Modelling Beliefs in Games with Generalized Preferences and Other Essays

A Dissertation
Presented to the Faculty of the Graduate School
Of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Rupa Athreya

Dissertation Director: Professor Stephen Morris

May 2002

UMI Number: 3046118

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Rupa Athreya
May 2002

In the first essay in this dissertation, "Modelling Beliefs in Games with Generalized Preferences," I bring together two strands of recent research in the decision theoretic foundations of games. One strand deals with the extension to games of models of behavior where players' preferences deviate from subjective expected utility maximization. The other strand investigates the epistemic foundations of non-expected utility behavior, studying the notions of belief generated from primitive consistency axioms on preferences. By exploring a minimal preference structure compatible with games from a subjective viewpoint, I provide a framework within which extensions of single-person decision theory to games can be compared and evaluated.

In the second essay, "Price Dispersion in the Wholesale Market for Steel," I look at the factors underlying price discrimination from an empirical perspective. I study the pricing decision of a durable commodity intermediary. The firm, a U.S. steel supplier, has provided daily data on purchases, sales, and inventory holdings on the products it carries. I find considerable dispersion in prices, both in the short term and over the long run. Apart from quantity discounts and local market power, the identities of the buyers affect the pricing decision of the seller, and there also seems to be a premium to information. The analysis suggests that there are substantial inefficiencies in the market for steel, and that buyers might benefit considerably from increased transparency in the market as a whole and lower barriers to search.

In the final essay, titled "Motivation, Perception and Morale," I explore some of the ways in which motivation depends on one's perception of the actions and intentions of others. I sketch a simple example in a two-person signalling game where an agent who cares intrinsically about the principal's perception of him extracts information from her actions. In the unique separating equilibrium, the principal is able to signal her beliefs about the agent's motivation through her choice of wage. Evidence indicates that the psychological effects of monetary rewards are important in their own right. A richer treatment of the inner lives of economic agents is therefore important to understanding a number of issues not explained fully by standard economic theory.

Contents

List of Tables vi
List of Figures vii

1 Modelling Beliefs in Games with Generalized Preferences **1**
1.1 Introduction 1
1.2 Preferences and Beliefs 4
1.2.1 Preference preliminaries 4
1.2.2 Beliefs, possibility and supports 8
1.3 Interactive beliefs 16
1.4 Beliefs in Games 21
1.4.1 Strict Dominance 22
1.4.2 Admissibility 26
1.5 Equilibrium 30
1.6 Conclusion 37

2 Price Dispersion in the Wholesale Market for Steel **38**
2.1 Introduction 38
2.2 The Steel Market 45
2.3 Data 50
2.3.1 Sales data 50
2.3.2 Customer information 55
2.4 Sources of Price Dispersion 57
2.5 Conclusion 67

3 Motivation, Perception and Morale **79**
3.1 Introduction 79
3.2 Theories and Evidence 82
3.3 A Simple Example 89
3.4 The Economics of Motivation 95
3.5 Conclusion 101

Bibliography **102**

List of Tables

2.1 A Snapshot View of Daily Price Dispersion for Product P2 41
2.2 Summary Statistics for Frequent Buyers, Product P2 43
2.3 Mill Shipments (million net tons) 46
2.4 Imports by Country of Origin (million net tons) 46
2.5 Locations of Top 100 Metal Service Centers in the U.S. and Canada, 1997, percentage of sites 47
2.6 Imports by Customs District (million net tons) 47
2.7 Mill Shipments to Major Markets (million net tons) 48
2.8 Value Added by Steel Service Centers 48
2.9 Cost of Possession Worksheet 49
2.10 Summary Statistics for Sales of Product P2 54
2.11 Major SIC code groups, Product P2 56
2.12 Customer Firm Size by Number of Employees 56
2.13 OLS with Time Dummies 58
2.14 OLS with Customer Characteristics and Time Dummies 59
2.15 OLS with Distance and Customer Characteristics, Product P2 61
2.16 OLS with Distance and Customer Characteristics, Product P2 62
2.17 OLS with Distance and Location Dummies, Product P2 63
2.18 Summary Statistics for Frequent Buyers, Group P Products 65

List of Figures

2.1 Prices and Costs for Product P2, July 1997 to December 1998 69
2.2 Prices and Costs for Product P2, December 1998 to June 2000 70
2.3 Prices and Margins for Product P2, July 1997 to July 1999 71
2.4 Prices and Margins for Product P2, July 1999 to June 2000 72
2.5 Distribution of Net Margins for Product P2 73
2.6 Distribution of Sales by Quantity, Product P2 74
2.7 Distribution of Customers by Distance, Product P2 75
2.8 Distribution of Sales by Distance, Product P2 76
2.9 Summary Statistics for Frequent Buyers by Average Quantity Purchased . . 77
2.10 Summary Statistics for Frequent Buyers by Distance from Company S . . . 78

Acknowledgements

As all dissertations do, this one owes a great deal to a number of people. Stephen Morris encouraged me to pursue my interests in the foundations of games, and proved a patient listener. George Hall and John Rust were very supportive of my attempts at working with real data. Chris Timmins found the time for many helpful conversations. David Pearce has been a constant source of ideas, support and friendship. My biggest intellectual debt is to Ben Polak, for innumerable conversations and a willingness to help at every stage; Ben's unwavering encouragement is responsible in no small part for this dissertation.

I am grateful to my colleagues Hiu Man Chan, Katja Seim and Sofia Sheidvasser for teaching me much of the economics I know and for being there at crucial moments. Outside the department, John Anderson, Jennifer Arzt, Martin Benjamin and Lynn Jones have broadened my horizons while being friends and anchors in this exciting (and occasionally overwhelming) journey. Special thanks to Jennie and Lynn for staying to the bitter end. Pam O'Donnell and Dee Smitkowski provided invaluable friendship throughout my time at Yale: my graduate school experience would have much less fun without their particular perspective on everyday events in the life of the department. Many thanks also to Shareen Joshi for uplifting tomfoolery, and to Gracy for more coffee than I should have drunk.

My twin, Rishi Raj, has kept me on my toes with his lively scepticism, while always conveying his unspoken support. Lastly, although no words will really suffice, I'd like to thank my parents, Geeta and Mrityunjay Athreya. Their love and wisdom have been my greatest source of strength, and they continue to embody an ideal for me to strive towards. This dissertation is dedicated to them.

Chapter 1

Modelling Beliefs in Games with Generalized Preferences

1.1 Introduction

Traditional decision theory treats risk (situations where probabilities of events are known) as formally equivalent to uncertainty (situations in which the probabilities of events are unknown). The subjective expected utility (SEU) model axiomatized by Savage (1954) has been the most important theory in analysing decision making under uncertainty. Expected utility (EU) has been the accepted basis for analysing decision making in games as well. Using the EU model to represent players' preferences, a number of solution concepts have been developed, most prominent among them being Nash Equilibrium.

However, the descriptive validity of the SEU model has been questioned. In recent years generalizations of, and alternatives to, the SEU framework have been developed, where the decision maker does not have point beliefs that can be represented as in Savage. A growing literature has attempted to extend this theory to interactive situations[1]. Players in games are modelled as facing uncertainty regarding opponents' strategies and the standard solution concepts are generalized to account for deviations from subjective expected utility

[1]See, for example, Dow and Werlang (1994), Eichberger and Kelsey (1998), Epstein (1997), Klibanoff (1996), Lo (1996), Marinacci (2000), Mukerji (1997), Ryan (2000).

maximization.

Savage (1954) shows how properties of a decision maker's probabilistic beliefs can be deduced from primitive consistency axioms on preferences. Different restrictions on conditional preferences generate different notions of qualitative belief. Under the SEU axioms, these are equivalent, so that the "subjective probability" generated by SEU preferences has a natural epistemic interpretation: belief can be identified with "full belief" or belief with probability one. However, when individuals are not expected utilily maximizers, the alternative notions of belief generated by assumptions on preferences are not necessarily equivalent. That is, qualitative notions of belief arising out of different restrictions on preferences vary in their implications. For example, when players' preferences belong to the maxmin expected utility (MEU) class, "full belief" can be thought of as corresponding to a situation where all beliefs in the agent's belief set place probability one on the event that is believed; weaker notions of belief might relax this notion to require, for instance, that beliefs only place some positive weight on the event that is believed, or that events which are not believed are assigned zero weight by at least one belief, but not necessarily by all of them.

Morris (1997) provides a general framework for preference-based belief which allows comparison of different notions of belief in order to determine which notions are 'stronger' or 'weaker' than others. He studies belief operators, which provide a semantic characterization of belief by identifying, for each state of the world, the events that the decision maker believes. Belief operators defined from the decision maker's preferences can be characterized using standard preference representations, and provide epistemic models for decision making in both expected utility and non-expected utility environments.

When different notions of belief are no longer equivalent, the choice of which notion of belief to employ becomes critical in extending decision theory to games. Alternative representations of preferences, and notions of belief, will mean that restrictions on the rationality of opponents lead to different epistemic conditions for solution concepts.

In this chapter, I bring together these developments in the decision-theoretic foundations of games by setting out a general framework for characterising dominance and equilibrium in normal form games with players whose preferences may deviate from the expected utility model. The standard framework considers players with SEU preferences; knowledge is defined as belief with probability one, and epistemic conditions for dominance and equilibrium are derived. In this chapter, I relax the axioms on preferences, allowing them to deviate from expected utility, and investigate the implications of varying belief of rationality for dominance and equilibrium.

