a match is $p_f = 1$. Deviation is not profitable as long as $\delta_h \geq \underline{\delta}^f_1$. It is easy to verify that $\underline{\delta}_1 < \underline{\delta}^f_{p_f}, \forall p_f < 1$, and thus the condition on the discount rate required for deviation to not be profitable is nested in the condition on the discount rate required for cooperation to be optimal in the first place. Intuitively, if a player finds it optimal to cooperate when he is uncertain about whether the other player will also cooperate, he certainly finds it optimal to cooperate when he is certain. Thus no player has a profitable deviation and condition *iii* is satisfied.

Given conditions *i*, *ii*, and *iii*, all low types will play D and all high types will play σ^{gt}. Thus the probability of being matched with a player who cooperates in

[10] Intuitively, this is the case because with only exogenous breakup, players of each type are equally likely to be unmatched. Formally, the probability that an unmatched player is the high type must be found from the stead-state distribution of players of each type and in each state, matched and unmatched. When no matches are ended endogenously (as occurs in this section, because endogenous breakup is not allowed, but which also occurs if the matching probability is sufficiently low even if endogenous breakup is allowed), the probability that an unmatched player is the high type reduces to π. However, when endogenous breakup is exercised (which occurs when endogenous breakup is allowed and the matching probability is higher), the probability no longer reduces to π, but rather depends on μ_f and β_f as well. This is discussed formally in Section 2.6.

equilibrium is the fraction of the population who are the high type, so $p_f = \pi$, and the consistency condition *iv* is satisfied. Thus $\delta_\ell < \underline{\delta}^f_\pi \leq \delta_h$ is necessary and sufficient for a separating equilibrium with one market to exist.

Proposition 4 (Welfare Gains from Separating Equilibrium with One Market)**.** *When the separating equilibrium with one market exists, it Pareto dominates repetition of the stage game Nash equilibrium. These welfare gains exist no matter how impatient are the low type, existing for any δ_ℓ such that $0 \leq \delta_\ell < \underline{\delta}^f_\pi$.*

When a separating equilibrium with one market exists, all players prefer it to repetition of the stage game Nash equilibrium because it provides higher ex ante expected utility for players of each type. For each type, $W_{t_i}(D|p_f) > W_{t_i}(D|0), \forall p_f > 0$. Thus the low type prefer the separating equilibrium because they receive $(c+\tau) > d$ in the first period of any match with a high type. The high type risk losing λ if matched with a low type in a separating equilibrium, which they could avoid by play D. However, a separating equilibrium only exists if $W^f_h(\sigma^{gt}|\pi) \geq W^f_h(D|\pi)$, and since $W_h(D|p_f) > W_h(D|0), \forall p_f > 0$, it must be the case that $W^f_h(\sigma^{gt}|\pi) > W_h(D|0)$, the value expected from matches in an uncooperative equilibrium, so the high type also expect higher utility. Thus when a separating equilibrium with one market exists, it Pareto dominates repetition of the stage game Nash equilibrium.

For the separating equilibrium with one market to exist, the high type must be sufficient patient, requiring $\underline{\delta}^f_\pi \leq \delta_h$. The low type, however, can be very impatient, $0 \leq \delta_\ell < \underline{\delta}^f_\pi$. Thus the welfare gains achievable with the separating equilibrium with one market exist even if the low type are completely myopic ($\delta_\ell = 0$). Figure 2.2 shows the relationship between the fraction of the population who are the high type, π, and the temptation to defect, τ, for different values of δ_h. Each line shows the minimum value of π for which the one-market separating equilibrium exists for that level of δ_h, as a function of τ. No matter how impatient are the low type, cooperation can still be sustained among the high types as long as enough of the population are the high type, providing higher expected utility for all.

2.5 Two Markets

In Section 2.4 we saw how when the first-best outcome is not obtainable, sorting may provide a Pareto improvement opportunity over repetition of the stage game Nash equilibrium. Low types expect higher utility because they receive the temptation payment whenever they are matched with a high type. The other side of

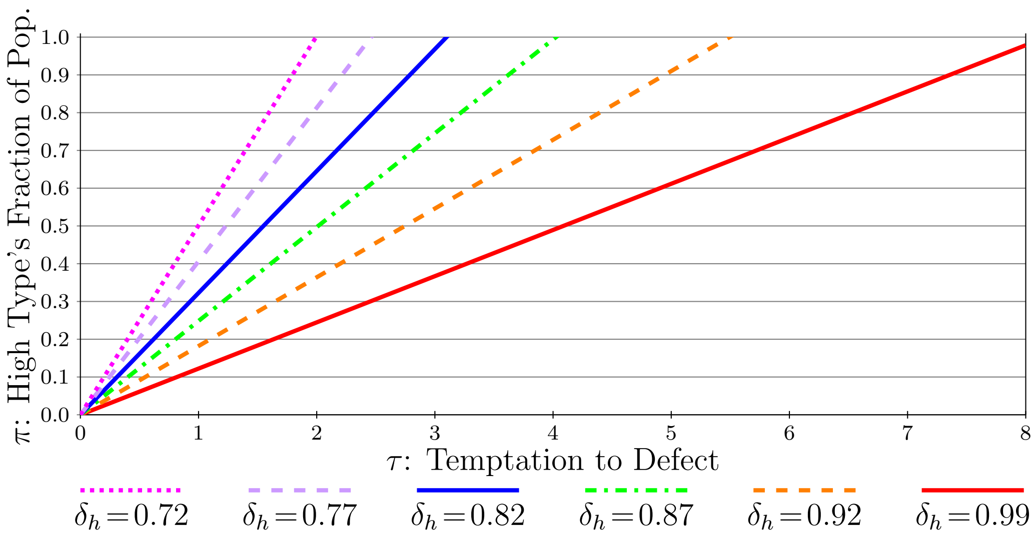

Figure 2.2: Minimum fraction of the population who are the high type (π) for existence of one-market separating equilibrium vs. temptation to defect (τ) for different values of δ_h with $c - d = 1$ and $\beta = 0.1$. For (τ, π) pairs above and left of a line, the one-market separating equilibrium exists for that value of δ_h.

the low types' gain is of course losses experienced by high types. However, in spite of these losses expected from being matched with low types, the high types also expect higher utility when the expected gains from cooperation from being matched with other high types exceed the expected losses.

Proposition 4 naturally leads to the question, can overall welfare be improved by providing opportunities for further sorting? Such gains would likely come from reducing the losses expected by high types, but in order for this to be a Pareto improvement, it must not come at the expense of the low types. In this section we will see how in some circumstances the addition of a second market can make the high types better off without decreasing the ex ante expected value of the low types.

2.5.1 Two Market Setting

The environment presented in Section 2.3 contained two markets, $m \in \{s, f\}$. However, in Section 2.4, the slow market was shut down by assuming $\mu_s = 0$ so that if a player entered the slow market, he would never receive a match and would thus be indifferent between staying unmatched and entering the slow market. We now want to consider what is possible if $\mu_s > 0$ so that the slow market is a viable option for players.

Just as was done in the one-market setting, we will assume that all players of the same type will always choose to follow the same stage game strategy once matched and to enter not only a market, but the same market each time they are unmatched, all of which will be verified shortly as being individually rational equilibrium behavior, as defined by Definition 2.

In the one-market setting, we considered the existence of three equilibria. In the two-market setting, continuing to limit our focus to symmetric equilibria, conditional on type, and two possible stage game strategies, the number of possible equilibria has increased to at least seven. There are four equilibria with full pooling in which all players enter the same market and follow the same stage game strategy, two equilibria with partial pooling in which all players enter the same market but follow separate stage game strategies once matched, and one equilibrium with full separation.

The equilibria in which all players pool into the fast market are the same as the equilibria discussed in Section 2.4. Each has an analogous equilibrium with pooling into the slow market. However, for any two equilibria with identical strategies followed once matched, the value expected from the equilibrium with pooling into the fast market will always be higher than that expected from pooling into the slow market because $\mu_f > \mu_s$, and thus no player will ever prefer an equilibrium with pooling into the slow market to the analogous equilibrium with pooling into the fast market. Accordingly, we will not consider further any equilibria with pooling into the slow market.

2.5.2 Separating Equilibrium with Two Markets

The new case of interest is the equilibrium with full separation made possible by the existence of two markets. Given the large number of parameters and the resulting complexity of the conditions that must be satisfied, it will be useful to focus attention on cases in which the two markets are as comparable as possible. Thus for simplicity, we will assume that the breakup probability is identical in each market ($\beta_s = \beta_f \equiv \beta$) and that the temptation to defect is equal to the loss from cooperating but being defected against ($\tau = \lambda$). Furthermore, since our interest is in comparing overall welfare with the one-market setting, any conclusions we reach about possible welfare improvements achievable with the addition of the slow market will be strongest if the fast market is made as attractive as possible. In the one-market setting, the matching probability only affected the level of the ex ante expected value of unmatched players, with higher matching probabilities translating into higher expected values. The incentives facing the players, and thus the qualitative results, are identical $\forall \mu_f \in (0, 1]$. Thus it is without loss of generality to consider the case where $\mu_f = 1$, which we will do from now on in this Section in order for the comparison to be as favorable as possible for the one-market setting.

Proposition 5 (Separating Equilibrium with Two Markets)**.** *A Separating Equilibrium with Two Markets is an Equilibrium with Two Markets, as given by Definition 2, in which all unmatched high types enter the slow market and cooperate following the grim trigger strategy once matched while all unmatched low types enter the fast market and defect once matched. Given δ_ℓ and δ_h, if $\delta_h \geq \underline{\delta}_1^s$, where $\underline{\delta}_1^s$ is given* (2.16)*, and if there exist bounds on the matching probability in the slow market, $\underline{\mu}_s < \mu_s < \overline{\mu}_s$, and on the temptation to defect, $\tau < \overline{\tau}$, such that $V_h^s \geq V_h^f$ and $V_\ell^f \geq V_\ell^s$, then a Separating Equilibrium with Two Markets exists.*

Proposition 6 (Independence of π in Separating Equilibrium with Two Markets)**.** *Existence of and the value expected by players of each type in the separating equilibrium with two markets does not depend on the fraction of the population who are the high type, π.*

First we will show Proposition 5, existence of the two-market separating equilibrium. In doing so, the fraction of the population who are the high type, π, will never appear, thus proving Proposition 6.

If all players believe that all players matched in the fast market will defect, then $p_f = 0$. Given these beliefs, from (2.16), $\underline{\delta}_0^f = \frac{1}{1-\beta} > 1 > \delta_{t_i}, \forall t_i, \forall \beta \in (0, 1)$, and thus by (2.15), all players will play D if matched in the fast market. By (2.14), the value expected by matched players of each type in the fast market is $W_{t_i}^f(D|0)$.

If all players believe that all players matched in the slow market will follow the grim trigger strategy, then $p_s = 1$. Given these beliefs, the cutoff on the discount rate required for cooperation, as given by (2.16), is $\underline{\delta}_1^s = \frac{\tau}{(1-\beta)(c-d+\tau)}$. A necessary condition for a separating equilibrium with two markets to exist is $\delta_h \geq \underline{\delta}_1^s$. If the high type are sufficiently patient that this condition holds, then by (2.15), all high types will follow the grim trigger strategy once matched in the slow market and the value expected once matched is $W_h^s(\sigma^{gt}|1)$. If a player who is the low type enters the slow market, he will choose whichever strategy maximizes his expected value according to (2.15) with $\underline{\delta}_1^s$ and the value expected once matched, $W_\ell^s(1)$, is given by (2.14).

Given the value expected once matched for players of each type in each market as required by condition *i* of Definition 2, the value expected by unmatched players from entering either market can be specified according to condition *ii*. For $\mu_f = 1$, matching occurs with certainty in the fast market, and thus the value expected by an unmatched player choosing to enter the fast market is equal to the value expected by a matched player in the fast market. Combining (2.11) and (2.13) with the case of $\mu_f = 1$ yields

$$V_{t_i}^f(0) = \frac{\mu_f W_{t_i}^f(D|0)}{1-\delta_{t_i}(1-\mu_f)} = W_{t_i}^f(D|0) = \frac{\mu_f d}{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f)) - \delta_{t_i}\mu_f\beta_f}$$
$$= \frac{d}{(1-\delta_{t_i}(1-\beta_f)) - \delta_{t_i}\beta_f} = \frac{d}{1-\delta_{t_i}} \tag{2.18}$$

Since matching is certain in the fast market, players entering the fast market receive d each period, and thus the value expected is the same as the value in a standard repeated game without search probabilities.

The value expected by unmatched players for the slow market where $\mu_s < 1$ does not simplify to the same extent. Using (2.11), (2.12), and (2.13), we can derive expressions for the value expected by unmatched players entering the slow market. The value expected by an unmatched high type entering the slow market is

$$V_h^s(1) = \frac{\mu_s W_h^s(\sigma^{gt}|1)}{1-\delta_h(1-\mu_s)} = \frac{\mu_s c}{(1-\delta_h(1-\mu_s))(1-\delta_h(1-\beta)) - \delta_h\mu_s\beta} \tag{2.19}$$

For the low type, there are two cases for the value expected by an unmatched player if he were to enter the slow market, one if he is sufficiently patient that he would cooperate[11] and one if he is not:

$$V_\ell^s(1) = \begin{cases} \dfrac{\mu_s W_\ell^s(\sigma^{gt}|1)}{1-\delta_\ell(1-\mu_s)} = \dfrac{\mu_s c}{(1-\delta_\ell(1-\mu_s))(1-\delta_\ell(1-\beta))-\delta_\ell\mu_s\beta} & \text{if } \delta_\ell \geq \underline{\delta}_1^s \\ \dfrac{\mu_s W_\ell^s(D|1)}{1-\delta_\ell(1-\mu_s)} = \dfrac{\mu_s\left([1-\delta_\ell(1-\beta)](c+\tau)+\delta_\ell(1-\beta)d\right)}{(1-\delta_\ell(1-\mu_s))(1-\delta_\ell(1-\beta))-\delta_\ell\mu_s\beta} & \text{if } \delta_\ell < \underline{\delta}_1^s \end{cases} \tag{2.20}$$

This completes the specification of the values expected by unmatched players. Given these values, each unmatched player chooses to enter the market that provides the highest expected value according to (2.2). For a separating equilibrium with two markets to exist, it is necessary that $V_h^s(1) \geq V_h^f(0)$ and $V_\ell^f(0) \geq V_\ell^s(1)$ as given by (2.18), (2.19), and (2.20). These conditions and the condition on the discount rate of the high type, $\delta_h \geq \underline{\delta}_1^s$, complete the specification required for conditions *i* and *ii* of Definition 2 for an equilibrium with two markets.

To see that no player has incentive to deviate from their strategy as specified by conditions *i* and *ii*, it is useful to consider bounds on parameters implied by these conditions. The condition for the high type on the discount rate and values expected from each market can be solved for an upper bound on the temptation to defect and a lower bound on the matching probability in the slow market, such that

$$\delta_h \geq \underline{\delta}_1^s \equiv \frac{\tau}{(1-\beta)(c-d+\tau)} \Leftrightarrow \tau \leq \overline{\tau}_h \equiv \frac{\delta_h(1-\beta)(c-d)}{1-\delta_h(1-\beta)} \tag{2.21}$$

$$V_h^s(\sigma^{gt}|1) \geq V_h^f(D|0) \Leftrightarrow \mu_s \geq \underline{\mu}_s \equiv \frac{(1-\delta_h)(1-\delta_h(1-\beta))d}{(1-\delta_h(1-\beta))(c-\delta_h d)-\delta_h\beta(c-d)} \tag{2.22}$$

For players who are the high type, if $\tau \leq \overline{\tau}_h$, there is no profitable deviation when matched, and if $\mu \geq \underline{\mu}_s$, the slow market is attractive enough to be worth the wait and there is no profitable deviation when unmatched. Thus high types do not have an incentive to deviate when these conditions hold.

To satisfy condition *iii* we also need to establish that the low type does not have a profitable deviation. There are two cases to consider. Consider first the case

[11] The case when $\delta_\ell \geq \underline{\delta}_1^s$ also implies that the first-best is possible, making the second market unnecessary from a welfare perspective. This is discussed in Section 2.5.3. However, it is still necessary to examine this case here in order to fully characterize when the two-market separating equilibrium exists and to see Proposition 6, that this equilibrium does not depend on π.

when $\delta_\ell \geq \underline{\delta}^s_1$ and the low type is sufficiently patient that they find it optimal to cooperate if matched in the slow market. This implies an alternate upper bound on the temptation parameter, as well as a lower bound on the matching probability in the slow market implied by the condition required for the low type to find it optimal to enter the fast market, such that

$$\delta_\ell \geq \underline{\delta}^s_1 \equiv \frac{\tau}{(1-\beta)(c-d+\tau)} \Leftrightarrow \tau \leq \overline{\tau}_\ell \equiv \frac{\delta_\ell(1-\beta)(c-d)}{1-\delta_\ell(1-\beta)} \tag{2.23}$$

$$V^f_\ell(D|0) \geq V^s_\ell(\sigma^{gt}|1) \Leftrightarrow \mu_s \leq \overline{\mu}_s \equiv \frac{(1-\delta_\ell(1-\beta))d}{c-\delta_\ell(1-\beta)d} \tag{2.24}$$

For players who are the low type, if $\tau \leq \overline{\tau}_\ell$ they would cooperate in the slow market. As long as the slow market is sufficiently slow, with $\mu_s \leq \overline{\mu}_s$, they do not find it worth the wait and entering the slow market is not a profitable deviation. Thus if $\tau \leq \overline{\tau}_\ell = \min\{\overline{\tau}_\ell, \overline{\tau}_h\}$ and $\underline{\mu}_s \leq \mu_s \leq \overline{\mu}_s$, no player has a profitable deviation and condition *iii* is satisfied.

Now consider the other case when the low type is not sufficiently patient that they would cooperate in the slow market, or when $\delta_\ell < \underline{\delta}^s_1$. This case implies a lower bound on the temptation to defect, or

$$\delta_\ell < \underline{\delta}^s_1 \equiv \frac{\tau}{(1-\beta)(c-d+\tau)} \Leftrightarrow \tau > \underline{\tau}_\ell \equiv \frac{\delta_\ell(1-\beta)(c-d)}{1-\delta_\ell(1-\beta)} \tag{2.25}$$

The low type in this case will only find it optimal to enter the fast market if $V^f_\ell(D|0) \geq V^s_\ell(D|1)$. Using (2.18) and (2.20) in the case where $\delta_\ell < \underline{\delta}_1$ yields

$$\frac{d}{1-\delta_\ell} \geq \frac{\mu_s\left([1-\delta_\ell(1-\beta)](c+\tau)+\delta_\ell(1-\beta)d\right)}{(1-\delta_\ell(1-\mu_s))(1-\delta_\ell(1-\beta))-\delta_\ell\mu_s\beta} \tag{2.26}$$

Since $\mu_f = 1$, the low type can get at least d without wait by entering the fast market. Since he will also play D if matched in the slow market, he will get d in the second and subsequent periods of any match in the slow market. Thus he only finds entering the slow market to be a profitable deviation if $\mu_s(c+\tau) > d$. The underlying intuition for this condition is transparent in the case presently being considered with $\mu_f = 1$ and $\beta_f = \beta_s \equiv \beta$, but it comes directly from simplifying (2.26) and similar conditions with corresponding intuition can be obtained in the more general case. This joint condition on the temptation to enter the slow market implicitly defines bounds on τ and μ_s as follows

$$\tau < \overline{\tau} \equiv \frac{d}{\mu_s} - c \qquad \text{and} \qquad \mu_s < \overline{\mu}_s(\overline{\tau}) \equiv \frac{d}{c + \tau'}, \forall \tau' \leq \overline{\tau} \tag{2.27}$$

Thus when $\delta_\ell < \underline{\delta}_1^s$ and the low type would play D in the slow market, if $\tau < \overline{\tau}$ the temptation payment received once matched in the slow market is not high enough to make it worth the wait for any $\mu_s < \overline{\mu}_s(\overline{\tau})$, and the deviation of entering the slow market is not profitable. Thus in this case if $\underline{\tau}_\ell < \tau \leq \min\{\overline{\tau}, \overline{\tau}_h\}$ and $\underline{\mu}_s \leq \mu_s < \overline{\mu}_s(\overline{\tau})$ no player has a profitable deviation and condition *iii* is satisfied.

Conditions *i*, *ii*, and *iii* of Definition 2 for an equilibrium with two markets to exist are thus satisfied. With both types following strategies consistent with the separating equilibrium with two markets, all players cooperate in the slow market and defect in the fast market, $p_s = 1$ and $p_f = 0$, and the consistency condition *iv* is satisfied. Thus a separating equilibrium with two markets that satisfies all the conditions of an equilibrium with two markets as in Definition 2 exists as claimed in Proposition 5.

Proposition 6 is that existence of and the value expected by players of each type in the separating equilibrium with two markets does not depend on the fraction of the population who are the high type, π. This follows immediately from the fact that π never appears in any expressions in the preceding proof of existence of the separating equilibrium with two markets.

Figure 2.3 shows the range of matching probabilities in the slow market for which the two-market separating equilibrium exists in the case where the low type is impatient enough that they would defect if they entered the slow market ($\delta_\ell < \underline{\delta}_1^s$). The solid line is $\overline{\mu}_s(\overline{\tau})$ given by (2.27), the upper bound on μ_s above which the low type will be sufficiently tempted that they will enter the slow market.[12] The horizontal dashed line is $\underline{\mu}_s$, as given by (2.22), for $\delta_h = 0.99$. When μ_s is above this horizontal dashed line, the high type finds it optimal to enter the slow market. The area between these two lines (shaded with northwest to southeast lines, both to the left and right of the vertical dashed line) are the matching probabilities in the slow market that, given τ, are in the range $\underline{\mu}_s \leq \mu_s < \overline{\mu}_s(\overline{\tau})$ for which the high type want to enter the slow market and the low type do not, which is the range of μ_s for

[12] Analogously, for a given μ_s, values of τ to the right of the solid line are sufficiently tempting for the low type that they will enter the slow market.

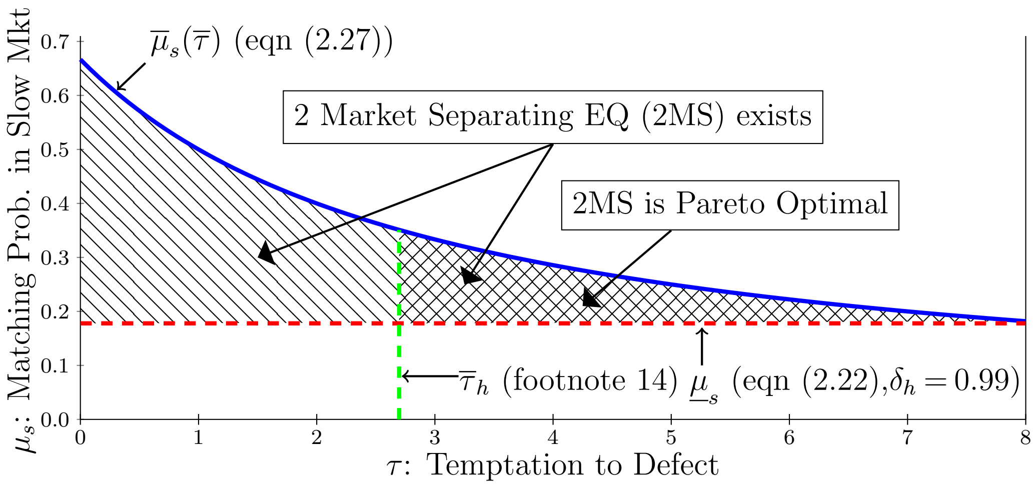

Figure 2.3: Range of matching probabilities in slow market (μ_s) for existence of two-market separating equilibrium vs. temptation to defect (τ) for $\delta_h = 0.99$ when the low type is impatient enough that they would defect if they entered slow market (with $\pi = 0.33$, $\beta = 0.1$, and $c - d = 1$). The region shaded with northwest to southeast lines are the (μ_s, τ) pairs for which the two-market separating equilibrium exists, while the subset crosshatched with both northwest to southeast and northeast to southwest lines are the (μ_s, τ) pairs that make up the Pareto improvement region of Proposition 7.

which the two-market separating equilibrium exists. The subset of this area to the right of the vertical dashed line (crosshatched with both northwest to southeast and northeast to southwest lines) is the subset of these (μ_s, τ) pairs for which the two-market separating equilibrium provides a Pareto improvement opportunity, which is discussed further in the next section.

2.5.3 Welfare Gains from Second Market

We have just seen that the two-market separating equilibrium can exist both when the low type would cooperate if matched in the slow market and when they would not. However, the case when the low type is sufficiently patient that they would cooperate in the slow market is the case when $\delta_\ell \geq \underline{\delta}_1^s$, and since the breakup probability is the same in each market, $\underline{\delta}_1^s = \underline{\delta}_1^f \equiv \underline{\delta}_1$, so a fully cooperative equilibrium with pooling in the fast market also exists.[13] Thus in this case the first-best outcome is feasible so there cannot be welfare gains from the addition of the second market. The sorting possibilities provided by the second market will only provide a Pareto improvement opportunity in the case when the first-best is not feasible.

Proposition 7 (Welfare with Two Markets)**.** *When the discount rates of the types are such that $\delta_\ell < \underline{\delta}_1 \leq \delta_h < \underline{\delta}_\pi$ and there exist bounds $\underline{\tau} < \tau < \overline{\tau}$, and $\underline{\mu}_s \leq \mu_s < \overline{\mu}_s$ such that $V_h^s \geq V_h^f$ and $V_\ell^f \geq V_\ell^s$, then*

7a. The first-best outcome in which all players cooperate is not feasible

7b. A Separating Equilibrium with One Market does not exist

7c. A Separating Equilibrium with Two Markets exists

7d. The Separating Equilibrium with Two Markets Pareto dominates all other equilibria feasible with this information structure and with all players following the grim trigger strategy σ^{gt} or always playing D

The proof of Proposition 7 follows directly from the preceding analysis. From Section 2.4 we know that the low type discount rate condition, $\delta_\ell < \underline{\delta}_1$, directly implies *7a*. The condition on the discount rate of the high type, $\delta_h < \underline{\delta}_\pi$, directly implies *7b* as was shown with Proposition 3. *7c* was shown in the previous section in the proof of Proposition 5, included in the case when $\delta_\ell < \underline{\delta}_1 \leq \delta_h$. In this case we found that the separating equilibrium with two markets exists if $\underline{\tau}_\ell < \tau \leq \min\{\overline{\tau}, \overline{\tau}_h\}$ and $\underline{\mu}_s \leq \mu_s < \overline{\mu}_s(\overline{\tau})$, where these bounds are given by (2.21), (2.22), (2.25), and (2.27). These bounds were defined such that *7d* is true. The only other equilibria feasible with the current information structure and strategies are repetition of the stage game Nash equilibrium with pooling into either the slow or fast market, with pooling into fast clearly preferred by all. Given *7c*, we know that the low type are indifferent between the two-market separating equilibrium and repetition of the stage game Nash equilibrium with pooling into the fast market. Since high types have the option to enter the fast market and play D but instead find it optimal to enter the slow market and cooperate, they prefer the two-market separating equilibrium. Other equilibria may be possible if we expanded our focus to include more complicated information structures or coordinated punishment. However, continuing our focus on

[13]The case of $\beta_s = \beta_f$ is for ease of comparison across markets but is not required. More generally the first-best is feasible for any $\beta_f \geq \beta_s$ as long as $\beta_f \leq \overline{\beta}_1(\tau)$ as given by (2.17), which corresponds to the condition that $\delta_\ell \geq \underline{\delta}_1^f$.

strategic simplicity, they are not, and the fact that separation can improve welfare despite the simplicity of what is expected of the players is central to the importance of this result.