I first show that common belief of rationality, where belief is defined in the sense of Savage, and weak monotonicity assumptions on preferences, lead players to play pure strategies that are iteratively undominated by other pure strategies. However, when the notion of belief is weakened, then it is not always possible to rule out play of iteratively strictly dominated strategies. I also show that when preferences are admissible, weakly dominated strategies can be eliminated, but that iterated weak deletion does not follow from common belief of admissibility. In a closely related paper, Epstein (1997) considers rationalizability and equilibrium when players' preferences deviate from expected utility maximization. However, he retains Savage belief and investigates the implications of varying the notion of rationality, relating a general notion of rationalizability to the standard model of expected utility and to pure strategy dominance. The exercise carried out in this chapter is to relate different notions of belief and to compare the implications of common belief of rationality, appropriately defined, across representations of preferences.

I also investigate epistemic conditions for pure-strategy Nash equilibrium, defined to be a pair of pure-strategy best responses. In the standard model, mutual belief of actions implies that actions form a Nash equilibrium (Aumann and Brandenburger 1995); I show first that with general preferences, mutual Savage belief of actions is sufficient for actions to be a Nash equilibrium, but that when belief is weakened to "weak" belief, as defined in Morris (1997), there exists a class of preferences for which mutual belief of actions does not imply that players play a Nash equilibrium.

The chapter is organized as follows. In section 2, I review briefly the decision-theoretic framework and recent extensions. In section 3, I extend this to two-player situations. I define an interactive belief system and prove some preliminary results regarding properties of interactive beliefs. Section 4 considers pure strategy dominance, and characterizes conditions for iterative strict and weak dominance. Section 5 explores epistemic conditions for equilibrium with generalized preferences.

1.2 Preferences and Beliefs

1.2.1 Preference preliminaries

The traditional approach to modelling choice under uncertainty in economics and decision theory combines an exogenous framework for modelling belief with a set of preferences. However, a given set of preferences over acts with state-contingent outcomes, and beliefs over those states will not in general be reconcilable without further restrictions. Hence a more natural approach in this context is that due to Savage (1954), where individual beliefs are derived from restrictions on preferences, so that beliefs can be thought of as being "revealed" by behavior rather than imposed *a priori* by the modeller.

I begin by outlining a framework for preferences, and then defining belief in those terms. Let Ω be a finite set of states of the world. Let the set of "acts" be the set of all functions from Ω to $\mathbf{R}$, the real line, so that acts can be thought of as vectors in $\mathbf{R}^\Omega$. Thus act $x \in \mathbf{R}^\Omega$ yields the prize x_ω in state $\omega \in \Omega$. For $E \subseteq \Omega$, x_E denotes the tuple $\{x_\omega\}_{\omega \in E}$. I denote by $-E$ the complement of E in Ω and use the notation $x_{\{\omega\}}$ or x_ω and $x_{-\{\omega\}}$ or $x_{-\omega}$ interchangeably. At each state of the world, ω, an individual is assumed to have a preference relation, $\succeq_\omega$, over acts, where $x \succeq_\omega y$ implies that, at ω, act x is at least as good as act y. Strict preference and indifference are defined in the usual fashion.

In the standard EU framework, an individual's beliefs are represented by a single additive probability measure on the set of states. Preference relations $\{\succeq_\omega\}_{\omega \in \Omega}$ on $\mathbf{R}^\Omega$ have an EU representation if there exists, for each state $\omega \in \Omega$, an increasing and continuous utility function $u_\omega : \mathbf{R} \to \mathbf{R}$ and an additive probability measure p_ω on Ω such that, for all $x, y \in \mathbf{R}^\Omega$,

$$x \succeq_\omega y \Leftrightarrow \sum_{\omega' \in \Omega} p_\omega(\omega') u_\omega(x_{\omega'}) \geq \sum_{\omega' \in \Omega} p_\omega(\omega') u_\omega(y_{\omega'}).$$

Two alternative representations of preferences which relax this restriction are the *maxmin expected utility* (MEU) framework and the *Choquet expected utility* (CEU) model. Preference relations $\{\succeq_\omega\}_{\omega \in \Omega}$ on $\mathbf{R}^\Omega$ have an MEU representation if there exists, for each state

$\omega \in \Omega$, an increasing and continuous utility function $u_\omega : \mathbf{R} \to \mathbf{R}$ and a unique, nonempty, closed and convex conditional belief set C_ω of additive probability measures on Ω such that, for all $x, y \in \mathbf{R}^\Omega$,

$$x \succeq_\omega y \Leftrightarrow \min_{p_\omega \in C_\omega} \left\{ \sum_{\omega' \in \Omega} p_\omega(\omega') u_\omega(x_{\omega'}) \right\} \geq \min_{p_\omega \in C_\omega} \left\{ \sum_{\omega' \in \Omega} p_\omega(\omega') u_\omega(y_{\omega'}) \right\}.$$

In the case of CEU preferences, the probability measure is non-additive. A non-additive probability measure (sometimes known as a *capacity*) is a function $\nu : 2^\Omega \to \mathbf{R}_+$, with $\nu(\emptyset) = 0$, $\nu(\Omega) = 1$ and $\nu(E) \leq \nu(F)$ if $E \subseteq F$. The *range* of act x is the set of values attained on Ω. Let T be any finite ordered subset of the real line containing the range of x, so that $\{r \in \mathbf{R} | \; r = x_\omega \text{ for some } \omega \in \Omega\} \subseteq T = \{r_1, \cdots, r_K\} \subseteq \mathbf{R}$ and $r_1 > r_2 > \cdots > r_K$. Letting $r_{K+1} = 0$, the expected value of x is

$$\mathbf{E}_\nu(x) = \sum_{k=1}^{K} (r_k - r_{k+1}) \; \nu(\{\omega \in \Omega | x_\omega \geq r_k\}).$$

This definition reduces to the usual notion of expected value if ν is additive. Write $u_\omega(x)$ for the vector $\{u_\omega(x_\omega)\}_{\omega \in \Omega}$. Preference relations $\{\succeq_\omega\}_{\omega \in \Omega}$ on $\mathbf{R}^\Omega$ have a CEU representation if there exists, for each state $\omega \in \Omega$, an increasing and continuous utility function $u_\omega : \mathbf{R} \to \mathbf{R}$ and a non-additive probability measure on ν_ω such that, for all $x, y \in \mathbf{R}^\Omega$,

$$x \succeq_\omega y \Leftrightarrow \mathbf{E}_{\nu_\omega}(u_\omega(x)) \geq \mathbf{E}_{\nu_\omega}(u_\omega(y)).$$

Lexicographic preferences allow decision makers to take into account events which are assigned probability zero *ex ante*. Say that x is lexicographically greater than y [$x \geq_L y$] if $y_i > x_i$ implies $x_h > y_h$ for some $h < i$. Preference relations $\{\succeq_\omega\}_{\omega \in \Omega}$ on $\mathbf{R}^\Omega$ have a lexicographic expected utility (LEU) representation if there exists (1) a positive integer J, (2) for each state $\omega \in \Omega$, an increasing and continuous utility function $u_\omega : \mathbf{R} \to \mathbf{R}$ and (3) for each $\omega \in \Omega$ and $j = 1, \cdots, J$, an additive probability measure p^j_ω on Ω such that, for all $x, y \in \mathbf{R}^\Omega$,

$$x \succeq_\omega y \Leftrightarrow \left\{ \sum_{\omega' \in \Omega} p^j_\omega(\omega') u_\omega(x_{\omega'}) \right\}_{j=1}^J \geq_L \left\{ \sum_{\omega' \in \Omega} p^j_\omega(\omega') u_\omega(y_{\omega'}) \right\}_{j=1}^J .$$

The following restrictions on preferences include conditions that ensure that the preference relation is a non-trivial ordering. I list them here for convenience, and will refer back to them as the occasion demands.

P1 : Preference relations are *reflexive* if $x \succeq_\omega x$ for all $x \in \mathbf{R}^\Omega$ and $\omega \in \Omega$.

P2 : Preference relations are *non-trivial* if not $x \succeq_\omega y$ for some $x, y \in \mathbf{R}^\Omega$ and all $\omega \in \Omega$.

P3 : Preference relations are *transitive* if $x \succeq_\omega y$ and $y \succeq_\omega z \Rightarrow x \succeq_\omega z$ for all $x, y, z \in \mathbf{R}^\Omega$ and all $\omega \in \Omega$.

P4 : Preference relations are *complete* if $x \succeq_\omega y$ or $y \succeq_\omega x$ for all $x, y \in \mathbf{R}^\Omega$ and $\omega \in \Omega$.

There are various monotonocity assumptions that may be made. Say that $x \geq y$ if $x_\omega \geq y_\omega$ for all $\omega \in \Omega$; $x > y$ if $x \geq y$ and $x_\omega > y_\omega$ for some $\omega \in \Omega$; $x \gg y$ if $x_\omega > y_\omega$ for all $\omega \in \Omega$.

P5 : Preference relations are *admissible* if $x > y \Rightarrow x \succ_\omega y$ and $x \geq y \Rightarrow x \succeq_\omega y$, for all $\omega \in \Omega$.

A weaker assumption, which is implied by [P5], is the following:

P5* : Preference relations are *monotone* if $x \gg y \Rightarrow x \succ_\omega y$ and $x \geq y \Rightarrow x \succeq_\omega y$, for all $\omega \in \Omega$.

P6 : Preference relations satisfy *non-null statewise monotonicity* if, for all $\omega, \omega' \in \Omega$, *either* $(x_{\omega'}, z_{-\omega'}) \sim_\omega (y_{\omega'}, z_{-\omega'})$ for all $x, y, z \in \mathbf{R}^\Omega$ *or* $(x_{\omega'}, z_{-\omega'}) \succ_\omega (y_{\omega'}, z_{-\omega'})$ for all $x \gg y, z$.

P7 : Preference relations are *continuous* if, for all for all $\omega \in \Omega$, the set $\{x \in \mathbf{R}^{\Omega} \mid x \succeq_{\omega} y\}$ is closed.

1.2.2 Beliefs, possibility and supports

Under certain conditions, results about belief on a finite state space can be represented by two equivalent formalisms: belief operators and possibility relations. A belief operator specifies, for each subset of the state space (or event), the set of states where the individual believes that event is true. A belief operator is thus a mapping $B : 2^{\Omega} \rightarrow 2^{\Omega}$, with the interpretation that the individual believes event $E \subseteq \Omega$ at state $\omega \in \Omega$ if and only if $\omega \in B(E)$. Belief operators are usually required to satisfy four logical properties[2]. For all $E, F \subseteq \Omega$:

B1 : $B(\Omega) = \Omega$.

B2 : $B(\emptyset) = \emptyset$.

B3 : $B(E) \cap B(F) \subseteq B(E \cap F)$.

B4 : $E \subseteq F \Rightarrow B(E) \subseteq B(F)$.

A possibility correspondence specifies, for each state of the world, which states the individual thinks are possible: it is a mapping from the uncertainty space to its subsets, $P : \Omega \rightarrow 2^{\Omega}$. The interpretation is that if $\omega' \in P(\omega)$, the state ω' is thought possible when the true state is ω. Under [B1], [B3] and [B4], there is a one to one relation between belief operators and possibility correspondences:

Definition 1 *The operator* $B : 2^{\Omega} \rightarrow 2^{\Omega}$ represents *the possibility relation* P *if*

[2]Note that [B3] and [B4] together imply that [B3] holds with equality.

$$B(E) = \{\omega \in \Omega | \ P(\omega) \subseteq E\}$$

$$P(\omega) = \bigcap\{F \subseteq \Omega | \ \omega \in B(F)\}$$

The operator B is a normal *belief operator if there exists P such that B represents P.*

Morris (1997, Lemma 1) shows that adding further restrictions yields the richer semantic structure for belief typically assumed by economists, which requires that possibility correspondences partition the state space.

B5 : $B(E) \subseteq B(B(E))$.

B6 : $-B(E) \subseteq B(-B(E))$.

These two conditions, when encountered in the knowledge literature, are usually called, respectively, *Positive Introspection* and *Negative Introspection* (or sometimes *knowing that you know* and *knowing that you don't know*). Together with [B1]-[B4], they imply that the possibility correspondence forms a partition of the state space. Note that in this context, this does not imply that the decision maker's beliefs are correct, only that his beliefs are internally consistent.

We are now in a position to define belief in terms of preferences. Savage's notion of belief captures the idea that if the individual's preferences never depend on anything that happens when event E does not occur, then the individual believes E. If the individual is ever concerned about what happens when E does not occur, then the individual cannot believe E. This is closely related to Savage's notion of *null* events: an individual (Savage) believes an event E if the complement of E is Savage-null.

Definition 2 *Belief operator B^{**} represents* Savage belief *of preference relations $\{\succeq_\omega\}_{\omega\in\Omega}$*

if

$$B^{**}(E) = \{\omega \in \Omega | (x_E, z_{-E}) \succeq_\omega (x_E, y_{-E}) \;\; \forall x, y, z \in \mathbf{R}^\Omega\}.$$

The associated possibility correspondence is defined as:

$$P^{**}(\omega) = \{\omega' \in \Omega | \exists x, y, z \in \mathbf{R}^\Omega \text{ such that } (x_{\omega'}, z_{-\omega'}) \succ_\omega (y_{\omega'}, z_{-\omega'})\}.$$

This notion of belief arises naturally in the standard expected utility model, where an event is believed if and only if it is assigned probability one by the decision maker's (subjective) probability measure. It is also possible to define alternative notions of belief.

Definition 3 *Belief operator* $\tilde{B}$ *represents* strong belief *of preference relations* $\{\succeq_\omega\}_{\omega\in\Omega}$ *if*

$$\tilde{B}(E) = \{\omega \in \Omega | (x_E, z_{-E}) \succeq_\omega (y_E, v_{-E}) \;\; \forall x \gg y, x, y, v, z \in \mathbf{R}^\Omega\}.$$

The associated possibility correspondence is defined as:

$$\tilde{P}(\omega) = \{\omega' \in \Omega | \exists x \gg y \gg z \in \mathbf{R}^\Omega \text{ such that } (x_{\omega'}, z_{-\omega'}) \succ_\omega y\}.$$

Strong belief formalizes the intuition that strict dominance on the event that is believed should be sufficient to determine the direction of preference, regardless of the outcome on the rest of the space. When preferences have a lexicographic expected utility (LEU) representation, an event is strongly believed if it is assigned probability one by the decision maker's first-order beliefs. Savage belief, by contrast, requires that *each* of the decision maker's beliefs assign probability one to the event that is believed.

We may also consider the following strengthening of strong belief, which I will call strong* belief.

$$\tilde{B}^*(E) = \{\omega \in \Omega | (x_E, z_{-E}) \succ_\omega (y_E, v_{-E}) \;\; \forall x \gg y, x, y, v, z \in \mathbf{R}^\Omega\}.$$

The associated possibility correspondence then becomes:

$$\tilde{P}^*(\omega) = \{\omega' \in \Omega | \exists x \gg y \gg z \in \mathbf{R}^\Omega \text{ such that } (x_{\omega'}, z_{-\omega'}) \succsim_\omega y\}.$$

A notion of belief which utilizes a much stronger notion of possibility is weak belief.

Definition 4 *Belief operator B represents* weak belief *of preference relations $\{\succsim_\omega\}_{\omega\in\Omega}$ if*

$$B(E) = \{\omega \in \Omega |\ P(\omega) \subseteq E\}.$$

where

$$P(\omega) = \{\omega' \in \Omega | (\forall x \gg y)(\exists z \ll y) \text{ such that } (x_{\omega'}, z_{-\omega'}) \succ_\omega y\}.$$

When preferences have an MEU representation, a state is weakly possible if it is assigned positive probability by *each* of the probability measures in the decision maker's set of conditional beliefs. An event is weakly believed if it contains every state assigned positive probability by all the measures in the belief set. By contrast, Savage or strong belief would require that the event believed is assigned probability one by *all* measures in the set of beliefs, with a state being considered possible if there exists a measure in the belief set which assigns positive probability to it.

Note that weak belief is normal by definition, so that it satisfies [B1], [B3] and [B4], but that it might fail [B2]. If the possibility correspondence $P(\omega)$ is empty for some ω, then $\omega \in B(\emptyset)$. As Morris (1997, Example 3) shows, this can happen for MEU preferences if the intersection of the supports of the decision-maker's beliefs is empty.

Dow and Werlang (1994) provide a notion of the support of a convex capacity in a CEU representation. Say that the set $A \subseteq \Omega$ is a *DW-support* of ν if A is $\subseteq$-minimal with respect to the property that $\nu(-A) = 0$; that is, $\nu(-A) = 0$ and $\nu(-B) > 0$ whenever B is a proper subset of A. Ryan (2000) defines a notion of belief intermediate between strong and weak belief, motivated by the notion of a DW-support.

Definition 5 *Belief operator $B^{\vee}$ represents* firm belief *of preference relations $\{\succeq_\omega\}_{\omega\in\Omega}$ if*

$$B^{\vee}(E) = \{\omega \in \Omega | (\forall \omega' \in -E)(\forall F \subseteq \Omega)(\exists x, y, z \in \mathbf{R}, x > y)$$
$$\text{such that } (z_{F\backslash\omega'}, y_{-F\cup\omega'}) \succeq_\omega (x_F, y_{-F})\}.$$

$$P^{\vee}(\omega) = \{\omega' \in \Omega | \exists E \subseteq -\{\omega'\} \text{ such that } (x_{-E}, y_E) \succ_\omega (z_{-E\backslash\omega'}, y_{E\cup\omega'})$$
$$\forall x, y, z \in \mathbf{R} \text{ with } x > y\}.$$

Ryan shows that the possibility correspondence $P^{\vee}(\cdot)$ is equivalent to the union of DW-supports of the convex capacity in a CEU representation, and that the weak possibility correspondence equals the *intersection* of DW-supports in the CEU representation.

Example 1 *(1): For CEU preferences $\{\succeq_\omega\}_{\omega\in\Omega}$ with representation (u_ω, ν_ω), the firm belief operator takes the form*

$$B^{\vee}(E) = \{\omega \in \Omega | (\forall \omega' \in -E)(\forall F \subseteq \Omega)\ \nu_\omega(F\backslash\omega') = 0 \Rightarrow \nu_\omega(F) = 0\},$$

$$P^{\vee}(\omega) = \{\omega' \in \Omega | \exists E \subseteq -\{\omega'\} \text{ such that } \nu_\omega(-E) > 0 \text{ and } \nu_\omega(-E\backslash\omega') = 0\}.$$

In particular, $\omega \in B^{\vee}(E)$ implies that $\forall \omega' \in -E, \nu_\omega(-\{\omega'\}) > 0$.