The shaded region to the right of the vertical dashed line (crosshatched with both northwest to southeast and northeast to southwest lines) in Figure 2.3 corresponds with the (μ_s, τ) pairs that make up the Pareto improvement region of Proposition 7. The bounds on μ_s and the upper bound on τ were discussed in the previous section. The vertical dashed line is the largest τ for which the one-market separating equilibrium exists, co-existing with the two-market separating equilibrium. The welfare implications of this co-existence are discussed in the next section. Proposition 7 deals with the case where the high type are not sufficiently patient for the one-market separating equilibrium to exist. This case, $\delta_h < \underline{\delta}_\pi$, implies a lower bound on τ above which the one-market separating equilibrium does not exist,[14] which leaves the two-market separating equilibrium as the only possible Pareto improvement opportunity over repetition of the stage game Nash equilibrium. Any (μ_s, τ) pair to the right of the vertical dashed line in the shaded region (crosshatched) are values of (μ_s, τ) for which the second market increases welfare. Even when the temptation to defect is 8 times larger than the value of a cooperative relationship and no Pareto improvement is possible in one market, cooperation can still be sustained among the high types by allowing separation of the types, making the high types better off without making the low types worse off.

2.5.4 One versus Two Market Separating Equilibria

Item *7a* of Proposition 7 is not necessary for existence of either separating equilibrium. However, when the first-best is feasible, the maximum welfare is already achievable and other equilibria are of little interest. When the first-best is not feasible, without sorting, the worst case in terms of welfare is all that is possible, and potentials for sorting become of interest as a way of improving welfare. In Proposition 4 we saw how separation by action in the one-market setting increases the welfare of all players. From Proposition 7 we know that when this separation is not possible, a Pareto improvement opportunity still exists over repetition of the stage game Nash equilibrium if full separation is made possible by introduction of a second, slower

[14] The one-market separating equilibrium does not exist for any $\overline{\tau}_h \equiv \tau > \frac{\delta_h(1-\beta_m)\pi(c-d)}{(1-\delta_h(1-\beta_m))}$, which is $\delta_h < \underline{\delta}_\pi$ with $\lambda = \tau$.

market, increasing the welfare of the high type without decreasing the welfare of the low type.

When both the one and two-market separating equilibria exist, the low type will always strictly prefer the separating equilibrium with one market to both the separating equilibrium with two markets and repetition of the stage game Nash equilibrium, between which the low type is indifferent. This occurs because in all three equilibria, the low type is guaranteed at least d each period, while in the separating equilibrium with one market, a payoff of $c + \tau$ is received in the first period of each new match with a high type. The high type will always strictly prefer either separating equilibrium to repetition of the stage game Nash equilibrium. However, the high type does not always prefer the separating equilibrium with one market to the one in two markets as does the low type due to the trade-off facing the high type between potential losses associated with being matched with low types in the one-market separating equilibrium and the wait expected in the two-market separating equilibrium. The high type prefers the two-market separating equilibrium to the one-market separating equilibrium if and only if he finds it worth the wait. Formally, this result is as follows:

Proposition 8 (Comparison of Separating Equilibria). *If $\delta_\ell < \underline{\delta}_1 < \underline{\delta}_\pi \leq \delta_h$ and the conditions of both Propositions 3 and 5 are met such that both a separating equilibrium with one and two markets exist, if given π there exists a $\overline{\mu}_s(\pi)$ such that $\mu_s \leq \overline{\mu}_s(\pi)$, or equivalently, if given μ_s there exists a $\underline{\pi}$ such that $\pi \geq \underline{\pi}$, then the separating equilibrium with one market Pareto dominates all other equilibria feasible with this information structure and with all players following the grim trigger strategy σ^{gt} or always playing D.*

In order for the separating equilibrium with one market to Pareto dominate the separating equilibrium with two markets, it must provide higher ex ante expected value for the high type, or $V_h^f(\sigma^{gt}|\pi) \geq V_h^s(\sigma^{gt}|1)$. When the high type prefer separation in one market to two markets depends on the fraction of the population who are the high type as well as how attractive the second, slower market is. For a given population fraction π, the high type prefer the separating equilibrium with one market to that in two markets if

$$\mu_s \leq \overline{\mu}_s(\pi) \equiv \frac{(1-\delta_h(1-\beta))[\pi c + (1-\pi)((1-\delta_h(1-\beta))(d-\tau) + \delta_h(1-\beta)d)]}{c - \delta_h(1-\beta)[\pi c + (1-\pi)((1-\delta_h(1-\beta))(d-\tau) + \delta_h(1-\beta)d)]} \tag{2.28}$$

Alternatively, for a given matching probability in the slow market, μ_s, the high type prefer separation-by-action within one market rather than full separation across two markets if a large enough fraction of the population is the high type, that is if

$$\pi \geq \underline{\pi} \equiv \frac{\frac{\mu_s c(1-\delta_h)}{(1-\delta_h(1-\beta))(1-\delta_h(1-\mu_s))-\delta_h \mu_s \beta} - [(1-\delta_h(1-\beta))(d-\tau)+\delta_h(1-\beta)d]}{c-[(1-\delta_h(1-\beta))(d-\tau)+\delta_h(1-\beta)d]} \tag{2.29}$$

Since the low type always prefer the separating equilibrium with one market to that in two markets, if the slow market is slow enough, as given by (2.28), or equivalently, if enough of the population is expected to cooperate, as given by (2.29), then the gains for the high type found in the separating equilibrium with two markets that come from avoiding being matched with low types playing D in the separating equilibrium with one market are not worth the wait and the separating equilibrium with one market is Pareto optimal.

Proposition 8 is illustrated in Figure 2.4. The solid line and horizontal dashed lines are the bounds on μ_s and the vertical dashed line is the bound on τ from Figure 2.3. The shaded (crosshatched) region to the right of the vertical dashed line is the (μ_s, τ) pairs for which only the two-market separating equilibrium exists. This is the region for which the second market unambiguously provides a Pareto improvement because no other equilibrium with cooperation is possible, as discussed in the previous section with Proposition 7. The shaded (with northwest to southeast lines and with solid fill) to the left of the vertical dashed line is the (μ_s, τ) pairs for which the one and two-market separating equilibria both exist. When both exist, the low type always prefer the separating equilibrium with one market to that in two. The high type also prefer the one-market separating equilibrium when $\mu_s \leq \overline{\mu}_s(\pi)$, which is the downward-sloping dashed line (equation (2.28)) in Figure 2.4. This subset below the downward-sloping dashed line (the region with solid background) are (μ_s, τ) values for which both types prefer separation within one market to separation across two markets. This is the Pareto improvement region discussed in Proposition 8.

2.6 Endogenous versus Exogenous Breakup

A legitimate question is how applicable are the results presented in the previous sections derived from a setting with only exogenous breakup to situations in which players have the ability to dissolve matches? Surprisingly, allowing players to end matches at any point does not change the qualitative results. First, with fully

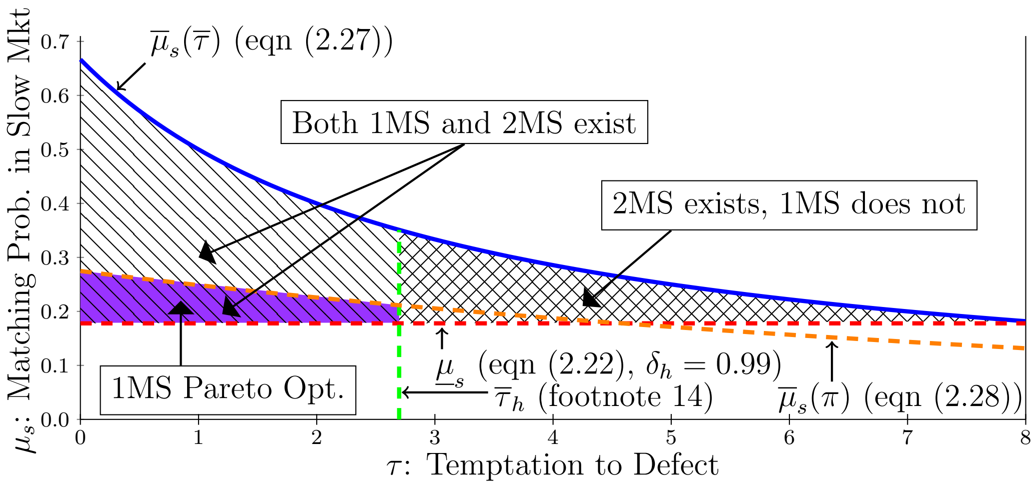

Figure 2.4: Range of matching probabilities in slow market (μ_s) for which the one-market separating equilibrium (1MS) and two-market separating equilibrium (2MS) exist vs. temptation to defect (τ) for $\delta_h = 0.99$ when the low type is impatient enough that they would defect if they entered slow market (with $\pi = 0.33$, $\beta = 0.1$, and $c-d = 1$). 2MS exists for (τ, μ_s) pairs in the entire shaded region, while 1MS only exists in the shaded region to the left of the vertical dashed line. The subset shaded solid below the downward-sloping dashed line is the Pareto improvement region of Proposition 8 in which both types prefer 1MS to 2MS.

endogenous breakup,[15] the first-best outcome is achievable in a similar set of circumstances to what was feasible with exogenous breakup, with the additional requirement that the matching probability cannot be too high. Second, separation, both within and across markets, can still improve welfare when the first-best is not feasible. In this section we will briefly explore the conditions required for the first-best outcome because these conditions capture much of the intuition of the endogenous breakup setting and highlight the one key difference. The additional steady-state requirements of the one-market separating equilibrium with endogenous breakup will also be discussed. Then, given the intuition developed here, the formal model with endogenous breakup will be presented in Section 2.7.

[15]Fully endogenous breakup means that players have the ability to end a match at any point. An exogenous breakup probability, β, remains in place to avoid absorbing states.

Definitions 1 and 2 of equilibria with exogenous breakup with one and two markets, respectively, require that stage game strategies and market entry be individually rational given beliefs, which must be consistent in equilibrium. To be an equilibrium in a setting with endogenous breakup, as given by Definition 3 in Section 2.7, we add to these conditions the requirement that endogenous breakup decisions be individually rational for all players, given their beliefs, as well as a consistency requirement on these beliefs. *No other changes or restrictions are made to the strategies or equilibrium conditions.*

Players, given their beliefs, choose to end a match if and only if the value expected from ending the match and entering the unmatched state is greater than the value expected from staying matched. The value expected when unmatched now depends not only on the matching probability and beliefs about the likelihood a new match will follow the grim trigger strategy or play D, but also on the decision a match will make about whether to stay matched after the stage game strategies have been revealed in the first period of a match.

2.6.1 First-Best Equilibrium with Endogenous Breakup

Consider the first-best outcome in which all players cooperate. Without endogenous breakup, the decision to cooperate or not depends on weighing the value of receiving the temptation payoff today followed by the payoff of an uncooperative match thereafter compared to receiving the payoff of a cooperative match each period, subject to the exogenous breakup possibility. Now players in deciding if defection is a profitable deviation must also take into account the possibility of endogenous breakup. Unlike when all breakup was exogenous, the matching probability now affects how profitable it is to deviate and thus when the first-best outcome is possible. The following proposition, which is proved in Section 2.7, summarizes when the first-best equilibrium is possible with endogenous breakup.

Proposition 9 (First-Best with Endogenous Breakup)**.** *With endogenous breakup, the first-best equilibrium in which all players cooperate exists if and only if either*

1. $\frac{d}{c+\tau} < \mu_f$, *and either* $\mu_f \leq \overline{\mu} \equiv 1 - \frac{\tau}{\delta_\ell(1-\beta)(c+\tau)}$ *or equivalently* $\frac{\tau}{(1-\mu_f)(1-\beta_f)(c+\tau)} \equiv \underline{\delta}(\mu_f) \leq \delta_{t_i}$, *or*
2. $\mu_f \leq \frac{d}{c+\tau}$ *and* $\frac{\tau}{(1-\beta)(c-d+\tau)} \equiv \underline{\delta}_1 \leq \delta_{t_i}$

The proof of Proposition 9 requires examining the endogenous breakup decision of players on and off the equilibrium path. The intuition for these endogenous breakup decisions is discussed below and summarized in Table 2.2, but the complete exploration that includes all the required notation is left for Section 2.7. To see the intuition for Proposition 9, consider the payoff a player expects to receive in the period after choosing to deviate. After deviating, the player will get d next period if the match stays intact and expects to get $\mu_f(c+\tau)$ next period if the match ends, receiving a new match with probably μ_f and receiving $c+\tau$ by deviating again against this new match. If $\mu_f(c+\tau) > d$, a player who defects expects to receive a higher payoff next period by ending the match and defecting again than by staying in the match with the player against whom he just defected. Thus if $\mu_f > \frac{d}{c+\tau}$ (case 1 of Proposition 9), to find cooperation optimal a player must be more patient when matches can be ended endogenously than when all breakup is exogenous, requiring $\delta_{t_i} \geq \underline{\delta}(\mu_f) > \underline{\delta}_1$ instead of just $\delta_{t_i} \geq \underline{\delta}_1$. If $\mu_f(c+\tau) \leq d$, then a player who defects expects to receive a higher payoff next period by staying in the match. Thus if $\mu_f \leq \frac{d}{c+\tau}$ (case 2 of Proposition 9), the condition required for cooperation to be optimal is the same with and without endogenous breakup, requiring $\delta_{t_i} \geq \underline{\delta}_1$.

The $(1-\mu_f)$ term in the denominator of $\underline{\delta}(\mu_f)$ captures the key change in incentives when players can end matches endogenously. Endogenous breakup allows for players to end a match if defected against, but it also increases the temptation for defection because now players can defect and then end the match to go on and defect again against a new match rather than staying in the uncooperative match. The intuition for when a defecting player would find it optimal to stay matched is the same as the intuition for when a player who is the low type would find it optimal to deviate and enter the slow market in the two-market separating equilibrium with exogenous breakup: is it better to take the payoff from an uncooperative match with certainty, or is the temptation payment worth the wait?

This change in incentives with the possibility of endogenous breakup means that there is an upper bound on the matching probability above which it is not possible to achieve the first-best because the temptation to deviate becomes too great when defecting players can end matches endogenously (the $\overline{\mu}$ in case 1 of Proposition 9). Figure 2.5 shows, as a function of the temptation to defect and for different levels of δ_ℓ, the highest the matching probability can be in order for the first-best equilibrium with endogenous breakup to be achievable. For all values of δ_ℓ, the relationship is depicted for $\delta_h = 0.99$, $c - d = 1$, and $\beta = 0.01$.[16] Consider the outermost line, which shows this relationship between μ_f and τ when $\delta_\ell = 0.97$. The matching probability can only be equal to 1 if the temptation to defect were 0, and then the maximum matching probability for which the first-best is achievable declines as the temptation to defect increases. For $\tau = 24.18$, the maximum matching probability for which the first-best is achievable is $\mu_f = 7.36\%$. For $\tau > 24.18$, the first-best is not achievable for any μ_f because the low types will no longer find cooperating optimal; that is, $\delta_\ell = 0.97$ is no longer patient enough to sustain cooperation. For the values of δ_ℓ depicted in Figure 2.5, the combinations of μ_f and τ for which the first-best is achievable are bounded above and to the right by the line corresponding with that value of δ_ℓ, with the highest μ_f for any value of τ prefered by both types. For values of μ_f above the line or values of τ to the right of the line, the temptation to defect is too great and the first-best is not achievable.

The maximum value of τ for which the first-best is achievable is the same with endogenous breakup as it was with exogenous breakup. What changes is that instead of the first-best being achievable with certain matching as it was with exogenous breakup, with endogenous breakup, the maximum matching probability is lower and declines with τ. When the temptation to defect is equal to the value of a cooperative relationship, the first-best is achievable for a matching probability around 70%. And when the temptation to defect is 24 times greater than the value of a cooperative relationship, the first-best is achievable for a matching probability around 7%. Certain matching remains possible, however, for both separating equilibria, and the situations in which separation provides Pareto improvement opportunities remain strikingly similar to the results found in the setting without endogenous breakup.

Further intuition for the model with endogenous breakup can be gained by examining when players choose to stay matched and when they do not. On the equilibrium path in the first-best equilibrium, all players cooperate and choose to stay matched. However, if a match defects, whether a player chooses to end the match in response depends on the matching probability. If the matching probability is low, both types would choose to stay matched despite their match's defection; while both types would prefer to end the uncooperative match and receive a new, cooperative match, the matching probability is too low so neither type finds it worth the wait. If the matching probability is higher, following a defection by a match, the

[16]Recall that in the environment with fully endogenous breakup, players have the ability to end a match at any point but there remains an exogenous breakup probability to avoid absorbing states.

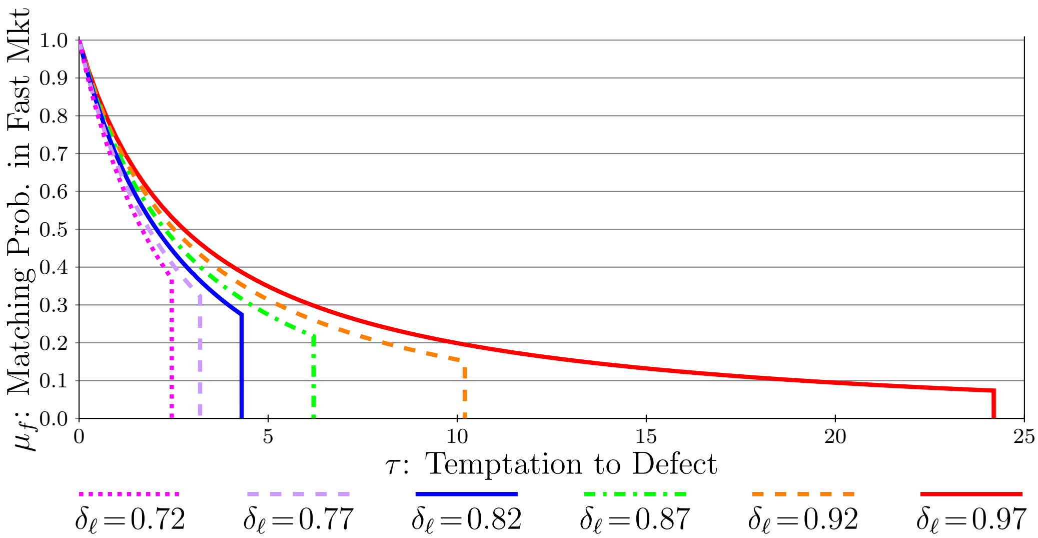

Figure 2.5: Maximum matching probability (μ_f) for achievable first-best with endogenous breakup vs. temptation to defect (τ) for different values of δ_ℓ with $c - d = 1$ and exogenous breakup probability $\beta = 0.01$. For (τ, μ_f) pairs below and left of a line, the first-best equilibrium with endogenous breakup exists for that value of δ_ℓ.

low type would still choose to stay matched but the high type would not; that is, the matching probability is now high enough that the high type do find it to be worth it to end the current uncooperative match and wait for a new, cooperative match. If the matching probability is higher still, the wait for a new match is sufficiently short so both types choose to end an uncooperative match. And if the matching probability is in the highest range, both types would choose to end an uncooperative match if the first-best equilibrium was feasible, but it is not, as previously discussed. The endogenous breakup decision made by players of each type in each of these four ranges for the matching probability are summarized in the Table 2.2.[17]

The fact that the matching probability must be lower for any level of τ to sustain the first-best equilibrium with endogenous breakup means that the utility

[17]A similar table is presented in Table 2.3, that includes the specific notation used for the endogenous breakup decision, which has been modified here because this notation is not presented until Section 2.7. The table there also includes the decisions that would be made by players if they were to defect (i.e., off the equilibrium path).

If matching prob. is	1^{st} best feasible if	Breakup decision if match defects
$\overline{\mu} < \mu$	Not feasible $\forall \delta_{t_i}$	Both types would end match (if EQ feasible)
$\widetilde{\mu}(\delta_\ell) < \frac{d}{c+\tau} < \mu \leq \overline{\mu}$	$\underline{\delta}(\mu) \leq \delta_{t_i}$	Both types end match
$\widetilde{\mu}(\delta_\ell) < \mu \leq \frac{d}{c+\tau} \leq \overline{\mu}$	$\underline{\delta}_1 \leq \delta_{t_i}$	Both types end match
$\widetilde{\mu}(\delta_h) < \mu \leq \widetilde{\mu}(\delta_\ell)$	$\underline{\delta}_1 \leq \delta_{t_i}$	Low type stay matched, high type end match
$\mu \leq \widetilde{\mu}(\delta_h)$	$\underline{\delta}_1 \leq \delta_{t_i}$	Both types stay matched

$\overline{\mu} \equiv 1 - \frac{\tau}{\delta_\ell(1-\beta)(c+\tau)}$, $\widetilde{\mu}(\delta_{t_i}) \equiv \frac{\left[1-\delta_{t_i}(1-\beta)\right]\frac{d}{c}}{\left[1-\delta_{t_i}(1-\beta)\frac{d}{c}\right]}$, $\underline{\delta}(\mu) \equiv \frac{\tau}{(1-\mu)(1-\beta)(c+\tau)}$, $\underline{\delta}_1 \equiv \frac{\tau}{(1-\beta)(c-d+\tau)}$

Table 2.2: Endogenous breakup decision of players in the first-best equilibrium with endogenous breakup for different values of the matching probability

expected in the first-best equilibrium by each type is lower when endogenous breakup is possible, even though endogenous breakup never occurs on the equilibrium path in the first-best equilibrium.

Proposition 10 (First-Best Utility Loss from Endogenous Breakup)**.** *The ratio of ex ante expected utility from the first-best equilibrium with endogenous breakup over the ex ante expected utility from the first-best equilibrium without endogenous breakup is*

$$\text{High type: } \frac{\delta_\ell(1-\beta_f)(c+\tau)-\tau}{\delta_\ell(1-\beta_f)(c+\tau)-\delta_h(1-\beta_f)\tau} \tag{2.30}$$

$$\text{Low type: } \frac{\delta_\ell(1-\beta_f)(c+\tau)-\tau}{\delta_\ell(1-\beta_f)c} \tag{2.31}$$

These expressions are obtained by comparing the highest possible expected utility for each type in each situation, with and without endogenous breakup. The ex ante utility expected by unmatched player i in the first-best equilibrium is

$$\frac{\mu_f c}{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f))-\delta_{t_i}\mu_f\beta_f}$$

This expected utility is highest when μ_f is as high as possible. Without endogenous breakup, certain matching is possible, so the highest ex ante expected utility occurs when $\mu_f = 1$. With endogenous breakup, the highest the matching probability

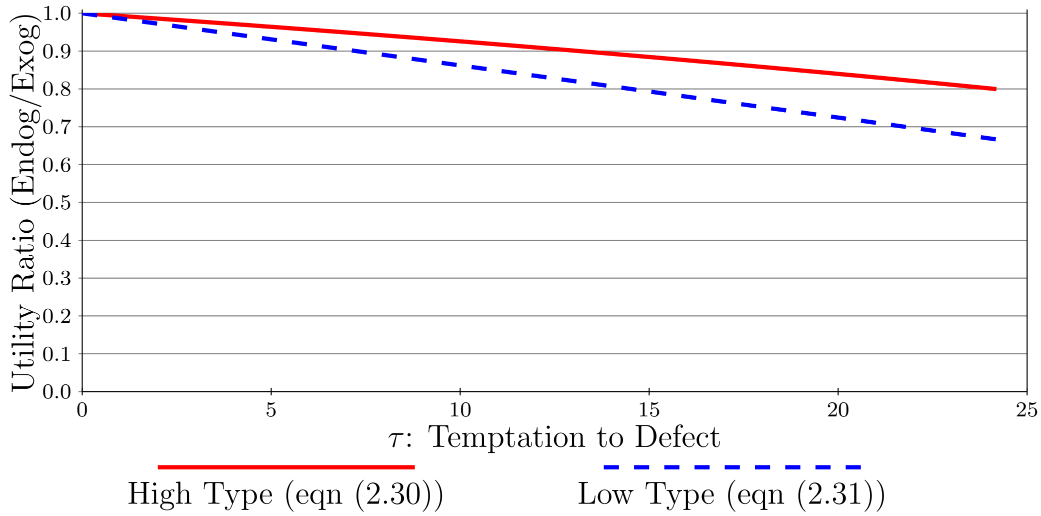

Figure 2.6: For first-best equilibrium, ratio of ex ante expected utility from first-best with endogenous breakup over ex ante expected utility from first-best with exogenous breakup vs. temptation to defect (τ) with $\delta_h = 0.99$, $\delta_\ell = 0.97$, $c - d = 1$, and exogenous breakup probability $\beta = 0.01$ (Proposition 10).

can be is $\overline{\mu}$, as given in Proposition 9. By substituting in $\mu_f = 1$ and $\mu_f = \overline{\mu}$ for the utility expected with and without endogenous breakup, respectively, and comparing the ratio, the expressions in Proposition 10 are obtained. Note that $\overline{\mu}$ includes δ_ℓ, which is why the expression for the high type includes both δ_h and δ_ℓ.

Figure 2.6 shows the ratio of the ex ante expected utility with endogenous breakup relative to the ex ante expected utility without exogenous breakup for the high type ($\delta_h = 0.99$) and low type ($\delta_\ell = 0.97$) when $c - d = 1$ and $\beta = 0.01$. The utility lost by allowing for endogenous breakup is low when the temptation to defect is low and increases as the temptation to defect increases, with the low type losing more utility than the high type. However, even when the temptation to defect is 24 times higher than the value of a cooperative relationship, the ex ante expected utility from the first-best equilibrium with endogenous breakup is 67% and 80% of that with exogenous breakup for the low and high type, respectively. Thus, while allowing agents the possibility of ending matches does lower the utility expected by agents, the penalty that comes with the increased flexibility of allowing for endogenous breakup is not that large.