(2): When preferences $\{\succeq_\omega\}_{\omega\in\Omega}$ have an MEU representation (u_ω, C_ω),

$$B^{\vee}(E) = \{\omega \in \Omega \mid (\forall \omega' \in -E)(\forall F \subseteq \Omega)\ p_\omega(F\backslash\omega') = 0 \text{ for some } p_\omega \in C_\omega$$
$$\Rightarrow p_\omega(F) = 0 \text{ for some } p_\omega \in C_\omega\},$$

$$P^{\vee}(\omega) = \{\omega' \in \Omega \mid \exists E \subseteq -\{\omega'\} \text{ such that } p_\omega(-E) > 0 \ \forall p_\omega \in C_\omega$$
$$\text{and } p_\omega(-E \backslash \omega') = 0 \text{ for some } p_\omega \in C_\omega\}.$$

In particular, $\omega \in B^{\vee}(E)$ *implies that* $\forall \omega' \in -E$, $p_\omega(\omega') = 0$ *for some* $p_\omega \in C_\omega$ *and* $p_\omega(-\{\omega'\}) > 0 \ \forall p_\omega \in C_\omega$. *Also,* $\omega' \in P^{\vee}(\omega)$ *implies that* $p_\omega(\omega') > 0$ *for some* $p_\omega \in C_\omega$, *so that* $P^{\vee}(\omega) \subseteq P^{**}(\omega) = \tilde{P}(\omega)$.

Proof: (1): Denote the proposed operator by $B^f(E)$, and suppose that $\omega \in B^f(E)$. Fix $\omega' \in -E$ and $F \subseteq \Omega$. Let $x, y \in \mathbf{R}$ with $x > y$. Since $\nu_\omega(F) > 0 \Rightarrow \nu_\omega(F \backslash \omega') > 0$, for any

$$z \geq \frac{\nu_\omega(F)x + (\nu_\omega(F \backslash \omega')) - \nu_\omega(F)y}{\nu_\omega(F \backslash \omega')}$$

we have

$$(z_{F \backslash \omega'}, y_{-F \cup \omega'}) \succeq^u_\omega (x_F, y_{-F}),$$

and hence $\omega \in B^{\vee}(E)$.

Conversely, if $\exists F \subseteq \Omega$ such that $\nu_\omega(F) > 0$ but $\nu_\omega(F \backslash \omega') = 0$, then for any $x, y, z \in \mathbf{R}, x > y$,

$$(x_F, y_{-F}) \succ^u_\omega (z_{F \backslash \omega'}, y_{-F \cup \omega'}),$$

so that $\omega \notin B^{\vee}(E)$. The equivalence of $P^{\vee}(\cdot)$ and $\hat{P}(\cdot)$ follows from Ryan (1999, Proposition 3.6).

(2): Suppose $\omega \in B^f(E)$. Fix $\omega' \in -E$ and $F \subseteq \Omega$. Consider $x, y, z \in \mathbf{R}$ with $x > y$, and let $f = (z_{F \backslash \omega'}, y_{-F \cup \omega'})$ and $g = (x_F, y_{-F})$. Then

$$\mathbf{u}_\omega(g) = \min_{p_\omega \in C_\omega} [y\, p_\omega(-F) + x\, p_\omega(F)] = \min_{p_\omega \in C_\omega} [y - y\, p_\omega(F) + x\, p_\omega(F)]$$

$$= y + \min_{p_\omega \in C_\omega} [(x - y)\, p_\omega(F)],$$

and

$$\mathbf{u}_\omega(f) = \min_{p_\omega \in C_\omega} \left[z\, p_\omega(F\backslash\omega') + y\, p_\omega(-F \cup \omega')\right] = \min_{p_\omega \in C_\omega} \left[z\, p_\omega(F\backslash\omega') + y - y\, p_\omega(F\backslash\omega')\right]$$

$$= y + \min_{p_\omega \in C_\omega} [(z - y)\, p_\omega(F\backslash\omega')].$$

If $\min_{p_\omega \in C_\omega} p_\omega(F\backslash\omega') = 0$, then $\mathbf{u}_\omega(f) = y$. But $\omega \in B^f(E)$, so that $\min_{p_\omega \in C_\omega} p_\omega(F) = 0$; then $\mathbf{u}_\omega(g) = y$. Hence, for $x, y \in \mathbf{R}$ with $x > y$, we have, for any $z \in \mathbf{R}$,

$$(z_{F\backslash\omega'}, y_{-F\cup\omega'}) \sim^{u}_{\omega} (x_F, y_{-F}),$$

and hence $\omega \in B^{\vee}(E)$.

If $p_\omega(F) > 0\ \forall p_\omega \in C_\omega$, then $p_\omega(F\backslash\omega') > 0\ \forall p_\omega \in C_\omega$, and

$$\mathbf{u}_\omega(f) \geq \mathbf{u}_\omega(g)$$

$$\Leftrightarrow y + \min_{p_\omega \in C_\omega} [(z - y)\, p_\omega(F\backslash\omega')] \geq y + \min_{p_\omega \in C_\omega} [(x - y)\, p_\omega(F)].$$

Then for any

$$z \geq \frac{\min_{p_\omega \in C_\omega} p_\omega(F)x + (\min_{p_\omega \in C_\omega} p_\omega(F\backslash\omega') - \min_{p_\omega \in C_\omega} p_\omega(F))y}{\min_{p_\omega \in C_\omega} p_\omega(F\backslash\omega')},$$

we have

$$(z_{F\backslash\omega'}, y_{-F\cup\omega'}) \succeq^{u}_{\omega} (x_F, y_{-F}),$$

which implies $\omega \in B^{\vee}(E)$.

Conversely, if $\omega \notin B^f(E)$, so that $\exists F \subseteq \Omega$ such that $p_\omega(F) > 0\ \forall p_\omega \in C_\omega$ but $p_\omega(F\backslash\omega') = 0$ for some $p_\omega \in C_\omega$, then for any $x, y, z \in \mathbf{R}, x > y$,

$$\mathbf{u}_\omega(g) = y + \min_{p_\omega \in C_\omega} [(x - y)\, p_\omega(F)] > \mathbf{u}_\omega(f) = y,$$

so that

$$(x_F, y_{-F}) \succ^u_\omega (z_{F\backslash\omega'}, y_{-F\cup\omega'}).$$

Hence $\omega \notin B^\vee(E)$.

Next, I show that $\omega \in B^\vee(E)$ if, and only if, $P^\vee(\omega) \subseteq E$. For the only if part, suppose $\omega \in B^\vee(E)$; I show that $P^\vee(\omega) \subseteq E$. Suppose not, so that $\exists\, \omega' \in P^\vee(\omega)$ such that $\omega' \notin E$. Then $\exists\, F \subseteq -\{\omega'\}$ such that $p_\omega(-F) > 0\ \forall\, p_\omega \in C_\omega$ and $p_\omega(-F\backslash\omega') = 0$ for some $p_\omega \in C_\omega$. But then $\omega \notin B^\vee(E)$.

Conversely, let $P^\vee(\omega) \subseteq E$, where $E \subset \Omega$. Fix $\omega' \in -E$ and $F \subseteq \Omega$. Suppose that $\min_{p_\omega \in C_\omega} p_\omega(F\backslash\omega') = 0$. Then we need to show that $\min_{p_\omega \in C_\omega} p_\omega(F) = 0$. If $\omega' \notin F$, then $\min_{p_\omega \in C_\omega} p_\omega(F\backslash\omega') = 0 \Rightarrow \min_{p_\omega \in C_\omega} p_\omega(F) = 0$, and $\omega \in B^\vee(E)$. If $\omega' \in F$, then $-F \subseteq -\{\omega'\}$; further, since $\omega' \notin P^\vee(\omega)$, it follows that $\forall\, G \subseteq \Omega$, if $\min_{p_\omega \in C_\omega} p_\omega(-G\backslash\omega') = 0$, then $\min_{p_\omega \in C_\omega} p_\omega(-G) = 0$. Letting $-F = G$ completes the proof.$\diamond$

In general (Morris 1997, Ryan 1999),

$$P(\omega) \subseteq P^\vee(\omega) \subseteq \tilde{P}(\omega) \subseteq P^{**}(\omega),$$

and therefore

$$B^{**}(E) \subseteq \tilde{B}(E) \subseteq B^\vee(E) \subseteq B(E).$$

The logical properties of belief can be related to restrictions on preferences. Morris (1997) shows that if preferences satisfy [P1]-[P4], Savage belief satisfies [B1]-[B4] and that if preferences satisfy [P1]-[P4] and [P5*], strong belief satisfies [B1]-[B4]. In addition, if preferences satisfy [P1]-[P4], [P5*] and [P7], then Savage and strong belief are equivalent.[3] If

[3]Proof: Let $\omega' \in P^{**}(\omega)$, so that $(x_{\omega'}, z_{-\omega'}) \succ_\omega (y_{\omega'}, z_{-\omega'})$ for some $x, y, z \in \mathbb{R}^\Omega$. By [P1] and [P5*], $x_{\omega'} \geq y_{\omega'}$. Letting $\hat{y} = (y_{\omega'}, z_{-\omega'})$, we then have $(x_{\omega'}, \hat{y}_{-\omega'}) \succ_\omega \hat{y}$. For $\hat{x}$ such that $\hat{x}_{\omega'} > x_{\omega'}$ and $\hat{x}_{-\omega'} \gg z_{-\omega'}$, by [P5*] and [P3], $(\hat{x}_{\omega'}, \hat{y}_{-\omega'}) \succ_\omega \hat{y}$. Then by [P7], there exists $\hat{z}$, with $\hat{y} \gg \hat{z}$, such that $(\hat{x}_{\omega'}, \hat{z}_{-\omega'}) \succ_\omega \hat{y}$. Hence $\omega' \in P(\omega)$.

preferences satisfy [P6] and [P7], then Savage, strong*, strong and weak belief are equivalent. I assume from now on that preferences satisfy [P1]-[P4].

1.3 Interactive beliefs

I now extend this framework to situations with two agents. I consider 2-player, finite action games, consisting of a finite set of players $I = \{1, 2\}$, finite pure strategy spaces S_i $(i = 1, 2)$, with typical element s_i, and payoff functions $g_i : S_1 \times S_2 \to \mathbf{R}$. Following Aumann and Brandenburger (1995), I define an interactive belief system as follows. Each player has a finite set of types, T_i, and

- for each type t_i, $i \neq j$, a preference relation $\succeq_{t_i}$ over the set of "acts" $\mathbf{R}^{T_j}$, and
- an action function $f_i : T_i \to S_i$, which maps each type of a player to his strategy space.

A type is thus a formal description of a player's actions and preferences (and thus beliefs). For simplicity, I restrict payoff functions to be the same across all types of a given player, and assume that these are "common knowledge" (loosely speaking). The subjective uncertainty faced by player i is then represented by player j's type space. Thus, strategy choice s_i by player i induces the act $[g_i(s_i, f_j(t_j))]_{t_j \in T_j}$, which is a vector of (utility) outcomes[4] in $\mathbf{R}^{T_j}$. I will sometimes denote the act $[g_i(s_i, f_j(t_j))]_{t_j \in T_j}$ by $G_i(s_i)$, where the t_jth component of $G_i(s_i)$ is $g_i(s_i, f_j(t_j))$.

[4]I assume that preferences can be represented in utility terms, without imposing more structure as to the precise nature of the representation at this stage.

I assume that each player knows his own type. This implies that from the point of view of type t_i of player i, while the state space consists of pairs $(t_i, t_j), t_j \in T_j$, the relevant uncertainty for this type of player i is summarized in the space T_j.

Definition 6 *A* belief operator *for player i maps subsets of the state space $\Omega = T_1 \times T_2$ to itself: i.e., $B_i : 2^\Omega \to 2^\Omega$. Consider events of the form $A = A_1 \times A_2$, where $A_1 \subseteq T_1$ and $A_2 \subseteq T_2$. Since each player knows his own type, we can define operators $\hat{B}_1$ and $\hat{B}_2$, where $\hat{B}_1 : 2^{T_2} \to 2^{T_1}$ and vice versa, such that $B_1(A_1 \times A_2) = (A_1 \cap \hat{B}_1(A_2)) \times T_2$ and $B_2(A_1 \times A_2) = T_1 \times (A_2 \cap \hat{B}_2(A_1))$.*

The belief operator $\hat{B}_1$ specifies the subsets of T_2 that player 1 believes: that is, $t_1 \in \hat{B}_1(A_2)$ if and only if type t_1 of player 1 believes that player 2 is some type in A_2. Then the event $B_1(A_1 \times A_2)$, which can be read as "player 1 believes the event A," consists of all pairs (t_1', t_2), where $t_1' \in A_1 \cap \hat{B}_1(A_2)$ and $t_2 \in T_2$. Note that this allows for the possibility that player 1 is wrong about player 2, but not that he is wrong about his own type.

We can define the analogues of the logical properties above for the "local" belief operator $\hat{B}_i$. For all $A_j, F_j \subseteq T_j$:

B1 : $B_i(T_j) = T_i$.

B2 : $B_i(\emptyset) = \emptyset$.

B3 : $\hat{B}_i(A_j) \cap \hat{B}_i(F_j) \subseteq \hat{B}_i(A_j \cap F_j)$.

B4 : $\hat{A}_j \subseteq F_j \Rightarrow \hat{B}_i(A_j) \subseteq \hat{B}_i(F_j)$.

The "composite" belief operator B_i inherits some logical properties from the "local" belief operator $\hat{B}_i$. For instance, if $\hat{B}_i$ satisfies [B1], so that $\hat{B}_i(T_j) = T_i$, then it follows that $B_i(\Omega) = (T_i \cap \hat{B}_i(T_j)) \times T_j = T_i \times T_j \equiv \Omega$. Also, if $\hat{B}_i$ satisfies [B3] and [B4], then so does B_i; B_i satisfies [B2] automatically. In addition, because players know their own type, if $\hat{B}_i$ satisfies [B1], then B_i satisfies [B5] and [B6] with equality. For [B6], note that $B_1(-B_1(A_1 \times A_2)) = B_1[-(A_1 \cap \hat{B}_1(A_2)) \times T_2]$. Letting $F_1 = -(A_1 \cap \hat{B}_1(A_2))$, we have $B_1(F_1 \times T_2) = [F_1 \cap \hat{B}_1(T_2)] \times T_2 = F_1 \times T_2 = -(A_1 \cap \hat{B}_1(A_2)) \times T_2 = -B_1(A_1 \times A_2)$.

Definition 7 *Consider an event $A \subseteq T_1 \times T_2$, where $A = A_1 \times A_2$. Denote the event "everyone believes A" by $B_*(A)$. Then*

$$B_*(A) = B_1(A) \cap B_2(A).$$

The event "everyone believes that everyone believes A" can then be denoted

$$[B_*]^2(A) \equiv B_*(B_*(A)) = B_1(B_*(A)) \cap B_2(B_*(A)) = B_1\left(B_1(A) \cap B_2(A)\right) \cap B_2\left(B_1(A) \cap B_2(A)\right),$$

and so on. The event "A is common belief" is thus

$$C(A) = \bigcap_{k=1}^{\infty} [B_*]^k(A).$$

Lemma 1 *Consider an event $A \subseteq T_1 \times T_2$, where $A = A_1 \times A_2$. If the belief operators $\hat{B}_1, \hat{B}_2$ satisfy [B4], then*

$$C(A) \subseteq \left[A_1 \cap \hat{B}_1(A_2) \cap \hat{B}_1(\hat{B}_2(A_1)) \cap \hat{B}_1(\hat{B}_2(\hat{B}_1(A_2))) \cap \cdots\right]$$
$$\times \left[A_2 \cap \hat{B}_2(A_1) \cap \hat{B}_2(\hat{B}_1(A_2)) \cap \hat{B}_2(\hat{B}_1(\hat{B}_2(A_1))) \cap \cdots\right],$$

with equality if both [B3] and [B4] hold.

Proof: By definition 7,

$$B_*(A) = \left[[A_1 \cap \hat{B}_1(A_2)] \times T_2\right] \cap \left[T_1 \times [A_2 \cap \hat{B}_2(A_1)]\right] = [A_1 \cap \hat{B}_1(A_2)] \times [A_2 \cap \hat{B}_2(A_1)].$$

Also,

$$\begin{aligned}B_*(B_*(A)) = B_1(B_*(A)) \cap B_2(B_*(A)) &= \left[A_1 \cap \hat{B}_1(A_2) \cap \hat{B}_1\left(A_2 \cap \hat{B}_2(A_1)\right)\right] \\ &\quad \times \left[A_2 \cap \hat{B}_2(A_1) \cap \hat{B}_2\left(A_1 \cap \hat{B}_1(A_2)\right)\right] \\ &\subseteq \left[A_1 \cap \hat{B}_1(A_2)) \cap \hat{B}_1(\hat{B}_2(A_1))\right] \times \left[A_2 \cap \hat{B}_2(A_1)) \cap \hat{B}_2(\hat{B}_1(A_2))\right],\end{aligned}$$

where the first line follows by definition and the second by [B4]; clearly, equality holds if $\hat{B}_1$ and $\hat{B}_2$ satisfy [B3] as well. Similarly,

$$\begin{aligned}B_*(B_*(B_*(A))) &= \left[A_1 \cap \hat{B}_1(A_2) \cap \hat{B}_1\left(A_2 \cap \hat{B}_2(A_1)\right) \cap \hat{B}_1\left(A_2 \cap \hat{B}_2(A_1) \cap \hat{B}_2(A_1 \cap \hat{B}_1(A_2))\right)\right] \\ &\quad \times \left[A_2 \cap \hat{B}_2(A_1) \cap \hat{B}_2\left(A_1 \cap \hat{B}_1(A_2)\right) \cap \hat{B}_2\left(A_1 \cap \hat{B}_1(A_2) \cap \hat{B}_1(A_2 \cap \hat{B}_2(A_1))\right)\right] \\ &\subseteq \left[[A_1 \cap \hat{B}_1(A_2)) \cap \hat{B}_1(\hat{B}_2(A_1)) \cap \hat{B}_1(\hat{B}_2(\hat{B}_1(A_2)))\right] \times \left[A_2 \cap \hat{B}_2(A_1)) \cap \hat{B}_2(\hat{B}_1(A_2)) \cap \hat{B}_2(\hat{B}_1(\hat{B}_2(A_1))\right],\end{aligned}$$

and so on. Hence,

$$\begin{aligned}C(A) &\equiv B_*(A) \cap B_*(B_*(A)) \cap B_*(B_*(B_*(A))) \cap \cdots \\ &\subseteq \left[A_1 \cap \hat{B}_1(A_2) \cap \hat{B}_1(\hat{B}_2(A_1)) \cap \hat{B}_1(\hat{B}_2(\hat{B}_1(A_2))) \cap \cdots\right] \\ &\quad \times \left[A_2 \cap \hat{B}_2(A_1) \cap \hat{B}_2(\hat{B}_1(A_2)) \cap \hat{B}_2(\hat{B}_1(\hat{B}_2(A_1))) \cap \cdots\right],\end{aligned}$$

given [B4], with equality given [B3] and [B4].◇

Corollary 1 $B_*(B_*(A)) \subseteq B_*(A)$.