2.6.2 One-Market Separating Equilibrium with Endogenous Breakup

In the first-best equilibrium, both with and without endogenous breakup, all players cooperate. This means that all players expect a new match to cooperate, so there is no need to examine anything regarding the distribution of unmatched players. In the two-market separating equilibrium, both with and without endogenous breakup, the types play different actions but are separated by market, again making it unnecessary to examine the distribution of unmatched players. In the one-market separating equilibrium, however, both cooperation and defection exist simultaneously. Thus expectations about whether a new match will cooperate depend on the probability that an unmatched player is the high type. If one type is more or less likely to be unmatched than the other type, then the probability that an unmatched player is of a particular type is not equal to that type's proportion in the population as a whole that includes both matched and unmatched players. Thus it is necessary to examine the distribution of types, high and low, as well as states, matched and unmatched, in the population.

In the one-market separating equilibrium with endogenous breakup, the endogenous breakup decisions of players of each type follow a similar pattern to that seen in the first-best equilibrium. If the matching probability is low, both types choose to stay matched. If the matching probability is increased, the high type find it optimal to end an uncooperative match in hopes of receiving a cooperative match, while the low type find it optimal to stay matched. If the matching probability is increased further, both the high and low types find it optimal to end uncooperative matches, the high type with the hope of receiving a cooperative match and the low type with the hope of being matched with a high type again and receiving the temptation payoff. We will examine each of these cases in turn.

In Section 2.4 when we examined the one-market separating equilibrium without endogenous breakup, each player's belief that a new match would be the high type was $p_f = \pi$. With the possibility of endogenous breakup, p_f is equal to the steady-state probability that a new match is the high type (defined as H below). This probability is found for each case in the following sections. Note that for the lowest range of the matching probability, players never choose to end matches endogenously even though they are allowed to do so. In this case, we expect to find that $p_f = \pi$ as was the case when endogenous breakup was not allowed. This is indeed what we find.

2.6.2.1 Distribution of types and states

Before examining the stationary distribution in each of the three cases, we first must introduce some notation. Each player in the population is either the high or low type and can be in one of three states: matched with a high type, matched with a low type, or unmatched. Thus the population is broken down into six type-states, which in vector form are given by

$$\Pi = (\pi_{hh}, \pi_{h\ell}, \pi_{hu}, \pi_{\ell h}, \pi_{\ell\ell}, \pi_{\ell u})$$

where the first subscript refers to the player's type and the second refers to his state, with h and ℓ indicating the player is currently matched with a high and low type, respectively, and u denoting that the player is unmatched. Players transition between states according to the following transition matrix

$$P = \begin{array}{c} \\ h_h \\ h_\ell \\ h_u \\ \ell_h \\ \ell_\ell \\ \ell_u \end{array} \begin{array}{c} \begin{array}{cccccc} h_{h'} & h_{\ell'} & h_{u'} & \ell_{h'} & \ell_{\ell'} & \ell_{u'} \end{array} \\ \begin{pmatrix} h_{hh'} & h_{h\ell'} & h_{hu'} & 0 & 0 & 0 \\ h_{\ell h'} & h_{\ell\ell'} & h_{\ell u'} & 0 & 0 & 0 \\ h_{uh'} & h_{u\ell'} & h_{uu'} & 0 & 0 & 0 \\ 0 & 0 & 0 & \ell_{hh'} & \ell_{h\ell'} & \ell_{hu'} \\ 0 & 0 & 0 & \ell_{\ell h'} & \ell_{\ell\ell'} & \ell_{\ell u'} \\ 0 & 0 & 0 & \ell_{uh'} & \ell_{u\ell'} & \ell_{uu'} \end{pmatrix} \end{array}$$

The column to the left of the matrix denotes the player's type and state (the subscript) at the start of the current period, while the row above the matrix denotes the player's type and state (the subscript) at the start of next period. The elements of the matrix denote the probability of transitioning from one type-state to another. For each element, the first subscript refers to the player's state at the start of this period while the second subscript (with the prime) refers to his state at the start of next period. For example, consider the element $\ell_{uh'}$. This is the probability that a low type starts this period unmatched (shown in the column to the left of the matrix as ℓ_u) and starts next period matched with a high type (shown in the row above the matrix as $\ell_{h'}$). Because players do not change type, the probability of h_h, h_ℓ, or h_u transitioning to $\ell_{h'}$, ℓ_ℓ, or $\ell_{u'}$ is 0, as is the probability of ℓ_h, ℓ_ℓ, or ℓ_u transitioning to $h_{h'}$, $h_{\ell'}$, or $h_{u'}$.

The high types, who make up a fraction π of the whole population, and the low types, who make up a fraction $1 - \pi$ of the whole population, can each be in one of the three states. Thus

$$\pi = \pi_{hh} + \pi_{h\ell} + \pi_{hu} \tag{2.32}$$

$$1 - \pi = \pi_{\ell h} + \pi_{\ell\ell} + \pi_{\ell u} \tag{2.33}$$

The probability that a new match will be the high type is equal to the likelihood that out of all unmatched players, a player is the high type. Thus the probability that a new match is the high type or the low type are given by

$$H \equiv \frac{\pi_{hu}}{\pi_{hu} + \pi_{\ell u}} \tag{2.34}$$

$$L \equiv \frac{\pi_{\ell u}}{\pi_{hu} + \pi_{\ell u}} = 1 - H \tag{2.35}$$

To find a stationary distribution, we need to find Π^* such that

$$\Pi^* P = \Pi^* \tag{2.36}$$

We are now ready to examine the stationary distribution in the one-market separating equilibrium in each of the three cases for the matching probability mentioned above.

2.6.2.2 Lowest matching probability: both types choose to stay matched

If the matching probability is low enough, all players prefer staying in an existing match rather than ending the match with hopes of receiving a new one. This means that a player who starts a period matched will be matched again next period with the same match unless exogenous breakup occurs. Thus $h_{hh'} = h_{\ell\ell'} = \ell_{hh'} = \ell_{\ell\ell'} = 1 - \beta_f$ and $h_{hu'} = h_{\ell u'} = \ell_{hu'} = \ell_{\ell u'} = \beta_f$. A player who starts the period unmatched receives a match with probability μ_f and stays matched until the start of next period with probability $1-\beta_f$. A new match is the high type with probability H and the low type with probability L. This means that $h_{uh'} = \ell_{uh'} = \mu_f H(1-\beta_f)$ and $h_{u\ell'} = \ell_{u\ell'} = \mu_f L(1-\beta_f)$. The probability that an unmatched player stays unmatched is $h_{uu'} = \ell_{uu'} = 1 - \mu_f - \mu_f\beta_f = 1 - \mu_f(1 - \beta_f)$, where $1 - \mu_f$ is the probability of not receiving a match and $\mu_f H\beta_f + \mu_f L\beta_f = \mu_f\beta_f$ is the probability of receiving a match but it ending exogenously. Thus, in this case, the transition matrix is

$$P = \begin{array}{c} \\ h_h \\ h_\ell \\ h_u \\ \ell_h \\ \ell_\ell \\ \ell_u \end{array} \begin{array}{c} \begin{array}{cccccc} h_{h'} & h_{\ell'} & h_{u'} & \ell_{h'} & \ell_{\ell'} & \ell_{u'} \end{array} \\ \begin{pmatrix} 1-\beta_f & 0 & \beta_f & 0 & 0 & 0 \\ 0 & 1-\beta_f & \beta_f & 0 & 0 & 0 \\ \mu_f H(1-\beta_f) & \mu_f L(1-\beta_f) & 1-\mu_f(1-\beta_f) & 0 & 0 & 0 \\ 0 & 0 & 0 & 1-\beta_f & 0 & \beta_f \\ 0 & 0 & 0 & 0 & 1-\beta_f & \beta_f \\ 0 & 0 & 0 & \mu_f H(1-\beta_f) & \mu_f L(1-\beta_f) & 1-\mu_f(1-\beta_f) \end{pmatrix} \end{array}$$

Using this transition matrix in (2.36) (and omitting the *'s on all the π terms), and substituting in for H and L using (2.34) and (2.35), we have the following system of equations

$$\begin{aligned}
\pi_{hh}(1-\beta_f) + \pi_{hu}\mu_f(1-\beta_f)\frac{\pi_{hu}}{\pi_{hu}+\pi_{\ell u}} &= \pi_{hh} \\
\pi_{h\ell}(1-\beta_f) + \pi_{hu}\mu_f(1-\beta_f)\frac{\pi_{\ell u}}{\pi_{hu}+\pi_{\ell u}} &= \pi_{h\ell} \\
\pi_{hh}\beta_f + \pi_{h\ell}\beta_f + \pi_{hu}(1-\mu_f(1-\beta_f)) &= \pi_{hu} \\
\pi_{\ell h}(1-\beta_f) + \pi_{\ell u}\mu_f(1-\beta_f)\frac{\pi_{hu}}{\pi_{hu}+\pi_{\ell u}} &= \pi_{\ell h} \\
\pi_{\ell\ell}(1-\beta_f) + \pi_{\ell u}\mu_f(1-\beta_f)\frac{\pi_{\ell u}}{\pi_{hu}+\pi_{\ell u}} &= \pi_{\ell\ell} \\
\pi_{\ell h}\beta_f + \pi_{\ell\ell}\beta_f + \pi_{\ell u}(1-\mu_f(1-\beta_f)) &= \pi_{\ell u}
\end{aligned}$$

Using equations (2.32) and (2.33) and simplifying, the stationary distribution of types and states, Π^*, is

$$\Pi^* = \left(\tfrac{\mu_f(1-\beta_f)\pi^2}{\mu_f(1-\beta_f)+\beta_f}, \tfrac{\mu_f(1-\beta_f)\pi(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}, \tfrac{\beta_f\pi}{\mu_f(1-\beta_f)+\beta_f}, \tfrac{\mu_f(1-\beta_f)\pi(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}, \tfrac{\mu_f(1-\beta_f)(1-\pi)^2}{\mu_f(1-\beta_f)+\beta_f}, \tfrac{\beta_f(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}\right)$$

The probabilities that a new match will be the high type or the low type are

$$H = \frac{\pi_{hu}}{\pi_{hu}+\pi_{\ell u}} = \frac{\frac{\beta_f\pi}{\mu_f(1-\beta_f)+\beta_f}}{\frac{\beta_f\pi}{\mu_f(1-\beta_f)+\beta_f} + \frac{\beta_f(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}} = \frac{\beta_f\pi}{\beta_f\pi + \beta_f(1-\pi)} = \frac{\pi}{\pi+1-\pi} = \pi$$

$$L = \frac{\pi_{\ell u}}{\pi_{hu}+\pi_{\ell u}} = \frac{\frac{\beta_f(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}}{\frac{\beta_f\pi}{\mu_f(1-\beta_f)+\beta_f} + \frac{\beta_f(1-\pi)}{\mu_f(1-\beta_f)+\beta_f}} = \frac{\beta_f(1-\pi)}{\beta_f\pi + \beta_f(1-\pi)} = \frac{1-\pi}{\pi+1-\pi} = 1-\pi$$

This is exactly what we expect to find. This case is equivalent to the case without endogenous breakup because no player chooses to exercise endogenous breakup. Thus the probability that a new match is the high type is the same as the fraction of high types in the entire population, and similarly for the low types, because players of each type are equally likely to be in the unmatched state. This is shown by the dotted line in Figure 2.7. The fraction of the population who are the high type (π) is shown on the x-axis, while the y-axis is the probability that a new match is the high type (H). Since $H = \pi$ for this lowest range of μ_f, this is the 45° line.

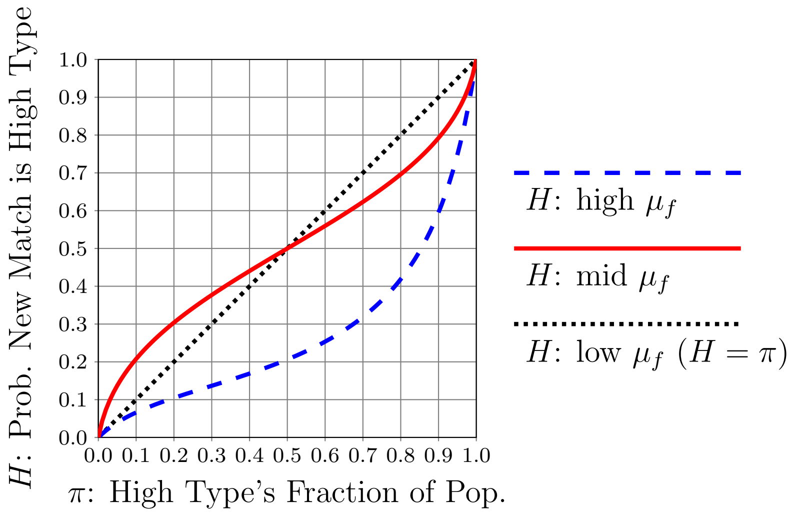

Figure 2.7: The probability that a new match is the high type (H, y-axis) vs. the fraction of the population who are the high type (π, x-axis), for the three ranges of the matching probability (μ_f) that determine the endogenous breakup decision of players in the one-market separating equilibrium with $\beta_f = 0.1$. For the lowest range of μ_f (dotted 45° line), all players choose to stay matched, so $H = \pi$. For the middle range of μ_f (solid line, shown for $\mu_f = 0.5$), the low type choose to stay matched but the high type end uncooperative matches, causing $H > \pi$ when $\pi < \frac{1}{2}$ and $H < \pi$ when $\pi > \frac{1}{2}$. For the highest range of μ_f (dashed line, shown for $\mu_f = 1$), both types choose to end uncooperative matches, causing $H < \pi$.

2.6.2.3 Middle matching probability: high type choose to end uncooperative matches

If the matching probability is increased, the high type find it optimal to end any match with a low type in favor of re-entering the unmatched state and hoping to receive a cooperative match with a high type. However, the low type, while they would prefer to receive a new match with a high type and receive the temptation payment, the wait is still too long so they find it optimal to stay matched. However, since if they are matched with a high type the high type will end the match, the only matches that stay together involving low types are those with two low types matched together. In this situation, the transition matrix is

$$P = \begin{array}{c} \\ h_h \\ h_\ell \\ h_u \\ \ell_h \\ \ell_\ell \\ \ell_u \end{array} \begin{array}{c} \begin{array}{cccccc} h_{h'} & h_{\ell'} & h_{u'} & \ell_{h'} & \ell_{\ell'} & \ell_{u'} \end{array} \\ \begin{pmatrix} 1-\beta_f & 0 & \beta_f & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ \mu_f H(1-\beta_f) & 0 & 1-\mu_f H(1-\beta_f) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1-\beta_f & \beta_f \\ 0 & 0 & 0 & 0 & \mu_f L(1-\beta_f) & 1-\mu_f L(1-\beta_f) \end{pmatrix} \end{array}$$

After solving the system for the steady state distribution, the probability of being matched with a high type is

$$H = \frac{\mu_f(1-\beta_f)\pi - \beta_f(\frac{1}{2}-2\pi) + \frac{1}{2}\sqrt{8\mu_f(1-\mu_f)\beta_f\pi(1-\pi)+4\mu_f^2\pi(1-\pi)+\beta_f^2(1-4\mu_f(2-\mu_f)\pi(1-\pi))}}{\beta_f+\mu_f(1-\beta_f)+\sqrt{8\mu_f(1-\mu_f)\beta_f\pi(1-\pi)+4\mu_f^2\pi(1-\pi)+\beta_f^2(1-4\mu_f(2-\mu_f)\pi(1-\pi))}}$$

This is shown by the solid line in Figure 2.7 for $\mu_f = 0.5$ and $\beta_f = 0.1$. As seen in the figure, $H > \pi$ if $\pi < \frac{1}{2}$ and $H < \pi$ if $\pi > \frac{1}{2}$. If less than half the population is the high type, then high types are more likely to be matched with low types, matches that the high types will end endogenously. This means that high types are relatively more likely to be unmatched than matched, and the probability of a new match being the high type is higher than the fraction of the population who are high types. Similarly, if more than half the population are high types, high types are more likely to be matched with other high types, matches that stay intact unless breakup occurs exogenously. This means that high types are relatively more likely to be matched than unmatched, meaning that the probability that a new match is the high type is less than the fraction of high types in the population. As μ_f increases and β_f falls, H departs further from π (the 45° line), and as μ_f decreases or β_f increases, H collapses to π (the 45° line). If exactly half of the population is the high type, then the probability of a new match being the high type is also one half. This is the point of intersection where H goes from greater than π to less than π, crossing the 45° line.

2.6.2.4 Highest matching probability: high and low type choose to end uncooperative matches

If the matching probability is increased further, both the high and low types find it optimal to end uncooperative matches, which is every match for low types and every match with a low type for high types. In this highest range for μ_f, both types prefer to end a current uncooperative match with hopes of receiving a new match who is the high type, the high types because they wish to be in a cooperative match and the low types because they would like to receive the temptation payoff. In this situation, the transition matrix is

$$P = \begin{array}{c} \\ h_h \\ h_\ell \\ h_u \\ \ell_h \\ \ell_\ell \\ \ell_u \end{array} \begin{array}{c} \begin{array}{cccccc} h_{h'} & h_{\ell'} & h_{u'} & \ell_{h'} & \ell_{\ell'} & \ell_{u'} \end{array} \\ \begin{pmatrix} 1-\beta_f & 0 & \beta_f & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ \mu_f H(1-\beta_f) & 0 & 1-\mu_f H(1-\beta_f) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

After solving for the steady-state distribution, the likelihood of a new match being the high type is

$$H = \frac{2\pi\sqrt{\beta_f}}{\sqrt{\beta_f} + \sqrt{\beta_f + 4\mu_f(1-\beta_f)\pi(1-\pi)}}$$

This is shown by the dashed line in Figure 2.7 for $\mu_f = 1, \beta_f = 0.1$. As seen in the figure, the probability of a new match being the high type is less than the fraction of high types in the population ($H < \pi$). This occurs because any high type who is matched with another high type chooses to stay matched, meaning relatively fewer high types are unmatched than are matched, reducing the odds of a new match being the high type below what it would be without endogenous breakup when all players are equally likely to be unmatched.

As μ_f falls or as β_f rises, H becomes closer to π (the 45° line). As μ_f rises or β_f falls, H bows down further below π (the 45° line). This occurs because the lower the probability of exogenous breakup, the longer high types stay in matches with each other and thus the lower the probability that high types will be unmatched, reducing the likelihood of a new match being the high type.

2.6.3 Hybrid Model: Between Fully Exogenous and Fully Endogenous Breakup

With one additional layer of notational complexity, it is possible to nest both the model with exogenous and endogenous breakup, and everything in between, in one model. The model with endogenous breakup presented in Section 2.7 and discussed in the previous sections introduces two new variables, one for an agent's endogenous breakup decision and one for his beliefs about the endogenous breakup decisions of others. These two variables always appear multiplied together; intuitively, a match only remains intact if both parties wish to stay matched and dissolves if either party wishes to split. These two variables can be multiplied by a third variable that indicates when endogenous breakup is possible that takes on a value of 1 if endogenous breakup is possible and 0 if it is not. By allowing this indicator of possible endogenous breakup to take on any value in $[0, 1]$, everything in between and including fully endogenous and fully exogenous breakup can be nested in one model.[18]

Players would have to consider not only if they want to end a match endogenously, but also when they expect to have the ability to do so. It would also be possible to have the ability for endogenous breakup to vary by type, which could be combined with the extension of the stage-game payoffs to depend on type, allowing the model to capture a wider range of two-sided interactions. While somewhat more complicated than the special cases of endogenous breakup always or never being allowed presented here, given the similarities between the model with fully exogenous and endogenous breakup, the intuition presented in this paper extends to the hybrid model.

2.7 Model with Endogenous Breakup

The previous section highlighted the similarities and differences between a setting in which all breakup is exogenous and in which endogenous breakup is permitted. This section presents the model discussed in the previous section in detail, extending the model developed in Section 2.3 to allow for fully endogenous breakup. The market structure remains the same, with probability of matching μ_m in market $m \in \{s, f\}$ where $m = s$ and $m = f$ still refer to the slow and fast market, respectively, with all one-market equilibria occurring in the fast market as before. We will assume that there is still an exogenous chance of breakup, β_m, just as before, in order to avoid absorbing states and the need to more explicitly specify the matching technology.

Players now have the option to dissolve a match endogenously each period when matched after playing the stage game. Let $\phi^m_{t_i}(\sigma^{gt}, p_m)$ be the endogenous

[18] More details are provided in Section 2.7.6.

breakup decision of player i type t_i following the grim trigger strategy in market $m \in \{s, f\}$ with beliefs $p_m \in \{0, 1\}$ about the likelihood that the other player will cooperate. Similarly, let $\phi^m_{t_i}(D, p_m)$ be the endogenous breakup decision but when player i is always playing D. Just as in the exogenous breakup case, since each player's strategy is revealed by the action taken by each in the first period of a match and all players are either following the grim trigger strategy or always playing D, the beliefs of each player after stage game play in the first period of a new match will be either $p_m = 1$ if each cooperated and $p_m = 0$ otherwise. To save on notation we will refer to $\phi^m_{t_i}$ without argument when the arguments are clear from context or when the statement applies to all of the possible combinations of strategies, σ^{gt} or D, and beliefs, $p_m \in \{0, 1\}$. Let $\phi^m_{t_i} = 0$ be the decision of player i type t_i to end a current match in favor of becoming unmatched, while $\phi^m_{t_i} = 1$ is his decision to stay matched. Player i's beliefs about the endogenous breakup decision to be made by other players is denoted $\theta^m_{t_j}$ with analogous arguments and definition.

In equilibrium, consistency will require that the endogenous breakup decision made by players, $\phi^m_{t_i}(\sigma^{gt}, p_m)$ and $\phi^m_{t_i}(D, p_m)$ matches with beliefs about what other players will do, $\theta^m_{t_j}(\sigma^{gt}, p_m)$ and $\theta^m_{t_j}(D, p_m)$, for each type, $t_i, t_j \in \{h, \ell\}$, market $m \in \{s, f\}$, and beliefs $p_m\{0, 1\}$, where j refers to all players with which player i could be matched, $j \neq i$. Consistency must be verified by checking that all decisions are individually rational given beliefs both on the equilibrium path and off the equilibrium path on paths that must be evaluated by players to evaluate potentially profitable deviations. For example, in the first-best outcome in which all players cooperate, each player knows what his own optimal decision would be if defected against and what it would be if choosing to deviate and defect himself, and thus using the same reasoning he can form beliefs about what other players' optimal decisions would be in each case. Only when a player knows what they will do in each case and has beliefs about what a match will do can a player evaluate the utility expected from following each strategy and thus determine if deviation is profitable or not in terms of providing higher expected utility. Thus while only $\phi^f_{t_i}(\sigma^{gt}, 1)$ will occur on the equilibrium path in the first-best equilibrium, players must determine their optimal decisions if defected against, $\phi^f_{t_i}(\sigma^{gt}, 0)$, and if choosing to defect themselves, $\phi^f_{t_i}(D, 0)$. Thus in equilibrium we will require consistency not only of beliefs about actions taken on the equilibrium path, $\theta^f_{t_j}(\sigma^{gt}, 1)$, but also of $\theta^f_{t_j}(\sigma^{gt}, 0)$ and $\theta^f_{t_j}(D, 0)$. Since in the first-best outcome all players cooperate in the fast market, beliefs about what might happen in the slow market can be ignored.

In a separating equilibrium the type of the other player is revealed by his action but in a pooling equilibrium it is not. Thus in a pooling equilibrium beliefs about the endogenous breakup decision of the other player are the population-weighted average of beliefs about the decision each type will find optimal. For example, suppose a pooling equilibrium in which all players cooperate exists. In deciding if deviating to D is profitable, a player must form beliefs about whether the other player will end the match after being defected against, $\theta^m_{t_j}(\sigma^{gt}, 0)$ for $t_j \in \{h, \ell\}$. It is not possible to observe the other player's type, so beliefs must be the average of the decision expected by each type weighted by the proportion of the population who are each type.[19] For example, suppose the high type will choose to end the match if defected against but low types will choose to stay matched. The expectations about the likelihood of the other player ending the match are $\theta^f(\sigma^{gt}, 0) \equiv \pi\theta^f_h(\sigma^{gt}, 0) + (1-\pi)\theta^f_\ell(\sigma^{gt}, 0) = 1 - \pi$. The beliefs $\theta^f(\sigma^{gt}, 1)$ and $\theta^f(D, 0)$ are defined similarly.

A match stays together only if both players choose to stay matched ($\phi^{p_m}_{t_i}\theta^m_{t_j} = 1 \cdot 1 = 1$). If either player chooses to end the match, the match dissolves, denoted ($\phi^{p_m}_{t_i}\theta^m_{t_j} = 0$). Since $\theta^m_{t_j}$ is player i's expectation about what player j will do and a match stays together only if both players choose to stay matched, the probability player i assigns to staying matched given that he does not choose to end the match himself is $1 \cdot \theta^m_{t_j} = \theta^m_{t_j}$.[20]

We now have all the notation necessary to develop a model similar to that in Sections 2.3 through 2.5 that allows for endogenous breakup.

2.7.1 Model with Endogenous Breakup

An unmatched player chooses to enter the market or to stay unmatched[21][22]

$$V_{t_i} = \max\{0, V^f_{t_i}(p_f), V^s_{t_i}(p_s)\} \tag{2.37}$$

[19] Recall that the fraction of the population who are the high type, π, and who are the low type, $1-\pi$, as well as the discount rates of each, δ_h and δ_ℓ, are commonly known.

[20] For example, in the example provided in the previous paragraph, this value is $1-\pi$. Also note that if player i decides to end the match, he could have an expectation about whether the other player will want to stay matched but it does not matter what this expectation is because $\phi^{p_m}_{t_i}\theta^m_{t_j} = 0 \cdot \theta^m_{t_j}, \forall \theta^m_{t_j}$.

[21] To consider the setting with only one market, let $\mu_s = 0$, in which case $V^s_{t_i} = 0$, which is equal to the outside option.

[22] As previously noted, while V_{t_i} depends on p_f and p_s, these arguments are omitted.