Corollary 2 $\{[B_*]^n(A)\}_{n=1}^{\infty}$ *is a decreasing sequence.*

It is also possible to derive a fixed-point characterization of common belief in the spirit of Monderer and Samet (1989), using the notion of an evident event. An event is evident if, whenever it is true, everyone believes it.

Definition 8 *An event $A = A_1 \times A_2$ is* evident *if $A \subseteq B_*(A)$.*

Lemma 2 *Let $E \subseteq \Omega$. Suppose $\hat{B}_1$ and $\hat{B}_2$ satisfy [B4]. Then $t \in C(E)$ if there exists an evident event $F = F_1 \times F_2$ such that $t \in F \subseteq B_*(E)$. If $\hat{B}_1$ and $\hat{B}_2$ satisfy [B3], then the reverse implication holds as well.*

Proof: Suppose $t \in F \subseteq B_*(E)$ for some evident event F, so that $F \subseteq B_*(F)$. By Lemma 1, $B_*(F) \subseteq F$, so that F is a fixed point of the B_* operator. Now by [B4], $B_*(F) \subseteq [B_*]^2(E)$. Then by induction on n,

$$t \in F = [B_*]^n(F) \subseteq [B_*]^{n+1}(E)$$

for all $n \geq 0$. Hence $t \in C(E)$.

Conversely, suppose that $\hat{B}_1$ and $\hat{B}_2$ satisfy [B3], and let $t \in C(E)$. Now

$$\begin{aligned}
C(E) &= \bigcap_{k=1}^{\infty} [B_*]^k(E) \\
&\subseteq \bigcap_{k=2}^{\infty} [B_*]^k(E) \\
&= \bigcap_{k=2}^{\infty} \left[B_1([B_*]^{k-1}(E)) \cap B_2([B_*]^{k-1}(E)) \right] \\
&= \left[\bigcap_{k=2}^{\infty} B_1([B_*]^{k-1}(E)) \right] \bigcap \left[\bigcap_{k=2}^{\infty} B_2([B_*]^{k-1}(E)) \right] \\
&\subseteq B_1 \left(\bigcap_{k=2}^{\infty} [B_*]^{k-1}(E) \right) \bigcap B_2 \left(\bigcap_{k=2}^{\infty} [B_*]^{k-1}(E) \right) \quad \text{by [B3]} \\
&= B_*(C(E)).
\end{aligned}$$

So $C(E)$ is evident, and $C(E) \subseteq B_*(E)$ by definition.◇

Belief operators defined on players' type spaces thus provide us with a framework for expressing players' beliefs independent of any particular assumptions on preferences. The analogues of Savage and strong belief in this framework are defined as follows:

Definition 9 *Belief operator $\hat{B}_i$ represents* Savage belief *of preference relations* $\{\succeq_{t_i}\}_{t_i \in T_i}$ *if, for* $A_j \subseteq T_j$

$$\hat{B}_i(A_j) = \{t_i \in T_i | (x_{A_j}, z_{-A_j}) \succeq_{t_i} (x_{A_j}, y_{-A_j}) \ \forall x, y, z \in \mathbf{R}^{T_j}\}.$$

Definition 10 *Belief operator $\hat{B}_i$ represents* strong belief *of preference relations* $\{\succeq_{t_i}\}_{t_i \in T_i}$ *if*

$$\hat{B}_i(A_j) = \{t_i \in T_i | (x_{A_j}, v_{-A_j}) \succeq_{t_i} (y_{A_j}, z_{-A_j}) \ \forall v, x, y, z \in \mathbf{R}^{T_j} \text{ s.t. } x \gg y\}.$$

Common belief is then defined simply as the iterated application of the appropriate belief operator.

1.4 Beliefs in Games

In this section, I investigate the epistemic foundations of iterative dominance. For now, I restrict attention to pure strategies, so that each player only plays pure strategies, and views the other player as doing so as well. In this framework, the event "player i is rational", as perceived by player j, corresponds to

$$R_i = \{t_i \in T_i : \ [g_i(f_i(t_i), f_j(t_j))]_{t_j \in T_j} \succeq_{t_i} [g_i(s_i, f_j(t_j))]_{t_j \in T_j} \ \forall \ s_i \in S_i\}.$$

That is, a rational type of player i will only take an action that is optimal in the sense of being at least as good as any other action available to him. A strategy s_i is then a best response for player i if there exists some type t_i of player i such that $f_i(t_i) = s_i$ and t_i is rational in the sense defined above.

The following recursive construction is useful:

Lemma 3 *Let $R_i^1 = R_i$ and, for $n \geq 2$,*

$$R_i^n = \{t_i \in T_i : \ t_i \in \hat{B}_i(R_j^{n-1})\},$$

where $\hat{B}_i$ denotes some notion of belief. Let

$$R_i^\infty = \bigcap_{k=1}^{\infty} R_i^k.$$

Then if $\hat{B}_1$ and $\hat{B}_2$ satisfy [B4], $C(R_1 \times R_2) \subseteq R_1^\infty \times R_2^\infty$. When both [B3] and [B4] hold, then

$$C(R_1 \times R_2) = R_1^\infty \times R_2^\infty = \left[R_1 \cap \hat{B}_1(R_2)) \cap \hat{B}_1(\hat{B}_2(R_1)) \cdots\right] \times \left[R_2 \cap \hat{B}_2(R_1)) \cap \hat{B}_2(\hat{B}_1(R_2)) \cdots\right].$$

1.4.1 Strict Dominance

With this framework in hand, we can investigate strict and weak dominance.

Definition 11 *A strategy s_i is* dominated *if there exists $s_i' \in S_i$ such that*

$$g_i(s_i', s_j) > g_i(s_i, s_j) \ \forall \ s_j \in S_j.$$

The set of iteratively undominated strategies is defined as follows:

Definition 12 *Let $U_i^0 = S_i$, and*

$$U_i^m = \{s_i \in U_i^{m-1} : \nexists s_i' \in U_i^{m-1} \text{ such that } g_i(s_i', s_j) > g_i(s_i, s_j) \ \forall \ s_j \in U_j^{m-1}\}.$$

Then

$$U_i^\infty = \bigcap_{k=1}^{\infty} U_i^k$$

is the set of strategies that survive iterated deletion of dominated strategies, where dominance is by pure strategies.

Finiteness of the original set of actions S_i implies the existence of an integer M such that $U_i^\infty = U_i^m$ for all $m \geq M$. The first result is that with strong assumptions on what players know, rational players will play iteratively undominated strategies, regardless of the particular representation of their preferences.

Proposition 1 *In the 2-player, finite action game* $< S_1, S_2, g_1, g_2 >$ *let* $(t_1, t_2) \in C(R_1 \times R_2)$, *where belief is defined as Savage or strong* belief. Suppose that preferences satisfy [P1]-[P4] and [P5*]. Then, for* $i = 1, 2$, $f_i(t_i) \in U_i^\infty$. *When preferences satisfy [P7] in addition, the result also holds if* $\hat{B}_1$ *and* $\hat{B}_2$ *are defined as strong belief.*

Proof: I use the characterization in Lemma 3. Let $F_i^0 = T_i$, and

$$F_i^m = \{t_i \in T_i : f_i(t_i) \in U_i^m\},$$

so that F_i^m denotes the "event" that player i plays a strategy that is undominated at stage m. Clearly, $t_i \in R_i^1$ implies $f_i(t_i) \in U_i^1$ – if not, there exists $s_i' \in S_i$ such that $G_i(s_i') \gg G_i(f_i(t_i))$ but $G_i(f_i(t_i)) \succeq_{t_i} G_i(s_i')$, which violates monotonicity of preferences. Hence $t_i \in R_i^1 \Rightarrow t_i \in F_i^1$. By the same logic, $R_j^1 \subseteq F_j^1$, so that, by [B4], $t_i \in R_i^2 \Rightarrow t_i \in \hat{B}_i(R_j^1) \Rightarrow t_i \in \hat{B}_i(F_j^1)$. Hence if player i knows that player j is rational, then he also knows that player j will not play a dominated strategy.

Now suppose, for $k = 1, \cdots, n$, that $t_i \in R_i^k \Rightarrow f_i(t_i) \in U_i^k$. I show that $t_i \in R_i^{n+1} \Rightarrow f_i(t_i) \in U_i^{n+1}$. Suppose not. Then there exists $s_i' \in U_i^n$ such that $g_i(s_i', s_j) > g_i(f_i(t_i), s_j)$ $\forall s_j \in U_j^n$, or, equivalently, $g_i(s_i', f_j(t_j)) > g_i(f_i(t_i), f_j(t_j))$ $\forall t_j \in F_j^n$. By definition, $t_i \in R_i^{n+1} \Rightarrow t_i \in \hat{B}_i(R_j^n)$. By the induction hypothesis for player j, $t_i \in \hat{B}_i(F_j^n)$, so that player i knows the event F_j^n. But Savage belief of F_j^n implies that for any $\epsilon > 0$.

$$(G_i(s_i')_{F_j^n}, G_i(s_i')_{-F_j^n}) \sim_{t_i} (G_i(s_i')_{F_j^n}, [G_i(f_i(t_i)) + \epsilon]_{-F_j^n})$$

$$\succ_{t_i} (G_i(f_i(t_i))_{F_j^n}, G_i(f_i(t_i))_{-F_j^n}),$$

where the second line follows by [P5*][5]. But this violates rationality for player i, so that we must have $f_i(t_i) \in U_i^{n+1}$.