The value expected by an unmatched player entering market $m \in \{s, f\}$ is

$$V^m_{t_i}(\cdot, \phi^m_{t_i} | \theta^m_{t_j}, p_m) = \overbrace{\mu_m W^m_{t_i}(\cdot, \phi^m_{t_i} | \theta^m_{t_j}, p_m)}^{\text{get matched}} + \overbrace{(1-\mu_m)(0 + \delta_{t_i} V_{t_i})}^{\text{stay unmatched}} \tag{2.38}$$

There is still an exogenous chance of breakup, β_m. We will continue to focus on symmetric strategies, conditional on type. We will also continue the assumption that players will either cooperate following the grim trigger strategy, denoted σ^{gt} and defined by (2.1), or always play D. Including the endogenous breakup decision, the value expected from a new match from following either the grim trigger strategy or always playing D is

$$\begin{aligned} & W^m_{t_i}(\sigma^{gt}, \phi^m_{t_i} | \theta^m_{t_j}, p_m) = \\ & \quad p_m \overbrace{\left(c + \delta_{t_i} \left[\begin{array}{c} \overbrace{\phi^m_{t_i}(\sigma^{gt}, 1)\theta^m_{t_j}(\sigma^{gt}, 1)[(1-\beta_m)W^m_{t_i}(\sigma^{gt}|1) + \beta_m V_{t_i}]}^{\text{no endogenous breakup}} \\ + \overbrace{(1 - \phi^m_{t_i}(\sigma^{gt}, 1)\theta^m_{t_j}(\sigma^{gt}, 1))V_{t_i}}^{\text{endogenous breakup}} \end{array} \right] \right)}^{\text{matched with cooperator}} \\ & \quad + (1-p_m) \overbrace{\left(d - \lambda + \delta_{t_i} \left[\begin{array}{c} \overbrace{\phi^m_{t_i}(\sigma^{gt}, 0)\theta^m_{t_j}(D, 0)[(1-\beta_m)W^m_{t_i}(\sigma^{gt}|0) + \beta_m V_{t_i}]}^{\text{no endogenous breakup}} \\ + \overbrace{(1 - \phi^m_{t_i}(\sigma^{gt}, 0)\theta^m_{t_j}(D, 0))V_{t_i}}^{\text{endogenous breakup}} \end{array} \right] \right)}^{\text{matched with non-cooperator}} \end{aligned} \tag{2.39}$$

$$W^m_{t_i}(D, \phi^m_{t_i} | \theta^m_{t_j}, p_m) =$$

$$\overbrace{p_m \left(c + \tau + \delta_{t_i} \left[\begin{array}{c} \overbrace{\phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0)[(1-\beta_m)W^m_{t_i}(D|0) + \beta_m V_{t_i}]}^{\text{no endogenous breakup}} \\ + \overbrace{(1 - \phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0))V_{t_i}}^{\text{endogenous breakup}} \end{array} \right] \right)}^{\text{matched with cooperator}}$$

$$+ (1-p_m) \overbrace{\left(d + \delta_{t_i} \left[\begin{array}{c} \overbrace{\phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0)[(1-\beta_m)W^m_{t_i}(D|0) + \beta_m V_{t_i}]}^{\text{no endogenous breakup}} \\ + \overbrace{(1 - \phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0))V_{t_i}}^{\text{endogenous breakup}} \end{array} \right] \right)}^{\text{matched with non-cooperator}} \quad (2.40)$$

It only makes sense to consider the value expected from continuing a match if the match continues. Accordingly, the $\phi^m_{t_i}$ and $\theta^m_{t_i}$ terms, which must always equal one for a match to continue, are omitted from the expressions for the value expected from staying matched, $W^m_{t_i}(\cdot|\cdot)$. Players have the option to end a match at any period after the stage game is played in the first period of a match. In the current framework, stage game payoffs are known and do not change, and thus if a player finds it optimal to end a match at any period in a match he will find it optimal to end the match at every period of the match. Thus any endogenous breakup will occur after the first period of a match when each players' strategies are revealed.[23] Accordingly, the second and later periods of a match are only reached when both players choose to stay matched, and in this case, the match will last until dissolved exogenously. Thus the value expected from a current match in the second and subsequent periods of the

[23]If payoffs were stochastic or changed over the duration of a match, then we would need to consider endogenous breakup after each period of a match. For example, it would be interesting to consider if the return for two players cooperating was not the same each period, but rather varied stochastically based on economic conditions, or varied monotonically over time. However, this type of extension is beyond the scope of the present work and is relegated to future consideration.

match is the same as it was before endogenous breakup was possible:

$$W^m_{t_i}(\sigma^{gt}|1) = c + \delta_{t_i}\left[(1-\beta_m)W^m_{t_i}(\sigma^{gt}|1) + \beta_m V_{t_i}\right] \;\Rightarrow\; W^m_{t_i}(\sigma^{gt}|1) = \frac{c + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.41}$$

$$W^m_{t_i}(\sigma^{gt}|0) = d + \delta_{t_i}\left[(1-\beta_m)W^m_{t_i}(\sigma^{gt}|0) + \beta_m V_{t_i}\right] \;\Rightarrow\; W^m_{t_i}(\sigma^{gt}|0) = \frac{d + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.42}$$

$$W^m_{t_i}(D|0) = d + \delta_{t_i}\left[(1-\beta_m)W^m_{t_i}(D|0) + \beta_m V_{t_i}\right] \;\Rightarrow\; W^m_{t_i}(D|0) = \frac{d + \delta_{t_i}\beta_m V_{t_i}}{1-\delta_{t_i}(1-\beta_m)} \tag{2.43}$$

Subbing these into (2.39) and (2.40), we get

$$\begin{aligned} W^m_{t_i}(\sigma^{gt}, \phi^m_{t_i}|\theta^m_{t_j}, p_m) = \\ p_m\left(c + \phi^m_{t_i}(\sigma^{gt},1)\theta^m_{t_j}(\sigma^{gt},1)\frac{\delta_{t_i}(1-\beta_m)c}{1-\delta_{t_i}(1-\beta_m)}\right) \\ + (1-p_m)\left(d - \lambda + \phi^m_{t_i}(\sigma^{gt},0)\theta^m_{t_j}(D,0)\frac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right) \\ + \delta_{t_i}\left(1 - \left(p\phi^m_{t_i}(\sigma^{gt},1)\theta^m_{t_j}(\sigma^{gt},1) + (1-p_m)\phi^m_{t_i}(\sigma^{gt},0)\theta^m_{t_j}(D,0)\right)\frac{(1-\delta_{t_i})(1-\beta_m)}{1-\delta_{t_i}(1-\beta_m)}\right)V_{t_i} \end{aligned} \tag{2.44}$$

$$\begin{aligned} W^m_{t_i}(D, \phi^m_{t_i}|\theta^m_{t_j}, p_m) = p_m(c+\tau) + (1-p_m)d \\ + \delta_{t_i}(1-\beta_m)\left(p_m\phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0) + (1-p_m)\phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0)\right)\frac{d}{1-\delta_{t_i}(1-\beta_m)} \\ + \delta_{t_i}\left(1 - \left(p_m\phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0) + (1-p_m)\phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0)\right)\frac{(1-\delta_{t_i})(1-\beta_m)}{1-\delta_{t_i}(1-\beta_m)}\right)V_{t_i} \end{aligned} \tag{2.45}$$

Just as we did with exogenous breakup, assuming symmetric strategies by type, we can solve to find

$$V^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m) = \frac{\mu_m W^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)}{1-\delta_{t_i}(1-\mu_m)} \tag{2.46}$$

Using this to sub in for V_{t_i} in (2.44) and (2.45), we get the ex ante expected value from a new match in terms of parameters only, now including the endogenous breakup decision.

$$\begin{aligned}
&W^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m) = \\
&\left(\frac{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_m))+\left(p_m\phi^m_{t_i}(\sigma^{gt},1)\theta^m_{t_j}(\sigma^{gt},1)+(1-p_m)\phi^m_{t_i}(\sigma^{gt},0)\theta^m_{t_j}(D,0)\right)\delta_{t_i}\mu_m(1-\beta_m)\right]}\right) \\
&\cdot\left(\begin{array}{c} p_m\left(c+\phi^m_{t_i}(\sigma^{gt},1)\theta^m_{t_j}(\sigma^{gt},1)\dfrac{\delta_{t_i}(1-\beta_m)c}{1-\delta_{t_i}(1-\beta_m)}\right) \\ +(1-p_m)\left(d-\lambda+\phi^m_{t_i}(\sigma^{gt},0)\theta^m_{t_j}(D,0)\dfrac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\right)\end{array}\right)
\end{aligned} \tag{2.47}$$

$$\begin{aligned}
&W^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m) = \\
&\left(\frac{(1-\delta_{t_i}(1-\mu_m))(1-\delta_{t_i}(1-\beta_m))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_m))+\left(p_m\phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0)+(1-p_m)\phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0)\right)\delta_{t_i}\mu_m(1-\beta_m)\right]}\right) \\
&\cdot\left(\begin{array}{c} p_m(c+\tau)+(1-p_m)d \\ +\left(p_m\phi^m_{t_i}(D,0)\theta^m_{t_j}(\sigma^{gt},0)+(1-p_m)\phi^m_{t_i}(D,0)\theta^m_{t_j}(D,0)\right)\dfrac{\delta_{t_i}(1-\beta_m)d}{1-\delta_{t_i}(1-\beta_m)}\end{array}\right)
\end{aligned} \tag{2.48}$$

Upon receiving a new match, players will choose the stage game strategy which maximizes their expected value from the match by comparing these values given their expectations about the strategy followed by the match, both for the stage game strategy (p_m) and for their endogenous breakup decision ($\theta^m_{t_j}$), as well as their own optimal endogenous breakup decision in each case (ϕ_{t_i}). The expected value when unmatched is then

$$V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m) = \frac{\mu_m W^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\mu_m)} \tag{2.49}$$

$$V^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m) = \frac{\mu_m W^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\mu_m)} \tag{2.50}$$

To verify that the model presented here is identical to the model presented in the paper with exogenous breakup except for the endogenous breakup component, consider what happens here if the endogenous breakup decision were made for the players rather than by the players themselves. If we set all of the endogenous breakup terms equal to one to correspond to all players deciding to stay matched ($\phi_{t_i}\theta^1_{t_j} = 1$), all equations simplify to the analogous equations from the exogenous breakup case. If we set all the endogenous breakup terms equal to zero ($\phi_{t_i}\theta^1_{t_j} = 0$), such that endogenous breakup occurs for sure and matches only last for one period, then the only possible equilibrium is one without cooperation, as expected from a one-shot prisoner's dilemma.

2.7.2 Equilibrium with Endogenous Breakup

The addition of endogenous breakup to the model has significantly increased the complexity of the expressions that characterize the value expected by matched and unmatched players, but the method of determining equilibria has remained the same. First determine the endogenous breakup decision each player will make in each situation, taking as given all combinations of beliefs. Next, for each combination of beliefs and the associated optimal endogenous breakup decision in each situation, determine the value expected once matched from following each stage game strategy. Then given the value expected from a new match, determine if entering the market is indeed preferable to staying unmatched, and determine which market to enter in the two-market setting. Last, for any collection of optimal endogenous breakup, stage game strategy, and market decisions for which no player has a profitable deviation, verify that the beliefs that gave rise to these optimal decisions are consistent with the optimal choices of all players.

2.7.2.1 Endogenous Breakup Decision of the Currently Matched

Each player i type t_i determines his optimal endogenous breakup decision by comparing the value expected following the first period of a match, given by (2.41), (2.42), and (2.43), with the value expected from being unmatched, given by (2.49) and (2.50). There are three cases in terms of combinations of stage game strategies revealed by each player in the first period of a match and resulting updated beliefs. Since all players make individually rational decisions, each time a player is faced with the same situation they will make the same decision, allowing us to consider only symmetric strategies by type. Using this assumption that each player will find it optimal to follow the same strategy each time they are unmatched and matched, a player will choose to stay matched or not in each case as follows:

$$\phi^m_{t_i}(\sigma^{gt},1)=\begin{cases}1 \text{ if } W^m_{t_i}(\sigma^{gt}|1)=\dfrac{c+\delta_{t_i}\beta_m V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}\geq V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)\\ 0 \text{ if } W^m_{t_i}(\sigma^{gt}|1)=\dfrac{c+\delta_{t_i}\beta_m V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}< V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)\end{cases} \tag{2.51}$$

$$\phi^m_{t_i}(\sigma^{gt},0)=\begin{cases}1 \text{ if } W^m_{t_i}(\sigma^{gt}|0)=\dfrac{d+\delta_{t_i}\beta_m V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}\geq V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)\\ 0 \text{ if } W^m_{t_i}(\sigma^{gt}|0)=\dfrac{d+\delta_{t_i}\beta_m V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}< V^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m)\end{cases} \tag{2.52}$$

$$\phi^m_{t_i}(D,0)=\begin{cases}1 \text{ if } W^m_{t_i}(D|0)=\dfrac{d+\delta_{t_i}\beta_m V^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}\geq V^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)\\ 0 \text{ if } W^m_{t_i}(D|0)=\dfrac{d+\delta_{t_i}\beta_m V^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)}{1-\delta_{t_i}(1-\beta_m)}< V^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)\end{cases} \tag{2.53}$$

Each player uses equations (2.51), (2.52), and (2.53) to determine his optimal endogenous breakup decision in each case for all different combinations of beliefs about the stage game strategies, given by p_m, and about endogenous breakup decisions, given by $\theta^m_{t_j}$. Given the assumption that all other players also make individually rational decisions in the same manner, since all aspects of the decision process are commonly known, including the discount rates of each type, using these equations each player is able to assess what the optimal decisions of the other players will be, identifying $\theta^m_{t_j}$ for each type and combination of initial beliefs, $p_m \in [0,1]$, and stage game strategies, σ^{gt} and D.

2.7.2.2 Stage Game Strategy Decision of the Newly Matched

Having determined the optimal endogenous breakup decisions player i knows he will make in each situation, $\phi^m_{t_i}$, and the optimal endogenous breakup decisions expected from others, $\theta^m_{t_j}$, player i can determine the optimal stage game strategy to follow. Given these and taking as given beliefs about the likelihood a new match will cooperate, $p_m \in [0,1]$, each player i type t_i determines the optimal stage game strategy to follow by weighing the value expected from following each as given by (2.47) and (2.48). The stage game strategy followed by player i and the value expected from a new match in market $m \in \{s,f\}$ are thus

$$\arg\max_{\sigma^{gt},D} W^m_{t_i}(\cdot,\phi^m_{t_i}|\theta^m_{t_j},p_m) = \begin{cases} \sigma^{gt} & \text{if } W^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m) \geq W^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m) \\ D & \text{if } W^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m) < W^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m) \end{cases} \tag{2.54}$$

$$W^m_{t_i}(\cdot,\phi^m_{t_i}|\theta^m_{t_j},p_m) = \max\{W^m_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},p_m), W^m_{t_i}(D,\phi^m_{t_i}|\theta^m_{t_j},p_m)\} \tag{2.55}$$

2.7.2.3 Market Decision of the Unmatched

Having determined the optimal strategy to follow once matched and the associated value expected from a new match, $W^m_{t_i}(\cdot,\phi^m_{t_i}|\theta^m_{t_j},p_m)$, each player i determines the value expected from entering market $m \in \{s,f\}$ according to (2.49) and (2.50). Unmatched players then choose to enter the market that provides the highest expected value, according to (2.37). In order to nest the equations necessary for an equilibrium when there is one market in the equations necessary for an equilibrium when there are two markets, consider $\mu_s = 0$. If $\mu_s = 0$ then $V^s_{t_i} = 0$, making the value expected in the slow market equivalent to the outside option, making the market decision identical to the case where there is only one market, the fast market.

2.7.2.4 Equilibrium with Endogenous Breakup

We now have all the pieces necessary to define an equilibrium similar to the equilibrium with two markets given by Definition 2 in the exogenous breakup setting.

Definition 3 (Equilibrium with Endogenous Breakup)**.** *An equilibrium with endogenous breakup is, for $m \in \{s,f\}$, a list of values, $(V_{t_i}, V^m_{t_i}, W^m_{t_i})$, and beliefs, $(p_m, \theta^m_{t_j})$, such that, given market probabilities of matching and exogenous breakup, (μ_m, β_m)*

- *i (Endogenous Breakup) For all possible combinations of beliefs, $\theta^m_{t_j}$ and p_m, and stage game strategies, σ^{gt} and D, player i uses equations* (2.51), (2.52), *and* (2.53) *to determine his own optimal endogenous breakup decisions, $\phi^m_{t_i}$, and to update his expectations about the endogenous breakup decisions of all players $j \neq i$, $\theta^m_{t_j}$, $\forall i$*

ii (Matched) Given his own optimal breakup decisions, $\phi^m_{t_i}$*, and those expected by others,* $\theta^m_{t_j}$*, player i chooses strategy* σ^{gt} *or D for any beliefs* $p_m \in [0,1]$ *according to* (2.54) *that maximizes* $W^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)$ *as given by* (2.55)*,* $\forall i$

iii (Unmatched) Given values $W^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)$*,* $m \in \{s, f\}$*, expected once matched, player i forms expectations* $V^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)$ *about the value of entering each market according to* (2.49) *and* (2.50) *and chooses to enter market* $m \in \{s, f\}$ *or to stay unmatched to maximize* V_{t_i} *according to* (2.37)*,* $\forall i$

iv (Individual Rationality) Given $V^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)$ *and* $W^m_{t_i}(\cdot, \phi^m_{t_i}|\theta^m_{t_j}, p_m)$*,* $m \in \{s, f\}$*, player i does not have an incentive to deviate from his strategy determined by conditions i, ii, and iii,* $\forall i$

v (Consistency) $\forall i$*, player i's beliefs* p_m *about the stage game strategies followed by other players and beliefs* $\theta^m_{t_j}$ *about the endogenous breakup decisions of other players are consistent with the strategies followed by all players* $j \neq i$

2.7.3 First Best with Endogenous Breakup

In Section 2.4 we determined the general conditions in which the first-best outcome in which all players cooperate is feasible when all breakup is exogenous. In this section we will examine the first-best outcome in the setting with endogenous breakup. This will allow us to observe the tractability of the framework and solution method along with the increase in complexity that results from allowing for endogenous breakup. We will see the similarities between the two settings, as well as the one notable difference that results from new temptations created by the ability of players to end matches endogenously.

2.7.3.1 Endogenous Breakup Decision of the Currently Matched

Suppose that all players believe that all other players will cooperate, or that $p_f = 1$. Given these beliefs we need to determine whether players will stay matched or not, both on the equilibrium path and in response to a deviation.

On the equilibrium path: both players cooperate

In the first-best equilibrium, all players cooperate. Thus on the equilibrium path after the first period of stage game play, the value expected from staying matched is $W^m_{t_i}(\sigma^{gt}|1)$. Player i type t_i will stay matched according to (2.51) in this case if

$$W^f_{t_i}(\sigma^{gt}|1) = \frac{c + \delta_{t_i}\beta_m V^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, 1)}{1 - \delta_{t_i}(1 - \beta_m)} \geq V^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, 1) \tag{2.56}$$

Intuitively we expect to find that a player in this situation will always prefer to stay matched. By staying matched he will receive c each period until the match dissolves exogenously with probability β_f. By ending the match he expects to start receiving c each period once matched again, and is thus indifferent if $\mu_f = 1$ and strictly prefers staying matched $\forall \mu_f < 1$.

We can replace $V^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, 1)$ with (2.49), yielding

$$c \geq \frac{(1-\delta_{t_i})\mu_f}{1 - \delta_{t_i}(1-\mu_f)} W^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, 1)$$

Using (2.47) with $p_f = 1$ to substitute in for $W^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, 1)$, we get

$$c \geq \tfrac{(1-\delta_{t_i})\mu_f}{1-\delta_{t_i}(1-\mu_f)} \left(\tfrac{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_f)) + \phi^f_{t_i}(\sigma^{gt},1)\theta^f_{t_j}(\sigma^{gt},1)\delta_{t_i}\mu_f(1-\beta_f)\right]} \right)$$
$$\cdot \left(1 + \phi^f_{t_i}(\sigma^{gt}, 1)\theta^f_{t_j}(\sigma^{gt}, 1)\frac{\delta_{t_i}(1-\beta_f)}{1-\delta_{t_i}(1-\beta_f)} \right) c$$

Canceling like terms and combining the last expression into one fraction, we are left with

$$1 \geq \frac{\mu_f \left[1 - \delta_{t_i}(1-\beta_f) + \phi^f_{t_i}(\sigma^{gt}, 1)\theta^f_{t_j}(\sigma^{gt}, 1)\delta_{t_i}(1-\beta_f)\right]}{1 - \delta_{t_i}(1-\beta_f) + \phi^f_{t_i}(\sigma^{gt}, 1)\theta^f_{t_j}(\sigma^{gt}, 1)\delta_{t_i}\mu_f(1-\beta_f)}$$

Since the denominator is clearly positive, this becomes

$$\begin{pmatrix} 1 - \delta_{t_i}(1-\beta_f) \\ +\phi^f_{t_i}(\sigma^{gt}, 1)\theta^f_{t_j}(\sigma^{gt}, 1)\delta_{t_i}\mu_f(1-\beta_f) \end{pmatrix} \geq \begin{pmatrix} \mu_f(1 - \delta_{t_i}(1-\beta_f)) \\ +\mu_f\phi^f_{t_i}(\sigma^{gt}, 1)\theta^f_{t_j}(\sigma^{gt}, 1)\delta_{t_i}(1-\beta_f) \end{pmatrix}$$
$$1 - \delta_{t_i}(1-\beta_f) \geq \mu_f(1 - \delta_{t_i}(1-\beta_f))$$
$$1 \geq \mu_f$$

Recall that whenever (2.56) holds, it means that player i will choose to stay matched whenever he is in a cooperative relationship, or $\phi^f_{t_i}(\sigma^{gt}, 1) = 1$. Since $1 \geq \mu_f$

is always true, (2.56) always holds, and $\phi^f_{t_i}(\sigma^{gt},1) = 1$, $\forall i$. Since this is true for all players, the only beliefs consistent with individual rationality are $\theta^f_{t_j}(\sigma^{gt},1) = 1$. That is, all players will choose to stay in a cooperative match, and all players expect all other players to stay in a cooperative match. Endogenous breakup will not be exercised on the equilibrium path in the first-best equilibrium. That does not mean, however, that endogenous breakup does not alter incentives and change when the first-best is feasible relative to the exogenous breakup setting, as we shall see shortly.

Off the equilibrium path: response to defection

In the first-best equilibrium the only endogenous breakup decision players will actually make is the one just discussed, where we found that $\phi^f_{ti}(\sigma^{gt},1) = 1$. However, all players need to determine what they would do in response to a deviation. If player i cooperates but his matched plays D, will player i stay matched? Using (2.52), player i will stay matched if

$$W^f_{t_i}(\sigma^{gt}|0) = \frac{d + \delta_{t_i}\beta_m V^f_{t_i}(\sigma^{gt},\phi^f_{t_i}|\theta^f_{t_j},1)}{1-\delta_{t_i}(1-\beta_m)} \geq V^f_{t_i}(\sigma^{gt},\phi^f_{t_i}|\theta^f_{t_j},1) \qquad (2.57)$$

The left hand side is the value expected from staying in an uncooperative match after reverting to the grim trigger in response to the other player's defection. The right hand side is the value expected from ending the match and becoming unmatched. The equation providing this expected value is still (2.49) as it was above in (2.56) since in the first-best equilibrium, expectations are that the future match will cooperate, despite the deviation experienced in the current match, yielding

$$d \geq \frac{(1-\delta_{t_i})\mu_f}{1-\delta_{t_i}(1-\mu_f)} W^f_{t_i}(\sigma^{gt},\phi^f_{t_i}|\theta^f_{t_j},1) \qquad (2.58)$$

However, we can now substitute in for the optimal endogenous breakup decision we just found for play on the equilibrium path. Using (2.47) with $\phi^f_{t_i}(\sigma^{gt},1) = 1, \theta^f_{t_j}(\sigma^{gt},1)) = 1$ for $W^f_{t_i}(\sigma^{gt},\phi^m_{t_i}|\theta^m_{t_j},1)$ in (2.58) yields

$$\frac{d}{c} \geq \frac{\mu_f}{1-\delta_{t_i}(1-\beta_f)(1-\mu_f)}$$

$$(1-\delta_{t_i}(1-\beta_f)(1-\mu_f))\tfrac{d}{c} \geq \mu_f$$

$$(1-\delta_{t_i}(1-\beta_f))\tfrac{d}{c} \geq \mu_f - \mu_f\delta_{t_i}(1-\beta_f)\tfrac{d}{c}$$

$$\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0)) \equiv \frac{\left[1-\delta_{t_i}(1-\beta_f)\right]\frac{d}{c}}{\left[1-\delta_{t_i}\left(1-\beta_f\right)\frac{d}{c}\right]} \geq \mu_f \qquad (2.59)$$

Above when we examined the optimal endogenous breakup decisions on the equilibrium path we found that players will always choose to stay in a cooperative match. When defected against, however, players will only choose to stay matched if the probability of receiving a new match is lower than the bound $\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0))$ given by (2.59). The fraction $\frac{d}{c}$ that appears in $\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0))$ is the ratio of the stage game payoff in an uncooperative match over that in a cooperative match. The higher this ratio the higher the matching probability for which players will find it optimal to stay matched despite being defected against.

$\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0))$ is the cutoff matching probability below which player i will choose to stay matched and above which player i will choose to end the match, and it depends on δ_{t_i}. Thus unlike on the equilibrium path where all players of both types found it optimal to stay in a cooperative match, off the equilibrium path, whether player i will find it optimal to stay matched when defected against depends on his type. $\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0))$ is decreasing in δ_{t_i}, so the cutoff matching probability is higher for the low type and lower for the high type. It follows that if the matching probability is higher than the cutoff determined by the low type's discount rate, both types will choose to end the match when defected against, and if the matching probability is lower than the cutoff determined by the high type's discount rate, both types will choose to stay matched when defected against. Summarizing, we have that $\forall i, t_i \in \{h, \ell\}$

$$\phi^f_{t_i}(\sigma^{gt},0)) = \begin{cases} 0 \text{ if } \mu_f > \widetilde{\mu}_f(\phi^f_{\ell}(\sigma^{gt},0)) \equiv \dfrac{\left[1-\delta_\ell(1-\beta_f)\right]\frac{d}{c}}{\left[1-\delta_\ell\left(1-\beta_f\right)\frac{d}{c}\right]} \\ 1 \text{ if } \mu_f \leq \widetilde{\mu}_f(\phi^f_{h}(\sigma^{gt},0)) \equiv \dfrac{\left[1-\delta_h(1-\beta_f)\right]\frac{d}{c}}{\left[1-\delta_h\left(1-\beta_f\right)\frac{d}{c}\right]} \end{cases} \tag{2.60}$$

If the matching probability is in the range between the cutoffs, such that $\widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))$, then the high type will choose to end a match if defected against but the low type will choose to stay matched. In the first-best equilibrium all players follow the grim trigger strategy meaning that types are not revealed by stage game actions. Thus even though for any value of μ_f each player i can determine what endogenous breakup decision will be made by each type if he defects against them, since he cannot tell the type of a match his expectations about the endogenous breakup decision the match will make if he defects against them must be the population-weighted average over the types, or $\theta^f(\sigma^{gt},0) \equiv \pi\theta^f_h(\sigma^{gt},0) + (1-\pi)\theta^f_\ell(\sigma^{gt},0)) = \pi \cdot 0 + (1-\pi)\cdot 1 = 1-\pi$. Thus expectations about the endogenous breakup decision of a match who is defected against are

$$\theta^f(\sigma^{gt},0) = \begin{cases} 0 & \text{if} \quad \mu_f > \widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0)) \\ 1-\pi & \text{if} \quad \widetilde{\mu}_f(\phi_h^f(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0)) \\ 1 & \text{if} \quad \mu_f \leq \widetilde{\mu}_f(\phi_h^f(\sigma^{gt},0)) \end{cases} \tag{2.61}$$

Off the equilibrium path: defection

We also need to determine the optimal endogenous breakup decision made by a player who chooses to defect. Cooperating is only optimal, and thus the first-best is only feasible as an equilibrium, if the value expected from cooperating is at least as great as the value expected from defecting. In order to evaluate if defecting is a profitable deviation, player i must determine the value expected from defecting, which requires determining the endogenous breakup decision he will make if he defects. Since we are considering the first-best equilibrium, each player expects a new match to follow the grim trigger strategy and cooperate in the first period of a match. We just determined what the other player will do when defected against, which depends on his type and is given by (2.61). Now we must determine what the defecting player will do himself after defecting against his match. If the matching probability is low enough, he will want to stay matched and receive payoff d each period instead of ending the match and receiving utility 0 each period waiting for a new match. However, if the matching probability is high enough, he may wish to end the match with hopes of receiving a new match and the associated temptation payment $c+\tau$, which he expects with probability $p_f = 1$ once matched again, rather than staying matched and receiving d. Player i, after defecting in the first period of the match, is faced with the decision presented in (2.53), finding it optimal to stay matched if

$$W_{t_i}^f(D|0) = \frac{d + \delta_{t_i}\beta_f V_{t_i}^f(D, \phi_{t_i}^f|\theta_{t_j}^f, 1)}{1-\delta_{t_i}(1-\beta_f)} \geq V_{t_i}^f(D, \phi_{t_i}^f|\theta_{t_j}^f, 1) \tag{2.62}$$

The purpose of determining the optimal endogenous breakup decision after defecting is to evaluate the value expected from deviating from the equilibrium behavior of cooperating. This is reflected in $V_{t_i}^f(D, \phi_{t_i}^f|\theta_{t_j}^f, 1)$ on the right hand side, which assumes that if player i found it optimal to deviate in the first period of the current match, he will find it optimal to deviate in the first period of every match.[24]

[24]This assumes that all parameters are constant each period, including the discount rates of each type, the stage game utility payoffs, as well as the search probabilities, as has been assumed throughout.

This also means that the endogenous breakup decision facing player i will always be $\phi^f_{t_i}(D,0)$, and expectations about the endogenous breakup decision of others are given by $\theta^f(\sigma^{gt},0)$. Using this and replacing $V^f_{t_i}(D,\phi^f_{t_i}|\theta^f_{t_j},1)$ with (2.50), (2.62) becomes

$$d \geq \frac{(1-\delta_{t_i})\mu_f}{1-\delta_{t_i}(1-\mu_f)} W^f_{t_i}(D,\phi^f_{t_i}(D,0)|\theta^f(\sigma^{gt},0),1)$$

Replacing $W^f_{t_i}(D,\phi^f_{t_i}(D,0)|\theta^f(\sigma^{gt},0),1)$ with (2.48) yields

$$d \geq \frac{(1-\delta_{t_i})\mu_f}{1-\delta_{t_i}(1-\mu_f)}\left(\frac{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_f))+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}\mu_f(1-\beta_f)\right]}\right) \cdot \left(c+\tau+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\frac{\delta_{t_i}(1-\beta_f)d}{1-\delta_{t_i}(1-\beta_f)}\right)$$

Canceling like terms and combining the last expression into one fraction, we are left with

$$d \geq \frac{\mu_f\left[(1-\delta_{t_i}(1-\beta_f))(c+\tau)+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}(1-\beta_f)d\right]}{1-\delta_{t_i}(1-\beta_f)+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}\mu_f(1-\beta_f)} \tag{2.63}$$

We now need to consider the three cases for $\theta^f(\sigma^{gt},0)$, depending on the matching probability. For each case we can replace $\theta^f(\sigma^{gt},0)$ with the appropriate beliefs as given by (2.61) and see when (2.63) holds. When it does, $\phi^f_{t_i}(D,0)=1$, and when it does not, $\phi^f_{t_i}(D,0)=0$. Intuitively, we expect the player to stay matched whenever the payoff received with certainty from staying matched, d, is at least as great as the payoff expected next period if ending the matched, $\mu_f(c+\tau)$. This intuition proves to be correct in each case.

When the matching probability is high enough, both types will end the match if defected against. Using $\theta^f(\sigma^{gt},0)=0$ for this case when $\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))<\mu_f$, (2.63) becomes

$$d \geq \frac{\mu_f(1-\delta_{t_i}(1-\beta_f))(c+\tau)}{1-\delta_{t_i}(1-\beta_f)}$$

$$\frac{d}{c+\tau} \geq \mu_f$$

If $\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0)) < \mu_f \leq \frac{d}{c+\tau}$, then $\phi^f_{t_i}(D,0) = 1$. Note, however, that in this case $\theta^f(\sigma^{gt},0) = 0$ and thus the match will end because the other player who has been defected against will choose to end the match even though the defecting player wants to stay matched.

When the matching probability is low enough, both types will choose to stay matched even though they have been defected against. Using $\theta^f(\sigma^{gt},0) = 1$ for this case when $\mu_f \leq \widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0))$, (2.63) becomes

$$\begin{aligned}
d &\geq \frac{\mu_f\left[(1-\delta_{t_i}(1-\beta_f))(c+\tau) + \delta_{t_i}(1-\beta_f)d\right]}{1-\delta_{t_i}(1-\beta_f)(1-\mu_f)} \\
(1-\delta_{t_i}(1-\beta_f)(1-\mu_f))d &\geq \mu_f\left[(1-\delta_{t_i}(1-\beta_f))(c+\tau) + \delta_{t_i}(1-\beta_f)d\right] \\
(1-\delta_{t_i}(1-\beta_f))d &\geq \mu_f\left[(1-\delta_{t_i}(1-\beta_f))(c+\tau) + \delta_{t_i}(1-\beta_f)d\right] - \mu_f\delta_{t_i}(1-\beta_f)d \\
\frac{d}{c+\tau} &\geq \mu_f
\end{aligned}$$

If $\mu_f \leq \min\{\frac{d}{c+\tau}, \widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0))\}$, then $\phi^f_{t_i}(D,0) = 1$, and since $\theta^f(\sigma^{gt},0) = 1$ as well in this case, the match will stay together unless dissolved exogenously.

If the matching probability is in the middle range, $\widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))$, then whether the other player will end the match after being defected against depends on the other player's type. Only the low type, or a fraction $1-\pi$ of the population, will choose to stay matched. With $\theta^f(\sigma^{gt},0) = 1-\pi$, (2.63) becomes

$$\begin{aligned}
d &\geq \frac{\mu_f\left[(1-\delta_{t_i}(1-\beta_f))(c+\tau) + (1-\pi)\delta_{t_i}(1-\beta_f)d\right]}{1-\delta_{t_i}(1-\beta_f) + (1-\pi)\delta_{t_i}\mu_f(1-\beta_f)} \\
\left[1-\delta_{t_i}(1-\beta_f) + (1-\pi)\delta_{t_i}\mu_f(1-\beta_f)\right]d &\geq \begin{pmatrix} \mu_f(1-\delta_{t_i}(1-\beta_f))(c+\tau) \\ +\mu_f(1-\pi)\delta_{t_i}(1-\beta_f)d \end{pmatrix} \\
(1-\delta_{t_i}(1-\beta_f))d &\geq \mu_f(1-\delta_{t_i}(1-\beta_f))(c+\tau) \\
\frac{d}{c+\tau} &\geq \mu_f
\end{aligned}$$

If $\widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0)) < \mu_f \leq \min\{\frac{d}{c+\tau}, \widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))\}$, then $\phi^f_{t_i}(D,0) = 1$. Note, however, that the match will only stay together if both players choose to stay matched, which only occurs if the other player is the low type. For both the low and middle ranges for μ_f, the upper bound is determined by $\min\{\frac{d}{c+\tau}, \widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))\}$. Whether the first or second expression is lower, and thus the binding upper bound, will be determined by conditions on the discount rate that will be examined next in relation to the conditions required for cooperation to be optimal.

Note that in all three cases, the cutoff value for the matching probability is determined by $\frac{d}{c+\tau}$, which is the same as the upper bound on the matching probability in the slow market given by (2.27). The upper bound given by (2.27) was the highest the matching probability could be in the slow market for the low types to still find it optimal to enter the fast market. Any higher, and they found entering the slow market and waiting until matched to obtain the temptation payment worth the wait. The intuition is the same here, except instead of deciding between entering the slow and fast market, the decision is between staying matched or ending the match and entering the market again.

We have now specified the optimal endogenous breakup decision each player will make in each case, as well as expectations about the optimal decisions others will make as required by condition i of Definition 3 for an equilibrium with endogenous breakup.

2.7.3.2 Stage Game Strategy Decision of the Newly Matched

We have seen in the first-best equilibrium how each player will make endogenous breakup decisions in each situation of a continued match, as well as the expectations each player has about the decision each other player will make, both on and off the equilibrium path. We now need to consider under what conditions players will find it optimal to cooperate given these expectations about endogenous breakup behavior. After receiving a new match, player i weighs the value expected from following the grim trigger strategy with that expected from defecting as in (2.54). Given beliefs $p_f = 1$ and using (2.47) and (2.48) for $W^f_{t_i}(\sigma^{gt}, \phi^f_{t_i}|\theta^f_{t_j}, p_f)$ and $W^f_{t_i}(D, \phi^f_{t_i}|\theta^f_{t_j}, p_f)$, respectively, player i finds it optimal to cooperate if

$$\left(\frac{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_f))+\phi^f_{t_i}(\sigma^{gt},1)\theta^f_{t_j}(\sigma^{gt},1)\delta_{t_i}\mu_f(1-\beta_f)\right]}\right)\left(1+\phi^f_{t_i}(\sigma^{gt},1)\theta^f_{t_j}(\sigma^{gt},1)\frac{\delta_{t_i}(1-\beta_f)}{1-\delta_{t_i}(1-\beta_f)}\right)c$$
$$\geq$$
$$\left(\frac{(1-\delta_{t_i}(1-\mu_f))(1-\delta_{t_i}(1-\beta_f))}{(1-\delta_{t_i})\left[(1-\delta_{t_i}(1-\beta_f))+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}\mu_f(1-\beta_f)\right]}\right)\left(c+\tau+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\frac{\delta_{t_i}(1-\beta_f)d}{1-\delta_{t_i}(1-\beta_f)}\right)$$

When we solved for the condition on the discount rate required for players to find cooperation optimal in the exogenous breakup setting with one market, the large multiplier terms in the equivalent expression were identical, and thus canceled out of the inequality. Here, the endogenous breakup terms are different on each side, so solving this for a condition on the discount rate is not as simple. We found above that $\phi^f_{t_i}(\sigma^{gt},1)\theta^f_{t_j}(\sigma^{gt},1) = 1$, while $\phi^f_{t_i}(D,0)$ and $\theta^f_{t_j}(\sigma^{gt},0)$ depend on the matching probability. Substituting $\phi^f_{t_i}(\sigma^{gt},1)\theta^f_{t_j}(\sigma^{gt},1) = 1$, canceling like terms, combining the second terms on each side, and canceling that denominator as well, this becomes

$$\frac{c}{1-\delta_{t_i}(1-\beta_f)+\delta_{t_i}\mu_f(1-\beta_f)} \geq \frac{(1-\delta_{t_i}(1-\beta_f))(c+\tau)+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}(1-\beta_f)d}{1-\delta_{t_i}(1-\beta_f)+\phi^f_{t_i}(D,0)\theta^f_{t_j}(\sigma^{gt},0)\delta_{t_i}\mu_f(1-\beta_f)} \quad (2.64)$$

There are several cases for the endogenous breakup decisions of defecting and defected against players that depend on the matching probability. One possibility of how to proceed is to solve for the matching probability and compare the results with the ranges of μ_f that imply the values of $\phi^f_{t_i}(D,0)$ and $\theta^f_{t_j}(\sigma^{gt},0)$ as given by (2.60) and (2.61). While this method does lead to the same results, the analysis is more transparent if we instead consider the cases for μ_f separately and derive the conditions required for cooperation to be optimal in each case.

Highest matching probability: $\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0)) < \mu_f$

In the case when $\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0)) < \mu_f$, we found that the player who has been defected against will end the match, or $\phi^f_{t_i}(\sigma^{gt},0) = 0$ for each type, and thus the defecting player expects $\theta^f(\sigma^{gt},0) = 0$. Because $\theta^f(\sigma^{gt},0)\phi^f_{t_i}(D,0) = 0$, $\forall \phi^f_{t_i}(D,0)$, the match will end regardless of whether the defecting player wishes to stay matched or not. Using $\theta^f(\sigma^{gt},0)\phi^f_{t_i}(D,0) = 0$ in (2.64), we find that player i, with discount rate δ_{t_i}, finds cooperating rather than defecting optimal if

$$\mu_f \leq 1 - \frac{\tau}{\delta_{t_i}(1-\beta_f)(c+\tau)} \equiv \overline{\mu}_f(t_i) \quad \text{where we define} \quad \overline{\mu}_f \equiv \overline{\mu}_f(\ell) \quad (2.65)$$

Since this upper bound is lower for lower δ_{t_i}, it is always binding based on the discount rate of the low type, making it useful to define $\overline{\mu}_f \equiv \overline{\mu}_f(\ell)$. This upper bound on the matching probability is the key difference between when the first-best is feasible with endogenous breakup compared to without endogenous breakup, and in fact, the only major difference between the incentives facing players in the settings with endogenous and exogenous breakup.[25] When players have the ability to end the

[25]The other main difference between the two settings occurs in the one-market separating equilib-

match, there is an upper bound on the matching probability above which defecting becomes a profitable deviation because a new match is expected too soon. Without endogenous breakup, players had to wait for the match to end exogenously, which made the prospects of being in a cooperative relationship more attractive. Now a player can defect, receive the temptation payment, and then just end the match so that he can receive the temptation payment again against a new match. The higher the matching probability, the sooner he expects to be able to be matched and receive the temptation payment again, and the less attractive is cooperating.

Just as was the case when all breakup was exogenous, with endogenous breakup the first-best is only feasible if players are patient enough. The upper bound on the matching probability given by (2.65) can be rearranged into a condition on the discount rate. For a given matching probability, player i finds cooperation optimal only if

$$\delta_{t_i} \geq \frac{\tau}{(1-\mu_f)(1-\beta_f)(c+\tau)} \equiv \underline{\delta}(\mu_f) \tag{2.66}$$

There are two differences between the lower bound on the discount rate required for the first-best to be feasible with endogenous breakup given by $\underline{\delta}(\mu_f)$ and the lower bound $\underline{\delta}_1^f \equiv \frac{\tau}{(1-\beta_f)(c-d+\tau)}$ found in Section 2.4 in the exogenous breakup setting. The first is the $(1-\mu_f)$ term in the denominator, which shows that the greater the likelihood of receiving a new match the more patient a player must be to find defection to not be a profitable deviation. In the exogenous breakup case the case of certain matching with $\mu_f = 1$ was often examined, but with endogenous breakup it can easily be seen from (2.66) that cooperation cannot be sustained between all players with certain matching when players can end a match.

The second difference between $\underline{\delta}(\mu_f)$ and $\underline{\delta}_1^f$ is the $(c+\tau)$ term which reflects the temptation payment only, and does not include the difference between a cooperative and uncooperative match, $(c-d+\tau)$, as is found in $\underline{\delta}_1^f$. The value of an uncooperative match does not enter into the lower bound on the discount rate given by (2.66) because the players do not stay in an uncooperative match in this case. The condition from the exogenous breakup setting, $\underline{\delta}_1^f$, still plays an active role here as well. Recall that the case we are currently considering is the case where $\widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0)) < \mu_f$, and in

rium. With endogenous breakup, it is necessary to examine explicitly the steady-state probability of being matched with a high type in the one-market separating equilibrium, as discussed in Section 2.6. However, this steady-state analysis is technical in nature, changing what beliefs p_m are equal to in equilibrium, rather than changing the incentives facing the players.

order for the first-best to be feasible in this case requires that $\mu_f \leq \overline{\mu}_f$. In order for $\widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt}, 0)) < \mu_f \leq \overline{\mu}_f$ to be a valid case clearly requires that $\widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt}, 0)) < \overline{\mu}_f$, which occurs whenever $\delta_{t_i} \geq \underline{\delta}_1^f$. Thus in order for the first-best to be feasible in this case requires both $\delta_{t_i} \geq \underline{\delta}_1^f$ and $\delta_{t_i} \geq \underline{\delta}(\mu_f)$.

Which lower bound is binding, $\underline{\delta}_1^f$ or $\underline{\delta}(\mu_f)$, depends on how the matching probability affects the incentive to deviate, which depends on where the matching probability is in the range $\mu_f \in [\widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt}, 0)), \overline{\mu}_f]$. If the matching probability is low enough that a defecting player would choose to stay matched, then the lower bound on the discount rate required for the player to not find deviation profitable is the same as in the exogenous breakup setting since he does not choose to exercise endogenous breakup. If the matching probability is high enough that a defecting player finds it optimal to end the match so that he can defect again, then the lower bound on the discount rate also includes a term involving the matching probability, requiring that players be more patient the higher the matching probability. The borderline case is $\mu_f = \frac{d}{c+\tau}$, which is the cutoff for the matching probability below which the defecting player wishes to stay matched. Specifically, if $\frac{d}{c+\tau} < \mu_f$, then $\underline{\delta}(\mu_f) < \underline{\delta}_1^f$, and if $\mu_f < \frac{d}{c+\tau}$ then $\underline{\delta}_1^f < \underline{\delta}(\mu_f)$. So while the endogenous breakup decision of the defecting player does not have an actual effect on the longevity of the match in this case because the player who is defected against chooses to end the match, it does have an effect on which condition on the discount rate is binding.

Lowest matching probability: $\boldsymbol{\mu_f \leq \widetilde{\mu}_f(\phi_h^f(\sigma^{gt}, 0))}$

In the case when $\mu_f \leq \widetilde{\mu}_f(\phi_h^f(\sigma^{gt}, 0))$, we found that the matching probability is low enough that players who are defected against will choose to stay matched since the expected wait before receiving a new match is too long, so $\phi_{t_i}^f(\sigma^{gt}, 0) = 1$ for each type, and thus the defecting player expects $\theta^f(\sigma^{gt}, 0) = 1$. As a result, it is now the endogenous breakup decision of the defecting player that determines if the match stays together or not. We found that the defecting player chooses to stay matched whenever the value he knows he will receive if he stays matched, d, is as great as the value expected from ending the match, $\mu_f(c + \tau)$. So the defecting player chooses to stay matched whenever $\mu_f \leq \frac{d}{c+\tau}$. It is easy to show that $\widetilde{\mu}_f(\phi_h^f(\sigma^{gt}, 0)) < \frac{d}{c+\tau}$ whenever $\underline{\delta}_1^f < \delta_h$. Assuming for a moment that this is the case and thus that $\phi_{t_i}^f(D, 0)) = 1$, we can use $\phi_{t_i}^f(D, 0)\theta^f(\sigma^{gt}, 0) = 1$ in (2.64) to find the condition required for cooperation to be optimal, which is that $\underline{\delta}_1^f < \delta_{t_i}$. The assumption that $\underline{\delta}_1^f < \delta_h$ is confirmed, and thus $\phi_{t_i}^f(D, 0)) = 1$ is the endogenous breakup decision of the defector. So in this case with $\mu_f \leq \widetilde{\mu}_f(\phi_h^f(\sigma^{gt}, 0))$, the condition required for the first-best to be feasible is $\underline{\delta}_1^f < \delta_{t_i}$, the same as is required without endogenous breakup. The condition required is the same in this case because no player chooses to exercise endogenous breakup when the matching probability is this low.

Middle matching probability: $\boldsymbol{\widetilde{\mu}_f(\phi_h^f(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0))}$

When $\widetilde{\mu}_f(\phi_h^f(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0))$, the high type find it optimal to end the match if defected against ($\phi_h^f(\sigma^{gt},0)=0$) but the low type find it optimal to stay matched($\phi_\ell^f(\sigma^{gt},0)=1$). Thus the defecting player has expectations $\theta^f(\sigma^{gt},0)=1-\pi$ as shown above. It is easy to show that whenever $\underline{\delta}_1^f < \delta_\ell$, then $\widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0)) < \frac{d}{c+\tau}$. It remains the case that $\phi_{t_i}^f(D,0) = 1$ if $\mu_f \leq \frac{d}{c+\tau}$. Thus whenever $\underline{\delta}_1^f < \delta_\ell$, the defecting player will find it optimal to stay matched. Assuming for the moment that this is the case, we can evaluate (2.64) using $\phi_{t_i}^f(D,0)\theta^f(\sigma^{gt},0) = 1-\pi$, which yields the condition that cooperation is optimal if $\underline{\delta}_1^f < \delta_{t_i}$. Because $\underline{\delta}_1^f < \delta_\ell$ is required for cooperation to be optimal, the assumption that $\mu_f \leq \frac{d}{c+\tau}$ and thus that $\phi_{t_i}^f(D,0)=1$ is in fact correct. Thus the first-best is feasible in this middle case with $\widetilde{\mu}_f(\phi_h^f(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi_\ell^f(\sigma^{gt},0))$ whenever $\underline{\delta}_1^f < \delta_{t_i}$.

Summary for all cases

We have found the conditions under which players find it optimal to cooperate, given the optimal endogenous breakup decisions from condition *i* of Definition 3. This completes the specification of condition *ii* required for an equilibrium with endogenous breakup. We know that on the equilibrium path all players will follow the grim trigger strategy and choose to stay matched as long as the required conditions are satisfied. The conditions required for the first-best to be feasible in each range of the matching probability, as well as the off-equilibrium endogenous breakup decisions in each case, are summarized in the following table.

2.7.3.3 Completing the Specification of the First Best Equilibrium with Endogenous Breakup

Condition *iii* of Definition 3 requires that we specify the decision of unmatched players. Because we are considering the first-best equilibrium in which all players cooperate in one market, the fast market, this condition is satisfied trivially in this case. The endogenous breakup decisions were derived such that each player is taking the action that maximizes his expected value, and each player's expectations about the decisions of others assume the same individually rational choices of others. The

If the matching probability is	1^{st} best feasible if	Endogenous Breakup Dec.
$\overline{\mu}_f < \mu_f$	Not feasible $\forall \delta_{t_i}$	$\phi^f_{t_i}(\sigma^{gt},0)=0, \phi^f_{t_i}(D,0)=0$
$\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0)) < \frac{d}{c+\tau} < \mu_f \leq \overline{\mu}_f$	$\underline{\delta}(\mu_f) \leq \delta_{t_i}$	$\phi^f_{t_i}(\sigma^{gt},0)=0, \phi^f_{t_i}(D,0)=0$
$\widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0)) < \mu_f \leq \frac{d}{c+\tau} < \overline{\mu}_f$	$\underline{\delta}^f_1 \leq \delta_{t_i}$	$\phi^f_{t_i}(\sigma^{gt},0)=0, \phi^f_{t_i}(D,0)=1$
$\widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0)) < \mu_f \leq \widetilde{\mu}_f(\phi^f_\ell(\sigma^{gt},0))$	$\underline{\delta}^f_1 \leq \delta_{t_i}$	$\phi^f_h(\sigma^{gt},0)=0, \phi^f_\ell(\sigma^{gt},0)=1$, $\phi^f_{t_i}(D,0)=1$
$\mu_f \leq \widetilde{\mu}_f(\phi^f_h(\sigma^{gt},0))$	$\underline{\delta}^f_1 \leq \delta_{t_i}$	$\phi^f_{t_i}(\sigma^{gt},0)=1, \phi^f_{t_i}(D,0)=1$

$\overline{\mu}_f \equiv 1 - \frac{\tau}{\delta_\ell(1-\beta_f)(c+\tau)}$, $\widetilde{\mu}_f(\phi^f_{t_i}(\sigma^{gt},0)) \equiv \frac{\left[1-\delta_{t_i}(1-\beta_f)\right]\frac{d}{c}}{\left[1-\delta_{t_i}(1-\beta_f)\frac{d}{c}\right]}$, $\underline{\delta}(\mu_f) \equiv \frac{\tau}{(1-\mu_f)(1-\beta_f)(c+\tau)}$, $\underline{\delta}^f_1 \equiv \frac{\tau}{(1-\beta_f)(c-d+\tau)}$

Table 2.3: Summary of results for first-best equilibrium with endogenous breakup

conditions we derived above required for cooperation to be optimal were found by weighing the value a player expects from cooperating against the value he expects from defecting. As a result, no player has a profitable deviation when the conditions are met, satisfying condition *iv* of Definition 3. All of the preceding was derived starting from the assumption that all players cooperate. When the conditions are satisfied all players cooperate, the probability of being matched with a cooperator is 1, and thus the beliefs $p_f = 1$ are consistent. This, together with the fact that the endogenous breakup decisions other players will make, $\theta^f_{t_j}$, $\forall j \neq i$, are consistent with the decision each player will make, $\phi^f_{t_i}$, $\forall i$, satisfies condition *v*. Thus all the conditions of Definition 3 are satisfied and an equilibrium with endogenous breakup in which all players cooperate exists. The first-best outcome is feasible with endogenous breakup for any combination of parameters satisfying the conditions laid out above. See Section 2.6.1 for examples and a discussion of the extent to which welfare is decreased by allowing for endogenous breakup.

2.7.4 Separating Equilibrium with Endogenous Breakup

We went through the process of deriving the conditions required for the first-best equilibrium to exist in great detail, allowing us to see the similarities and differences between the settings with exogenous and endogenous breakup. We saw the complicated yet straightforward process this entails and developed much of the intuition that is to be gained about the incentives facing the players in the endogenous breakup framework. Little is gained in terms of further intuition for the endogenous breakup setting by working through the derivations of the conditions required for existence of both the separating equilibrium with one and two markets with endogenous breakup. As a result the key features of each of the separating equilibria, as well as the intuition for the endogenous breakup decision in each, will be discussed but the formal analysis will not be presented.

2.7.4.1 Separating Equilibrium with One Market

A separating equilibrium with one market that meets all the conditions of an equilibrium with endogenous breakup as in Definition 3 exists under reasonable circumstances, as was the case in the exogenous breakup setting. Just as in the exogenous breakup setting, a separating equilibrium with one market[26] is an equilibrium in which all players who are the high type cooperate and all players who are the low type defect.