For strong* belief of F_j^n, we have

$$(G_i(s_i')_{F_j^n}, G_i(s_i')_{-F_j^n}) \succ_{t_i} (G_i(f_i(t_i))_{F_j^n}, G_i(f_i(t_i))_{-F_j^n}),$$

which, again, violates rationality for player 1. Hence for all $k \geq 1, t_i \in R_i^k \Rightarrow f_i(t_i) \in U_i^k$, so that $t_i \in R_i^\infty \Rightarrow f_i(t_i) \in U_i^\infty$.

When preferences are continuous, Savage and strong belief are equivalent, so the result above follows for strong belief as well.◇

When the notion of belief is weakened further, it is not always possible to rule out play of iteratively strictly dominated strategies. If players have MEU preferences, then common belief of rationality may still imply that they play iteratively dominated strategies.

Ryan (2000) argues that one reason to object to weak belief is that it sometimes fails to satisfy logical coherence, i.e., that it can fail [B2]. As noted earlier, this can happen for MEU preferences if the intersection of the supports of the decision maker's beliefs is empty, so that no state is considered possible. The following example shows that, even when weak belief satisfies [B2], it can be too weak to rule out very much.

Example 2 Consider the following 2-player game.

[5] I thank Matthew Ryan for suggesting this line of argument.

		2	
		l	r
1	U	2, 1	1, 0
	D	1, 2	2, 1

Let $T_1 = \{\hat{t}_1, t'_1\}$ and $T_2 = \{\hat{t}_2, t'_2, t''_2\}$. Suppose that players have MEU preferences, with utilities denoted by $\{\mathbf{u}_{t_i}\}_{t_i \in T_i}$, and belief sets $\{\Delta_{t_i}\}_{t_i \in T_i}$.

Now if player 2 is rational, then she will never play r, which is strictly dominated for her. Once r is eliminated, D is strictly dominated for player 1, so that the unique outcome that survives iterated deletion of strictly dominated strategies is (U, l). Let $R_2 = \{\hat{t}_2\}$ be the event "player 2 is rational"; then $f_2(\hat{t}_2) = l$. Suppose that $f_2(t'_2) = f_2(t''_2) = r$ and $\Delta_{t_2} = \{\hat{p}_2\}$, where $\hat{p}_2(\hat{t}_1) = 1$ for all $t_2 \in T_2$. I am interested in when player 1 will prefer the act induced by D, $(1_{R_2}, 2_{-R_2})$, over the one induced by U, $(2_{R_2}, 1_{-R_2})$. Let $f_1(\hat{t}_1) = D$, and $\Delta_{\hat{t}_1}$ be the convex hull of $\{\hat{p}_1, p'_1\}$, where the weights placed on the different states are:

	$\hat{t}_2$	t'_2	t''_2
$\hat{p}_1$	$\frac{1}{3}$	0	$\frac{2}{3}$
p'_1	$\frac{1}{3}$	$\frac{2}{3}$	0

Then at $(\hat{t}_1, \hat{t}_2)$, the only state that player 1 considers possible is $\hat{t}_2$, and the only state that player 2 considers possible is $\hat{t}_1$. Notice that although player 1 believes that player 2 is rational, the *maximum* weight he ever places on the event R_2 is $\frac{1}{3}$, whereas the *minimum* weight he places on its complement, which is *not* believed, is $\frac{2}{3}$. Hence

$$\mathbf{u}_{\hat{t}_1}(U) = \min_{p_1 \in \Delta_{\hat{t}_1}} [p_1(R_2)\, 2 + p_1(-R_2)\, 1] = \min_{p_1 \in \Delta_{\hat{t}_1}} [2\, p_1(R_2) + (1 - p_1(R_2))]$$

$$= 1 + \min_{p_1 \in \Delta_{\hat{t}_1}} p_1(R_2) = 1 + \frac{1}{3},$$

and

$$\mathbf{u}_{\hat{t}_1}(D) = \min_{p_1 \in \Delta_{\hat{t}_1}} [p_1(R_2)\ 1 + p_1(-R_2)\ 2] = 1 + \min_{p_1 \in \Delta_{\hat{t}_1}} p_1(-R_2) = 1 + \frac{2}{3},$$

so that $\mathbf{u}_{\hat{t}_1}(D) > \mathbf{u}_{\hat{t}_1}(U)$, and player 1 will (rationally) prefer playing strategy D to strategy U, despite weakly believing that player 2 is rational and will therefore not play r. Therefore, common weak belief of rationality does not rule out play of iteratively strictly dominated strategies.◇

1.4.2 Admissibility

When preferences are taken to satisfy admissibility, it is possible to rule out play of weakly dominated strategies. Consider two strategies $s_i, s_i' \in S_i$, such that $g_i(s_i', s_j) \geq g_i(s_i, s_j)$ $\forall s_j \in S_j$, with strict inequality for some $s_j' \in S_j$. If there exists $t_j' \in T_j$ such that $f_j(t_j') = s_j'$, so that $g_i(s_i', f_j(t_j)) \geq g_i(s_i, f_j(t_j))$ $\forall t_j \in T_j$, with strict inequality for t_j', admissibility implies $G_i(s_i') \succ_{t_i} G_i(s_i)$, so that type t_i of player i will not play the weakly dominated strategy, s_i. Assume from now on that every action is played by some type of each player. Denote the set of strategies that are not weakly dominated for player i by W_i.

When preferences are admissible, Savage belief is uninformative in that nothing non-trivial is ever believed (Morris 1997, Lemma 3). In fact, common Savage belief of rationality may lead to a contradiction in some games.

Proposition 2 *Suppose that the action functions f_i are onto (that is, that every strategy is played by some type of each player). Then there exist games in which common Savage belief of admissibility is impossible.*

Proof: Suppose that, for player 1, there exists $s_1 \in S_1$ which is weakly dominated. Then there exists $t_1' \in T_1$ such that $f_1(t_1') = s_1$, so that this type of player 1 violates admissibility. The only event that player 2 Savage believes is T_1, so that the only event that is commonly Savage believed is the whole state space. Since each player knows his own type, it follows from definition 7 that if an event is commonly believed, then it must in fact be true. Thus common belief of admissiblity would imply that *all* types of player 1 satisfy admissibility, which contradicts the hypothesis that type t_1' of player 1 violates admissibility.◇

This result parallels Theorem 8 in Samuelson (1992), where knowledge is defined as absolute certainty (or "probability 1" belief), corresponding to Savage belief in this framework. Since Savage belief is too strong in this context, we may therefore ask what the implications of common strong belief of rationality are for the play of undominated strategies, where "rational" players are now those whose preferences satisfy admissibility. Define

$$S_i^1 = \{s_i \in W_i : \nexists s_i' \in W_i \text{ such that } g_i(s_i', s_j) > g_i(s_i, s_j) \; \forall \; s_j \in W_j\},$$

and, for $m \geq 2$,

$$S_i^m = \{s_i \in S_i^{m-1} : \nexists s_i' \in S_i^{m-1} \text{ such that } g_i(s_i', s_j) > g_i(s_i, s_j) \; \forall \; s_j \in S_i^{m-1}\}.$$

$$S_i^\infty = W_i \cap \left[\bigcap_{k=1}^{\infty} S_i^k\right]$$

Thus S_i^∞ is the set of strategies which survive one round of deletion of weakly dominated strategies, followed by subsequent rounds of deletion of strictly dominated strategies, where domination is by pure strategies. As before, finiteness of the original set of actions S_i implies the existence of an integer M such that $S_i^\infty = S_i^m$ for all $m \geq M$.

Proposition 3 *In the 2-player, finite action game* $< S_1, S_2, g_1, g_2 >$ *when preferences satisfy [P1] - [P5], if* $(t_1, t_2) \in C(R_1 \times R_2)$, *where belief is defined as strong belief, then, for* $i = 1, 2,\ f_i(t_i) \in S_i^\infty$.