In the first-best equilibrium, no players will choose to end a match on the equilibrium path because on the equilibrium path all matches are cooperative and defecting is not a profitable deviation so there is no reason to end a match. In the separating equilibrium with one market, it remains the case that players will not find it optimal to end a cooperative match, which in this case is limited to players who are the high type who are matched with another high type.

There are, however, several cases in which players may find it optimal to end a match on the equilibrium path in the separating equilibrium with one market. One is high types who are in an uncooperative match. If the matching probability is high enough, a high type finds it optimal to end the match with hopes of receiving a new match who might also be a high type and cooperate. If the matching probability is too low, however, he finds it optimal to stay in the current uncooperative match because the hopes of receiving a new match that may be cooperative is not worth the expected wait. The other source of endogenous breakup on the equilibrium path is low types who hope to receive repeated temptation payments. If the matching probability is high enough, the low type may choose to end a match with hopes of being matched again soon with a high type and receiving the associated temptation payment. If the matching probability is low enough, however, he will find it optimal to stay matched rather than waiting for a new match.

Unlike in the first-best equilibrium with endogenous breakup, it is possible for the separating equilibrium with one market and endogenous breakup to exist with

[26]We will continue to assume that this one-market equilibrium occurs in the fast market as was done when we considered the first-best equilibrium.

certain matching. In the first-best equilibrium, if the matching probability was too high, the temptation to defect and end the match to defect again was too great. In the separating equilibrium, this temptation is reduced by the fact that defection does not result in the temptation payment with certainty, but rather only if matched with a high type. Thus the incentive for high types to end the match after defecting is not as strong, so cooperation by the high types can be sustained at higher matching probabilities than in the first-best with endogenous breakup, including the case of $\mu_f = 1$. However, the temptation created by a high matching probability now that players can end matches is still present. There is a trade-off between the level of temptation resulting from the value of τ and that resulting from a higher matching probability. If the matching probability is higher, the values of τ for which the high type still find it optimal to cooperate are lower. For example, the separating equilibrium with one market exists at values of the temptation payment ten times higher with $\mu_f = 0.3$ than with $\mu_f = 1$, showing the importance of expected wait times in facilitating cooperation.

For the separating equilibrium with endogenous breakup, it is necessary to consider explicitly the steady-state probability that a new match will be the high or low type. This is discussed in detail in Section 2.6.2.

2.7.4.2 Separating Equilibrium with Two Markets

A separating equilibrium with two markets meeting all the conditions of an equilibrium with endogenous breakup as in Definition 3 also exists under reasonable circumstances. Just as in the exogenous breakup case, in this equilibrium all high types enter the slow market and cooperate while all low types enter the fast market and defect. Just as in the first-best equilibrium, endogenous breakup is never exercised on the equilibrium path. As a result, many of the conditions characterizing when the separating equilibrium with two markets exists are identical to those in the exogenous breakup case. The conditions required for the high type to find it optimal to enter the slow market and for the low type to find it optimal to enter the fast market are still characterized by (2.22), (2.24), and (2.27).

The only real difference between the endogenous and exogenous breakup settings for the two-market separating equilibrium comes from the endogenous breakup decision of the high type if they are defected against in the slow market. For these players the endogenous breakup decision if defected against follows the now familiar pattern of stay matched if the matching probability, now μ_s, is low enough and end the match if it is high enough. This threat to exercise endogenous breakup if defected against changes the incentives for defecting, changing the conditions required for cooperation to be optimal for the high types in the slow market. As a result, (2.21) no longer characterizes the lower bound on the discount rate of the high type, and (2.23) no longer characterizes the discount rate cutoff above which the low type would cooperate if they were to enter the slow market. As was the case in separating equilibrium with one market, there is a trade-off for sustaining cooperation among the high types between τ and μ_s just as there was between τ and μ_f in the one-market separating equilibrium. However, for the two-market separating equilibrium the point at which this trade-off has bite is never realized because before that point is reached, the slow market has become too tempting for the low type, causing them to defect and enter the slow market.

Also of note is that it is possible to sustain the separating equilibrium with two markets with certain matching in the fast market, as was the case considered with exogenous breakup. Thus the only equilibrium with endogenous breakup in which $\mu_f = 1$ cannot occur is the first-best equilibrium.

2.7.5 Optimality of Equilibria with Endogenous Breakup

We have considered three equilibria in this setting with endogenous breakup. Just as was the case in the exogenous breakup setting, when the first-best is achievable, all players prefer this equilibrium to all others. When the first-best is not feasible, the separating equilibria with one and two markets, when they exist, are preferred by all players to repetition of the stage game Nash equilibrium. Given that the formal conditions characterizing the separating equilibria have been omitted, we will discuss the general patterns that emerge rather than formulating exact conditions that characterize the welfare properties of the separating equilibria as was done with Propositions 7 and 8 in the exogenous breakup setting.

In general, if the low type is significantly less patient than the high type, then both types benefit from the one-market separating equilibrium. If the types are close together in terms of their level of patience but the temptation to defect is too great to sustain the first-best equilibrium, then both types can benefit from the two-market separating equilibrium. The welfare implications are the same as they were in the setting with exogenous breakup. If neither the first-best equilibrium or the separating equilibrium with one market are feasible, then the separating equilibrium with two markets is preferred to all other equilibria feasible in this framework and with the stage game strategies of grim trigger or always defect. When the first-best is not feasible but both separating equilibria exist, the low type always prefer the one-market separating equilibrium while the high type sometimes also prefer the one-market separating equilibrium and in other circumstances prefer the separating equilibrium in two markets. The qualitative results about how separation can improve welfare found with fully exogenous breakup remain the same, even when matches can be ended endogenously.

2.7.6 Hybrid Model: Between Fully Exogenous and Fully Endogenous Breakup

For the hybrid model discussed in Section 2.6.3, simply replace every instance of $\phi^m_{t_i}\theta^m_{t_i}$ with $\rho^m_{t_i}\phi^m_{t_i}\theta^m_{t_i}$, where $\rho^m_{t_i} \in [0, 1]$ is the probability that an agent of type t_i is allowed to end a match in market m endogenously.[27] If $\rho^m_{t_i} = 1$, the model is identical to the model with fully endogenous breakup, and if $\rho^m_{t_i} = 0$, the model reduces to the model with fully exogenous breakup.

2.8 Conclusion

In Section 2.4 we examined a repeated prisoner's dilemma embedded in a search model with one market and considered conditions in which some or all players find it optimal to cooperate. While the original motivation for the paper was to determine when sorting can improve welfare when the first-best outcome is not feasible, the framework employed to address this question revealed a setting with several appealing properties for sustaining cooperation between all players.

First, cooperation can be sustained with minimal cognitive requirements on the players. The strategies followed by players in equilibrium are simple and only require remembering whether the current match has defected. Players simply determine if they find cooperation optimal and only punish players who defect against them directly, not requiring the desire to implement punishments against those who have not defected against them or the common understanding required for coordinated

[27]Everything presented in the paper has been the same for both types, with the exception of δ_{t_i}, but that is not required. Stage game payoffs can also be extended to differ by types, as discussed in footnote 3 in the paper. The possibility of ending matches endogenously in the hybrid model can also depend on type.

punishments. As a result, cooperation can be sustained with minimal informational requirements. There is no need for public randomization devices or signals and no need to identify individuals or keep track of histories beyond the scope of an active match. In such a setting, it would be reasonable to conclude that cooperation must require that matches be near-infinite in duration. Surprisingly, the probability of breakup each period can be relatively high and still sustain cooperation between players of both types.

When the first-best outcome is not achievable, players of both types benefit from sorting. All players prefer a separating equilibrium with one market, when it exists, to repetition of the stage game Nash equilibrium, and in some circumstances all prefer it to an equilibrium with full separation. These results are sensitive to the fraction of the population who are each type. Welfare gains for each type are greater the higher the fraction of the population who are the high type, and disappear when the fraction becomes too low.

In Section 2.5, we found that even if conditions are such that the first-best outcome and separation in one market are not feasible, a Pareto improvement opportunity may still exist if a second market is available to allow for full separation of types. This increase in the sustainability of cooperation that resulted from the introduction of a second market did not require changing the payoffs of the stage game. Overall utilities are changed as a result of the time cost imposed by the lower matching probability in the slow market. However, this type of change in utility is more desirable than altering stage game payoffs to facilitate cooperation. The cost incurred in the slow market is entirely born by players choosing to enter that market, something that is done entirely voluntarily and need not be done at all. If instead stage game payoffs are changed, the resources used to facilitate cooperation and insure against loss must come from somewhere.

The finding that introduction of a second market can provide a Pareto improvement opportunity has two immediate consequences for sustaining cooperation in a repeated prisoner's dilemma in a heterogeneous population. First, cooperation can be sustained for the more patient fraction of the population at significantly higher levels of temptation and potential loss than is feasible when full separation is not possible. Second, these welfare gains expected for the more patient players are no longer sensitive to the fraction of the population who are patient as was the case when full separation was not possible.

An implication of these findings is that facilitating separation of types in situations in which there exist prisoner's dilemma type tensions between benefits from cooperation and individual incentives has the potential to improve the welfare of the more patient fraction of the population without decreasing the welfare of others. When considering situations in which allowing for greater separation may provide Pareto improvement opportunities, we should consider separation opportunities consistent with the features of this model, such as ones with minimal informational requirements and costs that are only incurred by willing participants rather than requiring outside resources or transfers from unwilling participants. Time is a resource possessed by all, but that is more valuable to some than others. As explored here, time is one potential source of separation opportunities that can be leveraged to create Pareto-optimal outcomes. A necessary condition is just for some fraction of the population to find cooperation worth the wait.

Chapter 3

Effects of Spectrum Holdings on Equilibrium in the Wireless Market

3.1 Introduction

Providers of mobile wireless communication services (e.g., AT&T, Sprint Nextel, T-Mobile, Verizon Wireless) compete for customers on the basis of both price and service quality. A firm with higher service quality is able to attract more customers and set a higher price for its services. Network congestion plays a central role in determining a firm's service quality. By increasing its capacity, a firm is able to reduce network congestion and offer a higher level of service quality. Spectrum is an important determinant of capacity for a wireless provider. The effects of spectrum holdings on equilibrium outcomes is the focus of this paper.

A firm's spectrum holdings, commonly measured in megahertz (MHz), refer to what radio wave frequencies a firm has the right to use in order to transmit signals to and from customers' mobile devices to network infrastructure such as cell towers. The more MHz of spectrum a firm has the right to use to transmit signals, the more customers a firm is able to serve, and the higher the quality of service.

Various factors in addition to a firm's total spectrum holdings also affect service quality. For example, technological advances allow for more efficient use of existing spectrum. In addition, spectrum at different frequencies possesses different properties. These different properties affect the ability of signals sent at different frequencies to travel through different climates and topographies, over long distances, and to penetrate buildings and automobiles. However, without spectrum, a firm is unable to provide service, and whatever the technology employed, and whatever the frequency, more spectrum translates into greater bandwidth for transferring signals, and in turn, into higher capacity and service quality.

Spectrum is also interesting from a policy perspective. First, spectrum is a finite resource. The range of frequencies suitable for transmitting mobile wireless communications is limited. Thus, it is not possible for every firm to use as much spectrum as it might like. Consequently, various policies limit access to spectrum. In the United States, the Federal Communications Commission (FCC) controls the use of spectrum. The FCC affects the allocation of spectrum by regulating its use, by creating and implementing the mechanisms by which rights to its use are transfered, and through explicit regulations on spectrum accumulation.

The FCC designates different frequencies for different uses, such as radio and television transmission, radar, direct communication devices such as walkie-talkies, as well as mobile telecommunications, including talk, text, and data transmissions. From time to time, the FCC decides to change the use of certain frequencies. For a long time, frequencies between 698 and 806 MHz, also know as the "700 MHz band," were used for transmission of UHF television signals (channels 52 to 69). In the mid 2000s, broadcasters were required to transition from analog to digital transmissions, freeing up the 700 MHz band for other uses. The FCC auctioned rights to use spectrum in the 700 MHz band, most of which is now used for mobile wireless communications. The details of the allocation process are determined by the FCC and may or may not be optimal, for firms, customers, or society. Thus an examination of how different allocations of spectrum affect equilibrium in the wireless market is of interest.

In addition to designing the methods by which spectrum is made newly available for use in mobile wireless communications, the FCC also regulates the sale, lease, and other transfers of spectrum between firms and other parties. For example, a firm that possesses the right to use specific frequencies in a specific geographic region that wants to sell or lease the rights to this spectrum to another party must obtain the permission of the FCC. The FCC generally evaluates the effects of this spectrum transfer before giving its approval or not. As a general rule, the FCC places limitations on the amount of spectrum a firm may accumulate. These regulations affect how concentrated or equally distributed spectrum holdings are among firms.

These policies that regulate the transfer of existing spectrum, as well as the design of the mechanisms used for making spectrum newly available for use in mobile wireless communications, affect how spectrum is allocated among firms. A theoretical model of the industry, based on the model introduced by Pinto and Sibley (2013), will be calibrated to data from the wireless industry, and the effects of different allocations of spectrum will be simulated in order to evaluate how these allocations affect equilibrium outcomes including prices, output, and welfare.

Pinto and Sibley (2013) introduce an explicit formulation of network congestion into a standard oligopoly model of price competition between firms offering partially differentiated products. The formulation facilitates direct consideration of how capacity and service quality affect equilibrium output and prices. They show that in a linear model, if one firm, firm A, increases its spectrum holdings, its quality increases and it raises its price. Firm B cuts its price in response. However, the effects on firm B's quality in equilibrium are unclear. Ignoring equilibrium price effects for a moment, if firm A's congestion decreases, thereby increasing its service quality, some of firm B's customers will leave for firm A. Other things equal, this raises quality at firm B, but both B and A change their equilibrium prices as well. Since A raises its price and B lowers its price, both types of price changes work to keep B's customers from moving to firm A. Hence, the overall effect of the increase in firm A's capacity on firm B's service quality is unclear. However, there is the potential for an externality effect that could deter capacity increases by firm A.

In this paper a model based on that of Pinto and Sibley (2013) is calibrated to data from the wireless industry in order to examine the importance of these externalities. Simulations under three counterfactual allocations of spectrum holdings are conducted. In the first, one firm is given additional spectrum, holding constant the spectrum holdings of the other firms. The firm that receives the additional spectrum increases its service quality and raises its price. The other firms lower their prices in response, but lose some customers. Consequently, the service quality of the other firms increases, as they now serve fewer customers with the same capacity. The increase in service quality is highest for the firms whose spectrum remains constant when the firm that increases its spectrum holdings is a larger firm because the changes in quantities are larger in this case, compared to when it is a smaller firm that receives the additional spectrum.

The second scenario considered is the equalization of spectrum holdings among all firms. When an equal share of total spectrum is given to each firm, the firms that lose spectrum lose customers to the firms that gain spectrum. These changes to equilibrium quantities are accompanied by corresponding changes to equilibrium prices, with the firms that lose spectrum lowering their prices and the firms that gain spectrum raising prices. The effect of equalizing spectrum holdings on the quantity weighted average price and on welfare depends on the extent to which customers are sensitive to changes in service quality, relative to their sensitivity to changes in price. When customers care significantly more about price than service quality, customers experience a decrease in total consumer surplus due to the negative effects of price increases by the firms that receive additional spectrum outweighing the positive effects of price decreases by the firms that lose spectrum. Profits also decrease because the losses for the firms that lose spectrum outweigh the gains by the firms that gain spectrum. However, when customers give relatively equal weight to the importance of price and service quality, welfare is increased by the equalization of spectrum holdings. The quantity weighted average price is marginally higher and firms that lose spectrum lose profit, but this is outweighed by the increase in consumer surplus and profits at the firms who receive additional spectrum.

Empirical studies of the wireless industry have typically focused on substitution between wireless and wireline service, on the effect of taxation, and on the effect of capital investment on the economy as a whole. Hausman (2000) examines the efficiency of federal, state, and local taxes placed on wireless service. He concludes that these taxes are inefficient, costing about 50% more than they raise in revenues by surpressing demand for wireless services. Ingraham and Sidak (2004) examine the same issue several years later, and importantly, several years further into the adoption of wireless services, finding that the magnitude of the inefficiency of taxes on wireless service is even larger than the the level of inefficiency measured by Hausman (2000).

Caves (2011) examines the extent to which customers substitute wireless service for wireline service. Previous estimates found little substitution, but were conducted before widespread adoption of wireless service. Many households have both wireline and wireless telephones, making it unclear if the two services are substitutes or complements. He estimates a demand curve for wireline and wireless service and finds that a one percent decrease in the price of wireless service reduces demand for wireline service by between 1.2 and 1.3%.

Earlier literature has examined the link between investment in wireline telecommunications infrastructure and economic growth. For example, Röller and Waverman (2001) find a significant causal link between a country's investment in telecommunications infrastructure and the country's overall economic output. They find that infrastructure that allows for near universal telecommunications service leads to significantly higher economic growth. One focus of public policy throughout the development of the wireless industry in the United States has been expanding coverage in rural areas to allow for universal coverage by wireless telecommunications networks. As policies regulating the allocation of spectrum affect the growth of wireless service, spectrum policies are an important factor in reaching universal service.

The remainder of this paper proceeds as follows. The theoretical model is presented in Section 3.2, along with a simple, symmetric, two firm example useful for developing intuition for the sections that follow. The details of the data used to calibrate the model, as well as the calibration process, are discussed in Section 3.3. The results of simulations under counterfactual spectrum allocations are presented in Section 3.4. Section 3.5 includes discussion of several checks on the robustness of the results to the calibration process and counterfactual spectrum allocations. Section 3.6 concludes.

3.2 Model

In this section the model of the wireless industry that will be used in the simulations found in the following sections will be presented. The model is based on the more general model of the wireless industry found in Pinto and Sibley (2013).

There are $n \geq 2$ firms that sell a single, partially differentiated good. The demand for firm i is given by

$$q_i^D(p_i, p_{-i}, s_i, s_{-i}) = \alpha_i - \beta_{ii}p_i + \sum_{j \neq i} \beta_{ij}p_j + \lambda_{ii}s_i - \sum_{j \neq i} \lambda_{ij}s_j \tag{3.1}$$

where p_i is the price set by firm i and s_i is the quality of service provided by firm i's network. It is assumed that $\beta_{ii}, \beta_{ij}, \lambda_{ii}, \lambda_{ij} > 0$ so that firm i's demand is decreasing in its own price, increasing in the price of the other firms, increasing in its own service quality, and decreasing in the service quality of the other firms.

Firm i's service quality depends on the load placed on its network and on the network's capacity to handle that load. Let q_i^S be the load supplied to customers of firm i on its network and k_i be the capacity of the network. Then, based on the Kleinrock formula for expected packet delay, quality depends inversely on the difference between capacity and load. That is, firm i's quality is given by $s_i = k_i - q_i^S$.

For a firm's network to function, load must be balanced. That is, for a given service quality, s_i, and capacity, k_i, the load supplied by the firm must be equal to the load demanded by its customers: $q_i^S = q_i^D$. Using this-load balancing condition together with the expression for q_i^D given by equation (3.1), and rewriting service quality in terms of the load supplied (i.e., $q_i^S = k_i - s_i$), we get

$$k_i - s_i = q_i^S = q_i^D = \alpha_i - \beta_{ii} p_i + \sum_{j \neq i} \beta_{ij} p_j + \lambda_{ii} s_i - \sum_{j \neq i} \lambda_{ij} s_j \tag{3.2}$$

For load to be balanced for each firm, equation (3.2) must hold simultaneously for all firms. Define the following vectors for the capacities, service qualities, and prices of each firm, and the following vector and matrices of parameters:

$$K \equiv \begin{bmatrix} k_1 \\ k_2 \\ \vdots \\ k_n \end{bmatrix}, S \equiv \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{bmatrix}, P \equiv \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \end{bmatrix}$$

$$A \equiv \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix}, B \equiv \begin{bmatrix} \beta_{11} & -\beta_{12} & \cdots & -\beta_{1n} \\ -\beta_{21} & \beta_{22} & \cdots & -\beta_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ -\beta_{n1} & -\beta_{n2} & \cdots & \beta_{nn} \end{bmatrix}, L \equiv \begin{bmatrix} \lambda_{11} & -\lambda_{12} & \cdots & -\lambda_{1n} \\ -\lambda_{21} & \lambda_{22} & \cdots & -\lambda_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ -\lambda_{n1} & -\lambda_{n2} & \cdots & \lambda_{nn} \end{bmatrix}$$

Using these, the system of n equations made up of equation (3.2) for each firm can be written as follows

$$K - S = A - BP + LS \tag{3.3}$$

With each firm's demand given by equation (3.1), where each firm's demand depends on its own service quality and the service quality of the rival firms, when load is balanced for each firm, each firm's service quality is given by

$$\widehat{S} \equiv \begin{bmatrix} \hat{s}_1 \\ \hat{s}_2 \\ \vdots \\ \hat{s}_n \end{bmatrix} = (I + L)^{-1}(K - A + BP) \tag{3.4}$$

where I is the identity matrix.

Thus, for each firm, demand depends on its own service quality and the service quality of other firms, and its own service quality and the service quality of other firms depends on each firm's demand. Accordingly, firm i's load-balanced service quality is a function of the price and capacity of each firm: $\hat{s}_i(P, K)$. Substituting this in for the service quality in equation (3.1), firm i's ***load-balanced demand*** function is

$$\hat{q}_i(p_i, p_{-i}, k_i, k_{-i}) = \alpha_i - \beta_{ii} p_i + \sum_{j \neq i} \beta_{ij} p_j + \lambda_{ii} \hat{s}_i(p_i, p_{-i}, k_i, k_{-i}) - \sum_{j \neq i} \lambda_{ij} \hat{s}_j(p_i, p_{-i}, k_i, k_{-i}) \tag{3.5}$$

Equation (3.5) gives firm i's demand when load is balanced for each firm. Given this load-balanced demand function, and letting c_i denote firm i's marginal cost, firm i chooses its price in order to maximize profits, solving

$$\max_{p_i} \quad (p_i - c_i)\, \hat{q}_i(p_i, p_{-i}, k_i, k_{-i}) \tag{3.6}$$

The first order condition for firm i is

$$\hat{q}_i(p_i, p_{-i}, k_i, k_{-i}) + (p_i - c_i)\frac{\partial \hat{q}_i(p_i, p_{-i}, k_i, k_{-i})}{\partial p_i} = 0 \tag{3.7}$$

where

$$\frac{\partial \hat{q}_i(p_i, p_{-i}, k_i, k_{-i})}{\partial p_i} = -\beta_{ii} + \lambda_{ii}\frac{\partial \hat{s}_i(p_i, p_{-i}, k_i, k_{-i})}{\partial p_i} - \sum_{j \neq i} \lambda_{ij}\frac{\partial \hat{s}_j(p_i, p_{-i}, k_i, k_{-i})}{\partial p_i}$$

Equilibrium prices are found by solving the system of first order conditions for each firm. Note that firm i's first order condition can be rewritten as an inverse elasticity pricing rule:

$$\frac{p_i - c_i}{p_i} = -\frac{\hat{q}_i(p_i, p_{-i}, k_i, k_{-i})}{p_i\left(\frac{\partial \hat{q}_i(p_i, p_{-i}, k_i, k_{-i})}{\partial p_i}\right)} = -\frac{1}{\varepsilon_i} \tag{3.8}$$

3.2.1 Optimal Capacities

The discussion in the preceding section assumes that each firm capacity is given. In this section we will examine the optimal choice of capacity. Given the capacities and prices of rival firms, and given its own capacity, k_i, firm i's optimal price is given by equation (3.8). If we denote this optimal price by $p_i^*(k_i, k_{-i})$ and let $C_i(k_i)$ be the cost of capacity, firm i chooses its capacity to solve

$$\max_{k_i} \quad (p_i^*(k_i, k_{-i}) - c_i)\, \hat{q}_i(p_i^*(k_i, k_{-i}), p_{-i}^*(k_i, k_{-i}), k_i, k_{-i}) - C_i(k_i) \tag{3.9}$$

Firm i's optimal capacity is given by the solution to the first order condition (omitting the arguments of $\hat{q}_i(p_i^*(k_i, k_{-i}), p_{-i}^*(k_i, k_{-i}), k_i, k_{-i})$ and $p_i^*(k_i, k_{-i})$)

$$\left(\frac{\partial p_i^*}{\partial k_i} - c_i\right)\hat{q}_i + (p_i^* - c_i)\frac{\partial \hat{q}_i}{\partial k_i} = \frac{\partial C_i(k_i)}{\partial k_i} \tag{3.10}$$

Equilibrium capacities are obtained by solving the system of first order conditions for each firm. Then equilibrium prices are obtained by evaluating $p_i^*(k_i, k_{-i})$ at the equilibrium capacities. Thus the decisions of firms can be viewed as having two stages, with capacities selected in the first stage and prices selected in the second.

Our focus will be on equilibrium outcomes of the second stage of the model. Decisions about capacity are made over a longer time period than decisions about prices. Once decisions about changes to capacity are made, it can take firms substantial periods of time to implement those changes. For example, a firm may decide to deploy more cell towers to increase the capacity of its network. However, once this decision has been made, locations for the towers must be determined and rights to those locations obtained from what can be numerous governmental agencies,[1] and only then can the towers be constructed and integrated into the firm's network.

Similarly, a firm that wishes to increase its capacity by increasing its spectrum holdings also faces delays and potential obstacles. In the United States, the use of spectrum is regulated by the FCC. Spectrum is a finite resource. The vast majority of all spectrum is currently put to some form of use. Thus, the most common way for firms to obtain additional spectrum is to acquire it from other spectrum holders (who may or may not be rival firms). These transfers of spectrum require approval

[1] On the federal level, approval is required from the FCC and, depending on the location and height of the structure, the Federal Aviation Administration. In addition, approval must also be obtained from some combination of state, county, and local governments.

by the FCC, a process that can can take some time and does not always end with approval being given.

Occasionally the FCC chooses to change the use of a spectrum band and new spectrum becomes available for firms to use as part of their wireless networks. For example, in the late 2000s, television broadcasting was switched from analog to digital signals, making available spectrum in the 700 MHz band. In 2008, the FCC auctioned spectrum in the 700 MHz band for use by mobile wireless service providers.[2] Overall, new spectrum becomes available infrequently, and only in large blocks. Consequently, while firms might be able to formulate first order conditions similar to equation (3.10) and determine their optimal capacity k_i^*, in practice it is often not possible to implement exactly k_i^*. Thus, an examination of how different allocations of capacity affect equilibrium in the wireless market is of interest.

3.2.2 Two Firm Example

In the next section we will calibrate the model using data from the wireless industry and examine the effects of changes in capacity on second stage equilibrium outcomes such as prices, quantities, and service qualities. First, it is useful to examine a symmetric, two firm example in order to develop intuition for the effects of a firm increasing its capacity. With this in mind, suppose that $c_1 = c_2 = 1$ and the parameters are

$$A = \begin{bmatrix} 10 \\ 10 \end{bmatrix}, B = \begin{bmatrix} 8 & -4 \\ -4 & 8 \end{bmatrix}, L = \rho B = \begin{bmatrix} \rho 8 & -\rho 4 \\ -\rho 4 & \rho 8 \end{bmatrix}$$

The new parameter, ρ, is the relative importance customers place on price compared to service quality. When $\rho = 1$, customers care equally about price and

[2]For a discussion of the 700 MHz band, see `http://www.fcc.gov/encyclopedia/700-mhz-spectrum`.

service quality, when $\rho < 1$, customers care relatively more about price than service quality, and when $\rho > 1$, customers care relatively more about service quality than price.

Using these parameter values, we can solve the firms' second stage first order conditions, as given by equation (3.7), to find the equilibrium price for each firm, given capacities k_i and k_j for $i, j \in \{1, 2\}, i \neq j$. For $\rho = 1$, the equilibrium price of firm i is given by

$$p_i^*(k_i, k_j) = 0.4994k_i - 0.0179k_j + 1.7222 \tag{3.11}$$

Given this expression for the equilibrium price of each firm, $p_i^*(k_i, k_j)$, we can then solve for the equilibrium quantities and service qualities as given by equations (3.5) and (3.4). The resulting equilibrium price, output, and service quality for firm i are

$$\hat{q}_i^*(k_i, k_j) = 0.4302k_i - 0.0154k_j + 0.6222 \tag{3.12}$$

$$\hat{s}_i^*(k_i, k_j) = 0.5698k_i + 0.0154k_j - 0.6222 \tag{3.13}$$

To develop intuition for how changes in capacity affect each firm, it is useful to examine how changes in k_i and k_j affect firm i's price, quantity, and service quality for different values of ρ. Table 3.1 shows the coefficients on k_i and k_j in the function that gives firm i's equilibrium price ($p_i^*(k_i, k_j)$), quantity ($\hat{q}_i^*(k_i, k_j)$), and service quality ($\hat{s}_i^*(k_i, k_j)$) for $\rho = 0.75$, $\rho = 1$, and $\rho = 1.25$.

Holding everything else constant, firm i's equilibrium price, quantity, and service quality are all increasing in its capacity. However, the size of the effect of an increase in capacity on price, quantity, and service quality depends on the relative

Table 3.1: Effects of Capacities on Firm *i*

	Price		Quantity		Quality	
ρ	k_i	k_j	k_i	k_j	k_i	k_j
0.75	0.3742	-0.0171	0.4116	-0.0188	0.5884	0.0188
1.00	0.4994	-0.0179	0.4302	-0.0154	0.5698	0.0154
1.25	0.6245	-0.0184	0.4423	-0.0130	0.5577	0.0130

Table 3.1: Coefficients on k_i and k_j in firm i's second stage equilibrium price, load-balanced demand, and load-balanced service quality functions. Note that for $\rho = 1$, these coefficients correspond with equations (3.11) through (3.13).

importance of price and service quality for customers, as measured by ρ. As ρ increases, the influence of capacity on price and output increases, but the influence of capacity on service quality decreases. Recall that an increase in ρ is an increase in the relative importance of service quality for customers compared to price. As service quality becomes more important for customers, the more an increase in service quality increases demand and the more the firm is able to increase its price. This is why the coefficient on k_i is increasing with ρ for quantity and, in particular, for price.

Now consider the coefficient on k_i for firm i's service quality. The coefficients are positive because an increase in capacity increases service quality. However, the magnitude is decreasing as ρ increases. As service quality becomes relatively more important for customers, demand grows by more as capacity, and in turn, service quality, increases. However, an increase in demand has the opposite effect on service quality, decreasing it, which reverses some of the increase in service quality obtained by increasing capacity. Thus, while an increase in capacity increases service quality, it does so at a decreasing rate as service quality becomes relatively more important for customers.

Comparing the coefficients on k_i across price, quantity, and quality, for a given value of ρ, we notice several things. First, the effect of capacity on quality is larger than the effect of capacity on quantity because some of the increase in demand resulting from the increase in service quality is offset by the fact that the firm also increases its price. Second, the effect of an increase in capacity on price is larger than the effect on quality only when customers care more about service quality than price. The more customers care about service quality, the more the firm is able to increase its price as it increases service quality. When customers care relatively less about service quality, the same increase in service quality does not allow the firm to increase price by as much. At the same time, because the firm increases price more when customers care more about service quality, the larger the offsetting decrease in demand resulting from the price increase, and the larger the offsetting decrease in demand, the larger the increase in service quality resulting from the change in capacity.

Let us now consider the effects of an increase in k_j, the capacity of the rival firm, on firm i's price, quantity, and quality. As firm j increases its capacity, and thus its service quality, some customers leave firm i and go to firm j. This causes a decrease in firm i's quantity, but an increase in firm i's service quality.

The more customers care about service quality, the more firm i is forced to decrease its price when the rival firm increases its service quality. Thus, as ρ increases, the effect of an increase in k_j on firm i's price increases in magnitude. Conversely, the higher is ρ, the lower is the effect of an increase in k_j on firm i's quality and quantity. Recall that as a firm increases its capacity, its quality increases, but at a decreasing rate as ρ increases. Thus, the higher is ρ, the less of an effect an increase in k_j has on firm j's service quality, prompting relatively fewer customers to leave firm i for firm j. This effect is strengthened by the fact that the higher is ρ, when firm j increases its capacity, the more firm j increases its price and the more firm i decreases its price.

Overall, the effect of a firm's own capacity on its price, quantity, and quality, are much larger than the effects of a change in the rival firm's capacity. In the next section, these effects will be explored in a model calibrated to the wireless industry.

3.3 Calibration

In this section, the model presented in section 3.2 will be calibrated to data from the wireless industry. We will first discuss the data used in the calibrations before turning to the details of the calibration process. Then, in the next section the calibrated model will be used to examine the effects of alternative allocations of spectrum.

3.3.1 Data

We will focus on the four firms that provide wireless service on a national basis in the United States: AT&T, Sprint Nextel, T-Mobile, and Verizon Wireless. The model will be calibrated to match the quantities, prices, and capacities found in the data. All data used in the calibrations are available from the fourteenth, fifteenth, and sixteenth annual "Mobile Wireless Competition Reports" of the FCC,[3] covering approximately the years 2009 through 2011. The data are summarized in Table 3.2.

The model presented in the previous section assumes that quantity is measured at the level of a customer, and does not go in to detail about the usage patterns of specific customers. Generally, a customer selects a plan from a menu of available options offered by the firm, uses the firm's network for talk, text, and data, and then pays a bill at the end of the month, potentially including fees for exceeding the allotment of minutes, texts, or data that are included in the chosen plan. As these overage fees are typically quite high, customers generally choose a plan that allows them to avoid paying any overage fees, making the cost of marginal use effectively zero. Customers typically follow usage patterns that may vary based on location and

[3]See `http://www.fcc.gov/reports/mobile-wireless-competition-report-16th-annual`.

demographics (e.g., in a college town, in an urban businesses district), but that are otherwise fairly predictable. Accordingly, with the exception of a minority of outliers, the effect of an additional customer on a firm's service quality vis-à-vis the capacity of the firm's network is reasonably predictable, so while focusing on the customer as the applicable unit of measure is clearly an abstraction, it should not be one that substantively affects the results. Thus, quantity will be measured in customers.

Corresponding with a customer being the unit of measure for output, the price will be the monthly price of a standard plan that includes approximately 450 minutes with unlimited text messaging and data. While all of the firms offer various other plans, this combination was offered by all firms during the period considered. The only exceptions were that T-Mobile offered 300 minutes in 2009 and 500 in 2010 and 2011, AT&T limited data usage to 2GB in 2011, and "anytime" minutes started at 7 P.M. instead of 9 P.M. for Sprint. Plans that offered different features were fairly common across the firms and were generally matched with similar discrete changes in plan price. Selecting a different, but comparable across the firms, plan as the basis of the price did not substantively change the results.

The measure of capacity used for a firm was the firm's population-weighted "MHz-POPs" of spectrum across all frequency bands. A firm having one megahertz (MHz) of spectrum refers to the firm having the right to use radio frequencies that span one million hertz at a particular frequency, e.g., from 700 to 701 MHz. One "POP" refers to one person. The measure "MHz-POPs" refers to the total MHz of spectrum a firm has the right to use in a geographic region, multiplied by the population of that region. A firm's population-weighted MHz-POPs is the firm's average MHz-POPs, weighted by population.

While population-weighted MHz-POPs is clearly an imperfect measure of capacity, it does serve as a good proxy. It is not possible to offer wireless service without

Table 3.2: Data Used For Calibrations

	Year	AT&T	Sprint Nextel	T-Mobile	Verizon
	2009	85,120	48,133	33,790	85,445
Quantity	2010	95,536	49,910	33,734	87,535
	2011	103,247	55,021	33,185	92,167
	2009	$70.00	$70.00	$60.00	$80.00
Price	2010	$85.00	$70.00	$80.00	$90.00
	2011	$85.00	$80.00	$80.00	$90.00
	2009	82.0	52.5	50.4	87.7
Spectrum	2010	76.8	51.2	47.7	83.4
	2011	88.3	53.0	57.0	107.3

Table 3.2: Data from FCC's 14^{th},15^{th}, and 16^{th} Mobile Wireless Competition Reports. Quantity is the number of subscribers, measured in thousands. Price is the monthly price of a base plan with approximately 450 minutes and unlimited text and data. Spectrum is the nationwide population weighted MHz-POPs across all frequency bands.

access to spectrum. All else equal, the more spectrum a firm is able to use, the higher its capacity to offer service and the higher the potential quality of that service. The main limitations to this measure of capacity are that it ignores differences in frequencies and differences in technologies. Spectrum ranging from around 700 MHz up to around 2500 MHz[4] is currently used for mobile wireless communications. However, the properties of spectrum vary across this range of frequencies, including differences in the ability to transmit signals through adverse climates and topographies and over long ranges, and to penetrate buildings, automobiles, and other structures. While some frequency bands are generally viewed as having more favorable characteristics than other frequencies for offering mobile wireless service, some also argue that it is optimal to possess complementary frequencies in both lower and higher ranges. For our purposes here, we will ignore any differences in frequencies and simply consider

[4]2500 MHz is also commonly referred to as 2.5 GHz.

the total spectrum holdings of each firm.

Differences in technology also affect capacity and service quality for a given amount and frequency of spectrum. Several main technologies emerged in the second generation of mobile wireless communication following the original analog cellular technology, including CDMA (Code Division Multiple Access), TDMA (Time Division Multiple Access), GSM (Global System for Mobile Communications), and iDen (integrated Digital Enhanced Network). These "2G" technologies differ mostly in the method by which a signal is divided to allow for simultaneous use by multiple devices, e.g., division by "time", division by "code." Over time, various additional standards were developed to squeeze addition use out of the airwaves, progressing from 2G to 2.5G, 3G, 3.5G, and most recently, 4G. While the various abbreviations (e.g., GPRS, WCDMA, EDGE, 1xRTT, EV-DO, LTE) are likely both familiar and foreign to most customers, the common theme is that over time, firms have been able to transmit signals of increasing quality and reliability, at higher speeds, and to more customers using the available spectrum.

Regardless of the technology deployed by a firm to create the signals to be sent between devices and the frequencies bands used to transmit them, firms require spectrum to provide mobile wireless communications. The more spectrum held by a firm, the higher the firm's capacity to serve customers and the higher its potential service quality, making MHz-POPs a reasonable proxy for a firm's overall capacity. Using spectrum as the measure of capacity is also of interest because the allocation of spectrum is partially determined by public policy. The FCC controls the use of spectrum, and when it decides to change how certain frequencies are used, it controls the mechanism by which the spectrum is allocated. For example, when reallocating spectrum in the 700 MHz band from use in television transmission to mobile wireless communication, it designed and conducted the auctions used to allocate spectrum to its new holders. Thus, examining the effects of different allocations of spectrum sheds some light on the successes and failures of these allocation mechanisms.

3.3.2 Model Calibration

In order to calibrate the model to data from the wireless industry, parameters will be chosen to match the equilibrium quantities and prices observed in the data to those predicted by the model when the capacity of each firm is the capacity observed in the data. The load-balanced demand function,

$$\hat{q}_i(P, K) = \alpha_i - \beta_{ii}\left[p_i - \rho \hat{s}_i(P, K)\right] + \sum_{j=1, j \neq i}^{4} \beta_{ij}\left[p_j - \rho \hat{s}_j(P, K)\right] \tag{3.14}$$

gives the quantity for firm i when prices are given by the vector $P \equiv (p_1, p_2, p_3, p_4)$, capacities are given by the vector $K \equiv (k_1, k_2, k_3, k_4)$, and $\hat{s}_i(P, K)$ is the load-balanced service quality of firm i at P and K, given by equation (3.4). In addition to the load-balanced demand function, we will also make use of each firm's first order condition with respect to price

$$\hat{q}_i(P, K) + (p_i - c_i)\frac{\partial \hat{q}_i(P, K)}{\partial p_i} = 0 \tag{3.15}$$

The load-balanced demand function specified by equation (3.14) allows for consumer sensitivity to price and service quality to vary by firm, but assumes that sensitivity to service quality is a fixed fraction of sensitivity to price. The parameter ρ specifies this fixed fraction, giving the relative importance for customers of price compared to service quality. We will consider a range of values for ρ. Values of remaining parameters will be chosen to match the prices, quantities, and capacities observed in the data. The model is calibrated separately for each value of ρ.

For a value of ρ, the model is calibrated as follows. First, for a given vector of prices observed in the data, the vector of marginal costs is determined assuming a constant markup of forty percent. Second, each firm's first order condition with respect to its price is used to determine the betas, assuming that the first order conditions are satisfied at the equilibrium prices, quantities, and capacities observed in the data. Each firm's first order condition can be written in the form of an inverse elasticity pricing rule,[5] which, when combined with the assumed markup, yields the firm's own-price elasticity at the equilibrium price observed in the data. This, when combined with the equilibrium quantity observed in the data, provides the diagonal elements of the matrix of betas, B. The off-diagonal elements of B are then determined using diversion ratios calculated from the quantity market shares observed in the data. Given the B matrix, the matrix of lambdas is given by $L \equiv \rho B$ for the current value of ρ. Last, the vector of alphas is chosen such that the quantity predicted by the load-balanced demand function, when evaluated at the equilibrium prices and capacities observed in the data and using the matrices B and L just determined, matches the quantities observed in the data.

When the model is fully calibrated, solving the system of first order conditions evaluated at the observed capacities yields the observed equilibrium prices, and evaluating the load-balanced demand function for each firm at the equilibrium prices and observed capacities yields the observed equilibrium quantities. Using the load-balanced demand function and the firms' profit functions, it is also possible to calculate firm profits, as well as consumer and total surplus. Note that firm profits are calculated without any costs associated with capacity.

In order to evaluate the effects of alternative allocations of spectrum, this post-calibration process is repeated, except using counterfactual capacities rather than the

[5]See equation (3.8).

actual capacities observed in the data. That is, using the parameters calibrated to the actual capacities, the solution to the system of first order conditions evaluated at counterfactual capacities yields the counterfactual equilibrium prices. Then, also using the parameters calibrated to the actual capacities, the load-balanced demand function for each firm is evaluated at the counterfactual equilibrium prices and counterfactual capacities, yielding the counterfactual equilibrium quantities. The load-balanced demand function and the firms' profit functions, evaluated using the parameters calibrated to the actual capacities but with the counterfactual capacities and corresponding counterfactual equilibrium prices and quantities, can be used to calculate counterfactual firm profits, consumer surplus, and total surplus.

3.4 Effects of Spectrum Holdings

In this section the model calibrated in the manner discussed in the previous section will be used to simulate the effects of alternative allocations of spectrum. Three main counterfactual scenarios will be explored. In the first, one firm is given additional spectrum, holding constant the spectrum of the remaining firms. In the second, the total quantity of spectrum is held constant, but it is reallocated equally among the four firms. In the third, spectrum is transfered from one firm to another, holding constant the spectrum of the other two firms and the total quantity of spectrum. Each counterfactual allocation leads to similar conclusions about the effects of spectrum on equilibrium in the wireless market.

3.4.1 Additional Spectrum for One Firm

In this section we will consider the effects of giving one firm additional spectrum, holding constant the spectrum of the remaining firms. We will examine in detail two scenarios, one with a larger firm and one with a smaller firm being the one to increase its spectrum holdings. Let us first consider AT&T, one of the two largest firms. We will then consider Sprint, one of the two smaller firms.

Suppose AT&T receives 10% additional spectrum, holding constant the spectrum of the other firms. Table 3.3 shows counterfactual equilibrium quantities for different values of ρ for this counterfactuall allocation of spectrum. When AT&T receives additional spectrum, its quantity increases and the quantities of the other firms decrease. Because the spectrum of the other firms has not changed, they each have the same capacity but now have fewer customers using their networks, so their service qualities increase. AT&T now has more capacity, but it also has more load on its network. However, the increase in capacity exceeds the increase in quantity, so AT&T's quality increases as well. The effect on total quantity is small, but positive.

Now consider Table 3.4, which shows the effect on prices of AT&T's increase in spectrum. AT&T's additional spectrum increases its service quality, which allows it to increase its price. In response, the other firms lower their prices. Overall, because AT&T starts as one of the largest firms and further increases its quantity, the quantity weighted average price increases.

Suppose instead that Sprint receives 10% additional spectrum, holding constant the spectrum of the other firms. Table 3.5 shows the effect on equilibrium quantities of this allocation of spectrum. Just as was the case when it was Verizon that received the additional spectrum, the firm that receives the additional spectrum has a higher quantity, while the remaining firms have lower quantities. Sprint's quantity increases, but not by as much as its increase in spectrum, resulting in an increase in service quality. The remaining firms have the same capacity with which to serve a now lower quantity of customers, meaning that their service qualities increase as well. Quantity increases overall, though only slightly.

Table 3.6 shows the effect on equilibrium prices. Sprint's higher service quality allows it to increase its price, and in order to compete with Sprint, the other firms lower their prices. Sprint is one of the smaller firms, even after its increase in quantity. Thus, the effect of the other firms decreasing their prices dominates the effect of Sprint raising its price, resulting in a lower quantity weighted average price.

3.4.2 Equalizing Spectrum Holdings

In this section we will consider the effects of redistributing total spectrum holdings equally among the firms, holding constant the total quantity of spectrum. The intuition developed in the previous section in which the spectrum holdings of all firms except one were held constant should prove useful in this more complex scenario.

Initially, the two larger firms, AT&T and Verizon, hold larger shares of total spectrum than do Sprint and T-Mobile. Consequently, when the total quantity of spectrum is divided equally among all four firms, Sprint and T-Mobile receive additional spectrum while AT&T and Verizon experience a reduction in their spectrum holdings. In the previous section we saw how when a firm increases its spectrum holdings, it increases in quantity and its price. In response, the other firms decreased their prices and had lower quantities. This pattern will continue to hold here.

Consider first the effect of equalizing spectrum holdings on equilibrium quantities. As shown in Table 3.7, Sprint and T-Mobile, the two firms who increase their spectrum holdings as a result of the equalization, have higher quantities. AT&T and Verizon, the two firms who now hold less spectrum, have lower quantities. Cumulatively, total quantity decreases by around three quarters of one percent.

The effect of equalizing spectrum on the price of each firm follows a similar pattern to the effect on quantities. Table 3.8 shows the counterfactual equilibrium prices following the redistribution of spectrum. Sprint and T-Mobile, the beneficiaries

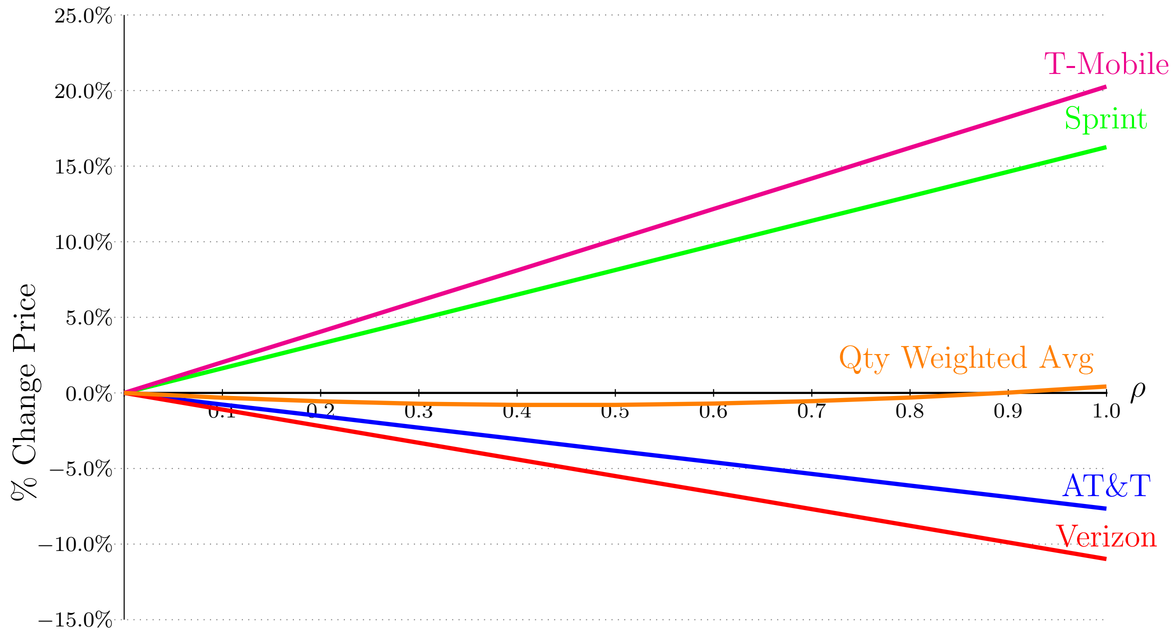

Figure 3.1: Percent change of equilibrium price for each firm and the quantity weighted average when all firms receive an equal share of total spectrum as a function of ρ

of the spectrum reallocation, both increase their price, while AT&T and Verizon both decrease theirs. The magnitude of the price changes are largest for Sprint and T-Mobile. Both firms increase their price to a greater extent, both in absolute and percentage terms, compared to the amount by which AT&T and Verizon lower their prices.

The overall effect on the quantity weighted average price depends on customers' sensitivity to change in service quality relative to price. This can be seen in Figure 3.1. When $\rho = 1$, customers care as much about service quality as they do about price. Thus, when ρ is closer to 1, Sprint and T-Mobile are able to increase their price to a greater extent than when ρ is closer to zero. When ρ is close to one, Sprint and T-Mobile both increase their price by more than AT&T and Verizon

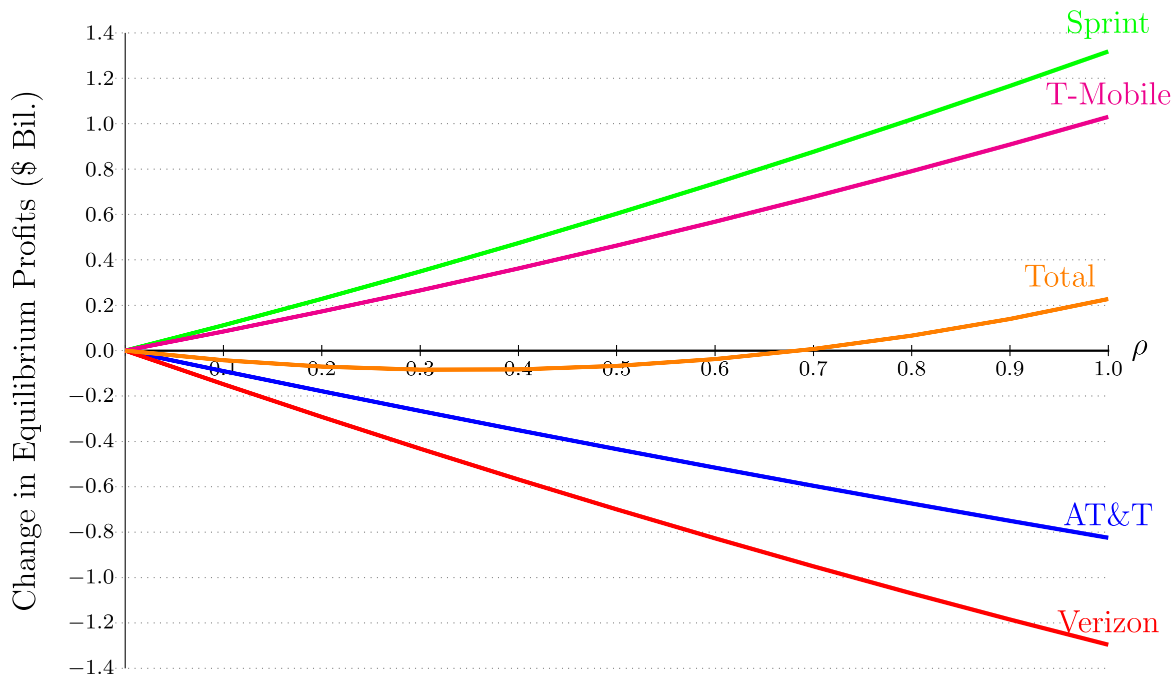

Figure 3.2: Change in equilibrium profits in $Billions when all firms receive an equal share of total spectrum as a function of ρ

decrease theirs, pulling up the quantity weighted average price. However, for lower values of ρ, the quantity weighted average price decreases. While AT&T and Verizon do not decrease their price by as much as Sprint and T-Mobile increase theirs, they have sufficiently higher quantities to pull the quantity weighted average price down.

A similar pattern can be seen in Figure 3.2 for firm profits. The reallocation of spectrum benefits Sprint and T-Mobile, who both have higher profits in equilibrium. AT&T and Verizon, who both lose spectrum, have lower profits in equilibrium. For low values of ρ, total profits decrease. The more weight customers give to service quality relative to price, the more profitable the additional spectrum is for Sprint and T-Mobile. As ρ becomes closer to one, the profit gains for Sprint and T-Mobile outweigh the decrease in profits for AT&T and Verizon, resulting in an overall increase in profits.

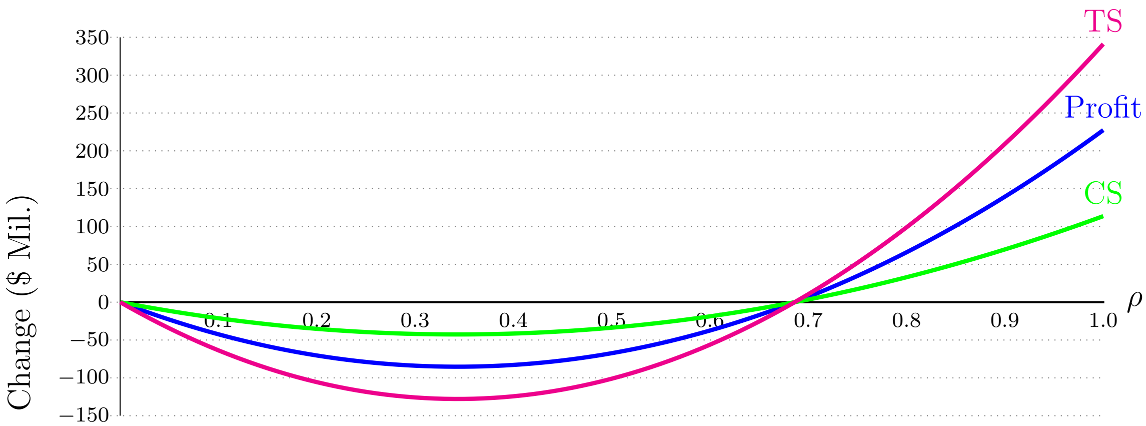

Figure 3.3: Change in equilibrium profits, consumer surplus, and total surplus (in $Millions) when all firms receive an equal share of total spectrum as a function of ρ

Figure 3.3 shows the welfare effects of equalizing spectrum holdings. For lower values of ρ, profits and consumer surplus decrease. However, when ρ becomes high enough, profits and consumer surplus begin to increase. As ρ approaches one, where customers care equally about service quality and price, welfare is higher as a result of the equalization of spectrum. However, this increase in welfare is far from a Pareto improvement, as some firms and customers lose while others gain.

3.4.3 Transfer of Spectrum Between Two Firms

When the proposed merger between AT&T and T-Mobile failed to succeed, AT&T transfered spectrum to T-Mobile in accordance with the breakup terms of the deal. In this section, we will consider the effects of AT&T transferring 10% of its spectrum to T-Mobile. The spectrum holdings of Sprint and Verizon remain unchanged.

The effects of this spectrum transfer from AT&T to T-Mobile on equilibrium

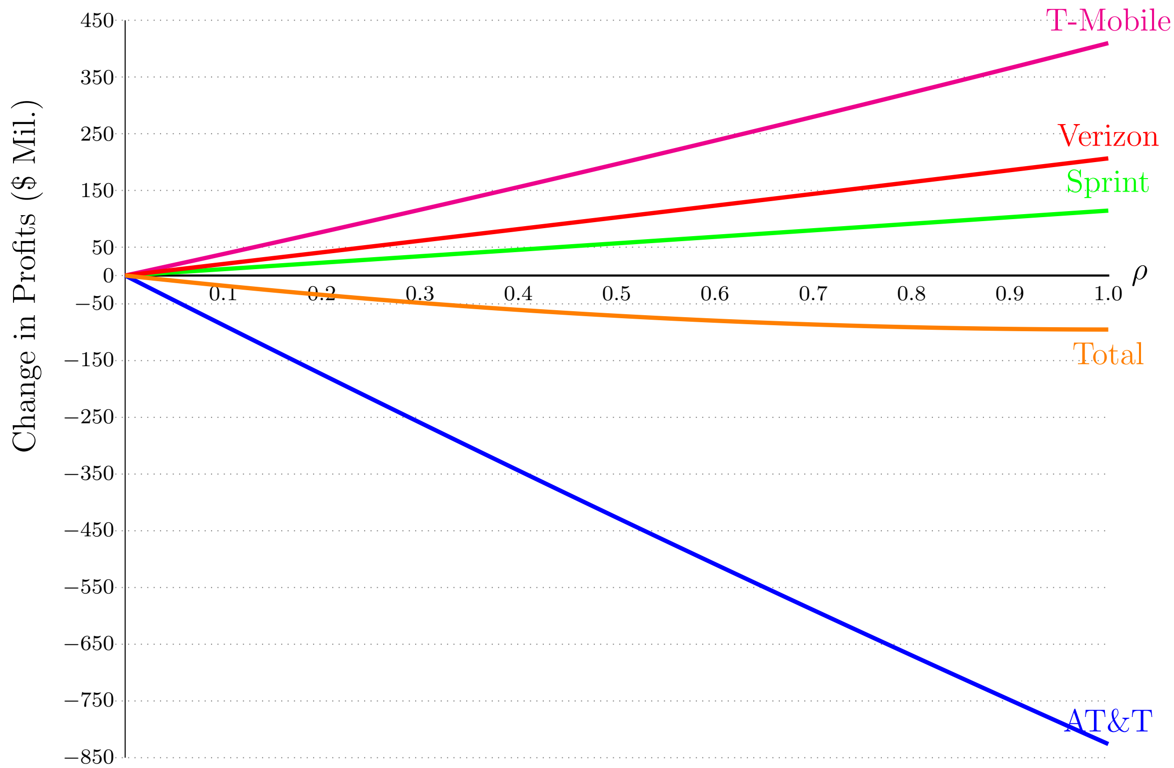

Figure 3.4: Change in equilibrium profits in $Billions when AT&T transfers 10% of its spectrum to T-Mobile as a function of ρ

quantities is shown in Table 3.10. T-Mobile is the largest beneficiary, increasing its quantity by the most. Both Sprint and Verizon also increase their quantities by similar percentages. AT&T experiences a significant decrease in quantity.

In order to help lessen this decrease in quantity, AT&T lowers its price, as seen in Table 3.9. Sprint and Verizon both increase their prices slightly. T-Mobile increases its price by more, particularly when customers give more weight to service quality.

The effects of this spectrum transfer on firm profits is shown in Figure 3.4. As expected, AT&T suffers a large decrease in profits, while T-Mobile gains substantially.

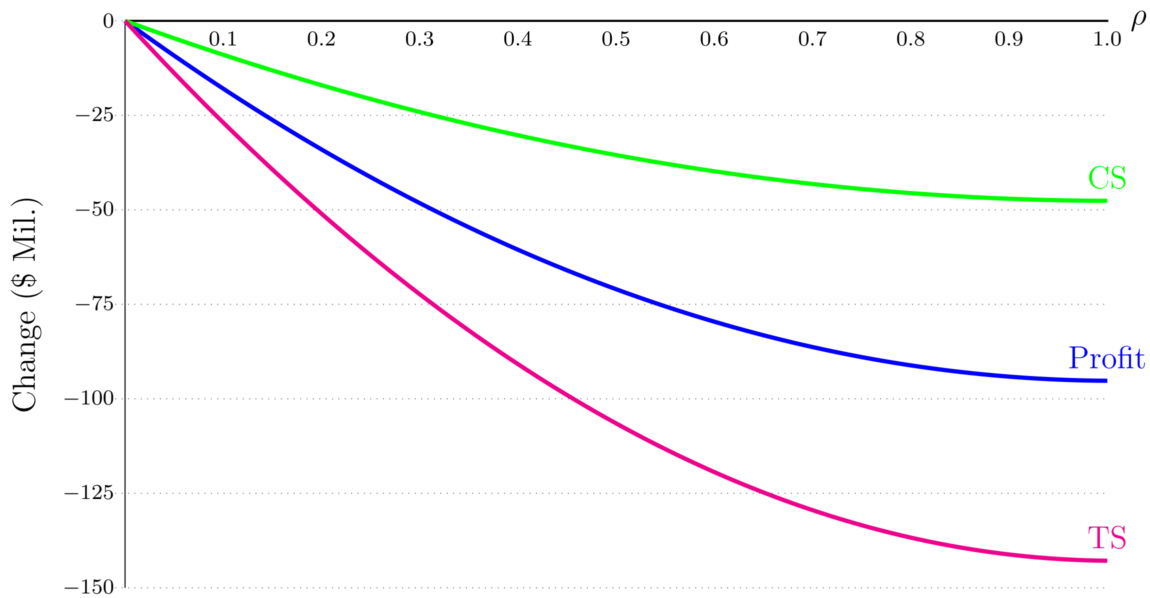

Figure 3.5: Change in equilibrium profit, consumer surplus, and total surplus (in $Millions) when AT&T transfers 10% of its spectrum to T-Mobile as a function of ρ

Verizon and Sprint also gain, each having slightly higher prices and quantities in the new equilibrium. Total profits, however, decline. The effects on consumer and total surplus are shown in Figure 3.5. Welfare decreases more steeply the lower is ρ, starting to level off as ρ approaches one. The decrease in total welfare is due to the large size of AT&T. If the transfer were to go in the other direction, the model would predict an increase in total welfare driven by a large increase in profits for AT&T. However, if this were to cause a reduction in T-Mobile's profits to the extent that T-Mobile was forced to leave the industry, this would not be captured by the present model.

3.5 Robustness

Numerous checks of the robustness of the model were conducted. The results presented in the previous section used data from a specific year, as noted. Identical simulations using data from other years were also conducted, with similar results. Simulations were also repeated using alternative values of the assumed markup over marginal cost, also yielding similar results.

The model was also calibrated to different measures of the price. The price for each firm used in the results discussed above was for a plan with 450 minutes, unlimited text messaging, and unlimited data, with the few exceptions previously noted. Simulations using the price of the following plans were also conducted

1. 450 minutes with text messages but no data
2. 950 minutes with text messages and data
3. 950 minutes with text messages but no data
4. unlimited minutes, text messages, and data

In addition to using the price of these alternative plans, average revenue per unit (ARPU), a common proxy for price used in studies of the wireless industry, was also used as the measure of price. In each case, the qualitative results remain unchanged.

As an additional robustness check on the results presented in the previous section, counterfactual allocations in which one firm was given 5% additional spectrum and 25% additional spectrum were also examined. Results with these counterfactual spectrum allocations were similar to those found when one firm was given 10% additional spectrum.

3.6 Conclusion

In this paper, a theoretical model similar to Pinto and Sibley (2013) was calibrated to data from the wireless industry. Using the calibrated model, simulations under various counterfactual allocations of spectrum were conducted in order to examine how a change in one firm's spectrum holdings affects other firms in equilibrium. If one firm receives additional spectrum, holding constant the spectrum of the other firms, that firm serves a higher quantity of customers and has a higher equilibrium price. The other firms lose customers and lower their prices. Service quality is higher at all firms.

Next, the effects of equalizing total spectrum holdings was simulated. When spectrum holdings are reallocated such that each firm receives an equal share of total spectrum, firms that start with a higher share of total spectrum lose spectrum to firms who start with a smaller share of spectrum. The firms that gain spectrum have higher service quality, attracting customers and allowing them to set higher prices in equilibrium. The firms that lose spectrum lower their prices. The welfare effects of the equalization of spectrum holdings depend on the degree to which customers care about service quality relative to price. If customers care relatively less about service quality, total welfare decreases. Only if customers care almost as much about service quality as they do about price is welfare increased by the equal distribution of spectrum.

Last, the transfer of spectrum from one firm to another was considered. When a larger firm with greater spectrum holdings transfers spectrum to a smaller firm that holds less spectrum, the firm that loses spectrum serves fewer customers, lowers its price, and has lower profits. The other firms, both the firm that receives the additional spectrum as well as the remaining firms whose spectrum remains unchanged, serve more customers, set higher prices, and have higher profits. The gains are largest for the firm that receives the additional spectrum. Total profits fall because the losses for the firm that loses spectrum outweigh the gains for the other firms.

In general, the effects of different allocations of spectrum depend on customer preferences over price and service quality. In this paper, customer sensitivity to price was calibrated to industry data, and then a range of possible preferences for service quality were considered. In future work, an explicit examination of customer preferences over both price and service quality would be valuable in examining the effects of spectrum policy.

Quantities When AT&T Receives 10% More Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	Total
	Counterfactual Quantities				
0	95,536	49,910	33,734	87,535	266,715
0.25	97,713	49,273	33,369	86,592	266,947
0.50	99,889	48,635	33,005	85,650	267,179
0.75	102,066	47,998	32,640	84,707	267,411
1.00	104,243	47,361	32,275	83,764	267,643
	Change (Counterfactual – Actual)				
0	0	0	0	0	0
0.25	2,177	-637	-365	-943	232
0.50	4,353	-1,275	-729	-1,885	464
0.75	6,530	-1,912	-1,094	-2,828	696
1.00	8,707	-2,549	-1,459	-3,771	928
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	2.28%	-1.28%	-1.08%	-1.08%	0.09%
0.50	4.56%	-2.55%	-2.16%	-2.15%	0.17%
0.75	6.84%	-3.83%	-3.24%	-3.23%	0.26%
1.00	9.11%	-5.11%	-4.33%	-4.31%	0.35%

Table 3.3: The top section has the counterfactual equilibrium quantities for each firm when AT&T receives 10% more spectrum, holding the spectrum of the other firms constant. The counterfactual equilibrium quantities equal the actual equilibrium quantities for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2010. All quantities are in thousands.

Prices When AT&T Receives 10% More Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	WgtAvg
	Counterfactual Prices				
0	$85.00	$70.00	$80.00	$90.00	$83.20
0.25	$85.77	$69.64	$79.65	$89.61	$83.27
0.50	$86.55	$69.28	$79.31	$89.22	$83.37
0.75	$87.32	$68.93	$78.96	$88.84	$83.48
1.00	$88.10	$68.57	$78.62	$88.45	$83.61
	Change (Counterfactual − Actual)				
0	$0	$0	$0	$0	$0
0.25	$0.77	-$0.36	-$0.35	-$0.39	$0.07
0.50	$1.55	-$0.72	-$0.69	-$0.78	$0.17
0.75	$2.32	-$1.07	-$1.04	-$1.16	$0.28
1.00	$3.10	-$1.43	-$1.38	-$1.55	$0.41
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	0.91%	-0.51%	-0.44%	-0.43%	0.09%
0.50	1.82%	-1.03%	-0.86%	-0.87%	0.20%
0.75	2.73%	-1.53%	-1.30%	-1.29%	0.33%
1.00	3.65%	-2.04%	-1.72%	-1.72%	0.49%

Table 3.4: The top section has the counterfactual equilibrium prices for each firm, as well as the quantity-weighted average price ("WgtAvg"), when AT&T receives 10% more spectrum, holding the spectrum of the other firms constant. The counterfactual equilibrium prices equal the actual equilibrium prices for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2010.

Quantities When Sprint Receives 10% More Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	Total
	Counterfactual Quantities				
0	95,536	49,910	33,734	87,535	266,715
0.25	95,093	50,914	33,588	87,158	266,753
0.50	94,650	51,918	33,443	86,782	266,793
0.75	94,207	52,922	33,297	86,405	266,831
1.00	93,763	53,926	33,151	86,029	266,869
	Change (Counterfactual – Actual)				
0	0	0	0	0	0
0.25	-443	1,004	-146	-377	38
0.50	-886	2,008	-291	-753	78
0.75	-1,329	3,012	-437	-1,130	116
1.00	-1,773	4,016	-583	-1,506	154
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-0.46%	2.01%	-0.43%	-0.43%	0.01%
0.50	-0.93%	4.02%	-0.86%	-0.86%	0.03%
0.75	-1.39%	6.03%	-1.30%	-1.29%	0.04%
1.00	-1.86%	8.05%	-1.73%	-1.72%	0.06%

Table 3.5: The top section has the counterfactual equilibrium quantities for each firm when Sprint receives 10% more spectrum, holding the spectrum of the other firms constant. The counterfactual equilibrium quantities equal the actual equilibrium quantities for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2010. All quantities are in thousands.

Prices When Sprint Receives 10% More Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	WgtAvg
	Counterfactual Prices				
0	$85.00	$70.00	$80.00	$90.00	$83.20
0.25	$84.84	$70.56	$79.86	$89.85	$83.12
0.50	$84.68	$71.13	$79.72	$89.69	$83.05
0.75	$84.53	$71.69	$79.59	$89.54	$82.99
1.00	$84.37	$72.25	$79.45	$89.38	$82.92
	Change (Counterfactual − Actual)				
0	$0	$0	$0	$0	$0
0.25	-$0.16	$0.56	-$0.14	-$0.15	-$0.08
0.50	-$0.32	$1.13	-$0.28	-$0.31	-$0.15
0.75	-$0.47	$1.69	-$0.41	-$0.46	-$0.21
1.00	-$0.63	$2.25	-$0.55	-$0.62	-$0.28
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-0.19%	0.80%	-0.18%	-0.17%	-0.09%
0.50	-0.38%	1.61%	-0.35%	-0.34%	-0.18%
0.75	-0.55%	2.41%	-0.51%	-0.51%	-0.26%
1.00	-0.74%	3.21%	-0.69%	-0.69%	-0.33%

Table 3.6: The top section has the counterfactual equilibrium prices for each firm, as well as the quantity-weighted average price ("WgtAvg"), when Sprint receives 10% more spectrum, holding the spectrum of the other firms constant. The counterfactual equilibrium prices equal the actual equilibrium prices for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2010.

Quantities With Equal Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	Total
	Counterfactual Quantities				
0	85,120	48,133	33,790	85,445	252,488
0.25	81,049	53,024	38,070	79,576	251,719
0.50	76,977	57,915	42,350	73,707	250,949
0.75	72,906	62,806	46,630	67,839	250,181
1.00	68,835	67,697	50,910	61,970	249,412
	Change (Counterfactual – Actual)				
0	0	0	0	0	0
0.25	-4,071	4,891	4,280	-5,869	-769
0.50	-8,143	9,782	8,560	-11,738	-1,539
0.75	-12,214	14,673	12,840	-17,606	-2,307
1.00	-16,285	19,564	17,120	-23,475	-3,076
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-4.78%	10.16%	12.67%	-6.87%	-0.30%
0.50	-9.57%	20.32%	25.33%	-13.74%	-0.61%
0.75	-14.35%	30.48%	38.00%	-20.61%	-0.91%
1.00	-19.13%	40.65%	50.67%	-27.47%	-1.22%

Table 3.7: The top section has the counterfactual equilibrium quantities for each firm when all firms receive an equal share of total spectrum. The counterfactual equilibrium quantities equal the actual equilibrium quantities for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium quantities to the counterfactual equilibrium quantities. The bottom section shows the percent change. Values shown use data from 2009. All quantities are thousands.

Prices With Equal Spectrum

ρ	AT&T	Sprint	T-Mobile	Verizon	WgtAvg
	Counterfactual Prices				
0	$70.00	$70.00	$60.00	$80.00	$72.05
0.25	$68.66	$72.85	$63.04	$77.80	$71.58
0.50	$67.32	$75.69	$66.08	$75.60	$71.47
0.75	$65.98	$78.54	$69.12	$73.41	$71.73
1.00	$64.64	$81.38	$72.16	$71.21	$72.35
	Change (Counterfactual – Actual)				
0	$0	$0	$0	$0	$0
0.25	-$1.34	$2.85	$3.04	-$2.20	-$0.46
0.50	-$2.68	$5.69	$6.08	-$4.40	-$0.57
0.75	-$4.02	$8.54	$9.12	-$6.59	-$0.31
1.00	-$5.36	$11.38	$12.16	-$8.79	$0.31
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-1.91%	4.07%	5.07%	-2.75%	-0.64%
0.50	-3.83%	8.13%	10.13%	-5.50%	-0.79%
0.75	-5.74%	12.20%	15.20%	-8.24%	-0.43%
1.00	-7.66%	16.26%	20.27%	-10.99%	0.42%

Table 3.8: The top section has the counterfactual equilibrium prices for each firm, as well as the quantity-weighted average price ("WgtAvg"), when all firms receive an equal share of total spectrum. The counterfactual equilibrium prices equal the actual equilibrium prices for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium prices to the counterfactual equilibrium prices. The bottom section shows the percent change. Values shown use data from 2009.

Prices When AT&T Transfers 10% of its Spectrum to T-Mobile

ρ	AT&T	Sprint	T-Mobile	Verizon	WgtAvg
	Counterfactual Prices				
0	\$85.00	\$80.00	\$80.00	\$90.00	\$85.07
0.25	\$83.93	\$80.26	\$81.42	\$90.28	\$84.99
0.50	\$82.87	\$80.51	\$82.84	\$90.55	\$84.94
0.75	\$81.80	\$80.77	\$84.25	\$90.83	\$84.94
1.00	\$80.73	\$81.02	\$85.67	\$91.10	\$84.98
	Change (Counterfactual − Actual)				
0	\$0	\$0	\$0	\$0	\$0
0.25	-\$1.07	\$0.26	\$1.42	\$0.28	-\$0.08
0.50	-\$2.13	\$0.51	\$2.84	\$0.55	-\$0.13
0.75	-\$3.20	\$0.77	\$4.25	\$0.83	-\$0.13
1.00	-\$4.27	\$1.02	\$5.67	\$1.10	-\$0.09
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-1.26%	0.33%	1.78%	0.31%	-0.10%
0.50	-2.51%	0.64%	3.55%	0.61%	-0.15%
0.75	-3.76%	0.96%	5.31%	0.92%	-0.15%
1.00	-5.02%	1.28%	7.09%	1.22%	-0.11%

Table 3.9: The top section has the counterfactual equilibrium prices for each firm, as well as the quantity-weighted average price ("WgtAvg"), when AT&T transfers 10% of its spectrum to T-Mobile. The counterfactual equilibrium prices equal the actual equilibrium prices for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2011.

Quantities When AT&T Transfers 10% of its Spectrum to T-Mobile

ρ	AT&T	Sprint	T-Mobile	Verizon	Total
	Counterfactual Quantities				
0	103,247	55,021	33,185	92,167	283,620
0.25	100,007	55,461	34,655	92,873	282,996
0.50	96,767	55,901	36,125	93,580	282,373
0.75	93,528	56,341	37,596	94,286	281,751
1.00	90,288	56,781	39,066	94,992	281,127
	Change (Counterfactual − Actual)				
0	0	0	0	0	0
0.25	-3,240	440	1,470	706	-624
0.50	-6,480	880	2,940	1,413	-1,247
0.75	-9,719	1,320	4,411	2,119	-1,869
1.00	-12,959	1,760	5,881	2,825	-2,493
	Percent Change				
0	0%	0%	0%	0%	0%
0.25	-3.14%	0.80%	4.43%	0.77%	-0.22%
0.50	-6.28%	1.60%	8.86%	1.53%	-0.44%
0.75	-9.41%	2.40%	13.29%	2.30%	-0.66%
1.00	-12.55%	3.20%	17.72%	3.07%	-0.88%

Table 3.10: The top section has the counterfactual equilibrium quantities for each firm when AT&T transfers 10% of its spectrum to T-Mobile. The counterfactual equilibrium quantities equal the actual equilibrium quantities for $\rho = 0$. The middle section has the absolute change when moving from the actual equilibrium to the counterfactual equilibrium. The bottom section shows the percent change. Values shown use data from 2011. All quantities are in thousands.

References

Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. Econometrica 58, 1041–1063.

Armstrong, M., 2006. Competition in two-sided markets. The RAND Journal of Economics 37, 668–691.

Bohnet, I., Kübler, D., 2005. Compensating the cooperators: Is sorting in the prisoner's dilemma possible? J. Econ. Beh. Organ. 56, 61–76.

Camerer, C., Loewenstein, G., Rabin, M. (Eds.), 2002. Advances in behavioral economics. Princeton University Press.

Caves, K.W., 2011. Quantifying price-driven wireless substitution in telephony. Telecommunications Policy 35, 984–998.

Ellison, G., 1994. Cooperation in the prisoner's dilemma with anonymous random matching. Rev. Econ. Stud. 61, 567–588.

Ely, J.C., Välimäki, J., 2002. A robust folk theorem for the prisoner's dilemma. J. Econ. Theory 102, 84–105.

Evans, D.S., 2003. The antitrust economics of multi-sided platform markets. Yale Journal on Regulation 20, 325.

Evans, D.S., Schmalensee, R., 2005. Paying with plastic: the digital revolution in buying and borrowing. The MIT Press.

Fudenberg, D., 2006. Advancing beyond *Advances in Behavioral Economics*. J. Econ. Lit. 44, 694–711.

Fudenberg, D., Levine, D., Maskin, E., 1994. The folk theorem with imperfect public information. Econometrica 62, 997–1039.

Fudenberg, D., Maskin, E., 1989. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 54, 533–554.

Gans, J.S., King, S.P., 2003. The neutrality of interchange fees in payment systems. The BE Journal of Economic Analysis & Policy .

Garcia-Swartz, D.D., Hahn, R.W., Layne-Farrar, A., 2006. The move toward a cashless society: A closer look at payment instrument economics. Review of Network Economics 5.

Ghosh, P., Ray, D., 1996. Cooperation in community interaction without information flows. Rev. Econ. Stud. 63, 491–519.

Guthrie, G., Wright, J., 2007. Competing payment schemes. The Journal of Industrial Economics , 37–67.

Hausman, J., 2000. Efficiency effects on the us economy from wireless taxation. National Tax Journal , 733–742.

Ingraham, A.T., Sidak, J.G., 2004. Do states tax wireless services inefficiently-evidence on the price elasticity of demand. Va. Tax Rev. 24, 249.

Kaiser, U., Wright, J., 2006. Price structure in two-sided markets: Evidence from the magazine industry. International Journal of Industrial Organization 24, 1–28.

Kandori, M., 1992. Social norms and community enforcement. Rev. Econ. Stud 59, 63–80.

Kranton, R.E., 1996. The formation of cooperative relationships. J. Law, Economics, and Organization 12, 214–233.

Lhost, J.R., Srinagesh, P., Sibley, D.S., 2012. The economics of surcharging credit card transactions. Working Paper .

Pinto, B.P., Sibley, D.S., 2013. Network congestion and the unilateral effects analysis of mergers. Working Paper .

Rochet, J.C., Tirole, J., 2002. Cooperation among competitors: Some economics of payment card associations. The Rand Journal of Economics , 549–570.

Rochet, J.C., Tirole, J., 2003a. An economic analysis of the determination of interchange fees in payment card systems. Review of Network Economics 2.

Rochet, J.C., Tirole, J., 2003b. Platform competition in two-sided markets. Journal of the European Economic Association 1, 990–1029.

Rochet, J.C., Tirole, J., 2006. Two-sided markets: a progress report. The RAND Journal of Economics 37, 645–667.

Röller, L.H., Waverman, L., 2001. Telecommunications infrastructure and economic development: A simultaneous approach. American Economic Review , 909–923.

Rysman, M., 2009. The economics of two-sided markets. The Journal of Economic Perspectives 23, 125–143.

Schwartz, M., Vincent, D.R., 2006. The no surcharge rule and card user rebates: Vertical control by a payment network. Review of Network Economics 5.

Weyl, E.G., 2010. A price theory of multi-sided platforms. The American Economic Review 100, pp. 1642–1672.

Wildfang, K.C., Marth, R.W., 2006. The persistence of antitrust controversy and litigation in credit card networks. Antitrust Law Journal 73, 675–707.

Wright, J., 2003. Optimal card payment systems. European Economic Review 47, 587–612.

Wright, J., 2004. The determinants of optimal interchange fees in payment systems. The Journal of Industrial Economics 52, 1–26.