Proof: Let $E_i^0 = T_i$,

$$E_i^1 = \{t_i \in T_i : f_i(t_i) \in W_i\},$$

and

$$E_i^{m+1} = \{t_i \in T_i : f_i(t_i) \in S_i^m\}.$$

By lemma 3, $t_i \in R_i^\infty$. We have already seen that $t_i \in R_i^1$ implies $f_i(t_i) \in W_i$; i.e., that $t_i \in E_i^1$. Now suppose, for $k = 1, \cdots, n$, that $t_i \in R_i^{k+1} \Rightarrow f_i(t_i) \in S_i^k$. I show that $t_i \in R_i^{n+2} \Rightarrow f_i(t_i) \in S_i^{n+1}$. Suppose not. Then there exists $s_i' \in S_i^n$ such that $g_i(s_i', s_j) > g_i(f_i(t_i), s_j)\ \forall s_j \in S_j^n$, i.e., that $g_i(s_i', f_j(t_j)) > g_i(f_i(t_i), f_j(t_j))\ \forall t_j \in E_j^{n+1}$. By definition, $t_i \in R_i^{n+1} \Rightarrow t_i \in \hat{B}_i(R_j^n)$. By the induction hypothesis for player j, $t_i \in \hat{B}_i(E_j^{n-1})$, so that player i knows the event E_j^{n+1}. Strong belief of E_j^{n+1} implies that, for any $\epsilon > 0$,

$$(G_i(s_i')_{E_j^{n+1}}, G_i(s_i')_{-E_j^{n+1}}) \succeq_{t_i} (G_i(f_i(t_i))_{E_j^{n+1}}, [G_i(f_i(t_i)) + \epsilon]_{-E_j^{n-1}})$$

$$\succ_{t_i} (G_i(f_i(t_i))_{E_j^{n+1}}, G_i(f_i(t_i))_{-E_j^{n+1}}),$$

where the second line follows by admissibility. But this violates rationality for player i. Hence for all $k \geq 1, t_i \in R_i^{k+1} \Rightarrow f_i(t_i) \in S_i^k$, so that $t_i \in R_i^\infty \Rightarrow f_i(t_i) \in S_i^\infty$. ◇

Brandenburger (1992) considers the case of preferences which satisfy lexicographic expected utility (LEU). He defines the set of "permissible" strategies as those which are chosen if there is common first-order knowledge of rationality (common strong belief), and shows that this is equivalent to the set of strategies which survive the iterated deletion procedure above, where dominance is defined to include dominance by mixed strategies[6]. Note that common (strong) belief of admissibility does not imply iterated deletion of *weakly* dominated strategies. Intuitively, this should not be surprising. Strong belief requires only that preferences reflect strict dominance on the event that is believed; admissibility requires that preferences respect weak dominance, but only with respect to the state space as a whole. In other words, admissibility does not provide sufficient structure to determine the direction of preferences when the weak dominance is over a *subset* of the state space. Hence common knowledge of admissibility does not allow players to rule out play of *iteratively* weakly dominated strategies by other players.

Stahl (1995) considers the case of LEU preferences (Blume, Brandenburger and Dekel 1991a) and defines the set of *lexicographically rationalizable* (LR) strategies where, at each stage, the just eliminated strategies are considered infinitely less likely than the currently viable strategies but infinitely more likely than previously eliminated strategies. He shows that when mixed strategy dominance is allowed, the set of LR strategies is equivalent to the set of strategies that survive iterated deletion of weakly dominated strategies, where all dominated strategies are removed at each stage. The required epistemic condition is 'iterated lexicographic coherence,' which formalizes the intuition that strategies ruled out at one step of reasoning should be considered infinitely less likely than the remaining viable strategies. In terms of axioms on preferences, this would require that preferences satisfy "iterated admissibility" in the sense that preferences would need to be admissible with respect to events that are *believed*. Admissibility with respect to the event that strategies

[6]He draws on a result of Blume, Brandenburger and Dekel (1991a) which shows that if a player's lexicographic probability system (LPS) has full support, then the player's optimal actions will be admissible.

inadmissible at stage $k-1$ are not played would allow a player to ignore the "null" event in determining the ranking of weakly dominated strategies at stage k (with respect to the other player's stage $k-1$ strategies). However, this is a very strong condition, one which seems almost to run counter to the logic of admissibility. Brandenburger and Keisler (2000) also provide epistemic conditions for iterated admissibility in an LEU framework. However, the belief operator that they require does not satisfy monotonicity. Moreover, they show that, generically, iterated admissibility requires that the type space be "belief-complete" - that is, that player i consider not merely all *strategies* of player j as possible, but also every *belief* that player j might conceivably have. In finite type spaces of the sort studied in this chapter, iterated admissibility will clearly not obtain in general.

1.5 Equilibrium

This section considers equilibrium concepts for players in 2-person games.

Definition 13 *A strategy profile $s^* = (s_1^*, s_2^*)$ is a* pure-strategy Nash Equilibrium *if, for $i = 1, 2$,*

$$g_i(s_i^*, s_{-i}^*) \geq g_i(s_i', s_{-i}^*) \;\; \forall s_i' \in S_i.$$

I now show that rationality and mutual Savage or strong* belief of actions imply that players' choices will be a Nash equilibrium.

Proposition 4 *Let $s^* = (s_1^*, s_2^*)$ be a pair of actions. For $i = 1, 2$, let*

$$Q_i = \{t_i \in T_i : f_i(t_i) = s_i^*\}.$$

Suppose for some $t^* = (t_1^*, t_2^*)$ *that* $t^* \in R_1 \times R_2$, *and* $t^* \in B_*(Q_1 \times Q_2)$, *where* $\hat{B}_1$ *and* $\hat{B}_2$ *denote Savage or strong* belief. Suppose that preferences satisfy [P1] - [P4] and [P5*]. Then* s^* *is a Nash equilibrium. If preferences satisfy [P1]-[P4] and [P5], then strong belief is sufficient to ensure that* s^* *is a Nash equilibrium.*

Proof: Consider player 1. Note first that, since $(t_1^*, t_2^*) \in B_1(Q_1 \times Q_2) = [Q_1 \cap \hat{B}_1(Q_2)] \times T_2$, it must be the case that $f_1(t_1^*) = s_1^*$. Now Savage belief of Q_2 implies that

$$([g_1(f_1(t_1^*), f_2(t_2))]_{t_2 \in Q_2}, [g_1(f_1(t_1^*), f_2(t_2))]_{t_2 \in -Q_2}) \equiv (g_1(s_1^*, s_2^*)_{Q_2}, [g_1(s_1^*, f_2(t_2))]_{t_2 \in -Q_2})$$
$$\sim_{t_1^*} (g_1(s_1^*, s_2^*)_{T_2})$$

and

$$([g_1(s_1, f_2(t_2))]_{t_2 \in Q_2}, [g_1(s_1, f_2(t_2))]_{t_2 \in -Q_2}) \equiv (g_1(s_1, s_2^*)_{Q_2}, [g_1(s_1, f_2(t_2))]_{t_2 \in -Q_2})$$
$$\sim_{t_1^*} (g_1(s_1, s_2^*)_{T_2}) \ \forall s_1 \in S_1.$$

Now, since $t_1^* \in R_1$,

$$[g_1(s_1^*, f_2(t_2))]_{t_2 \in T_2} \succeq_{t_1^*} [g_1(s_1, f_2(t_2))]_{t_2 \in T_2} \ \forall \ s_1 \in S_1,$$

or, equivalently,

$$(g_1(s_1^*, s_2^*)_{Q_2}, [g_1(s_1^*, f_2(t_2))]_{t_2 \in -Q_2}) \succeq_{t_1^*} (g_1(s_1, s_2^*)_{Q_2}, [g_1(s_1, f_2(t_2))]_{t_2 \in -Q_2}) \ \forall s_1 \in S_1.$$

Hence we have

$$(g_1(s_1^*, s_2^*)_{T_2}) \succeq_{t_1^*} (g_1(s_1, s_2^*)_{T_2}) \ \forall s_1 \in S_1.$$

This is possible if and only if, $\forall s_1 \in S_1$,

$$g_1(s_1^*, s_2^*) \geq g_1(s_1, s_2^*).$$

By similar reasoning,

$$g_2(s_1^*, s_2^*) \geq g_2(s_1^*, s_2) \; \forall s_2 \in S_2.$$

Hence s^* is a Nash equilibrium.

For strong* belief of actions, suppose that actions are not a Nash equilibrium. Note as before that $f_i(t_i^*) = s_i^*$ for $i = 1, 2$. Suppose that player 1, say, has a strictly profitable deviation. That is, there exists s_1' such that

$$g_1(s_1', s_2^*) > g_1(s_1^*, s_2^*).$$

Strong* belief of Q_2 implies that

$$(g_1(s_1', s_2^*)_{Q_2}, [g_1(s_1', f_2(t_2))]_{t_2 \in -Q_2}) \succ_{t_1^*} (g_1(s_1^*, s_2^*)_{Q_2}, [g_1(s_1^*, f_2(t_2))]_{t_2 \in -Q_2}).$$

But this contradicts the rationality of player 1. Hence player 1 has no unilateral deviation that yields a strictly higher payoff, and similarly for player 2, so that s^* is a Nash equilibrium.◇

Strong* and Savage belief thus serve as benchmark notions. In proposition 4, belief of actions is sufficiently strong to rule out either player deviating unilaterally given what the other player is playing. However, once we relax belief of actions to mean weak belief, we can no longer rule out non-Nash behavior.

Example 3 Consider the following 2-player game.

		2: l	r
1	t	1, 3	2, 4
	b	2, 2	1, 1

Players are assumed to have MEU preferences $\{\mathbf{u}_{t_i}\}_{t_i \in T_i}$, $\{\Delta_{t_i}\}_{t_i \in T_i}$. Let $[b]$ denote the event that player 1 plays the strategy b, and let $[r]$ denote the event that player 2 plays the strategy r. Strong or Savage belief of $[r]$ implies that, for any $t_1 \in \hat{B}_1([r]), p_1([r]) = 1 \; \forall \, p_1 \in \Delta_{t_1}$. Clearly, this implies that player 1 would prefer the act induced by t, $(1_{-[r]}, 2_{[r]})$ over the one induced by b, which is $(2_{-[r]}, 1_{[r]})$, so that it is not possible to support play of b when player 1 strongly believes that player 2 is playing r. However, suppose that type t_1' of player 1 weakly believes that player 2 is playing r. Weak belief imposes only the restriction that, for all $t_2 \notin [r]$, there exists $p_1 \in \Delta_{t_1'}$ such that $p_1(t_2) = 0$. Now

$$\mathbf{u}_{t_1'}(t) = \min_{p_1 \in \Delta_{t_1}} [2p_1([r]) + p_1(-[r])]$$

and

$$\mathbf{u}_{t_1'}(b) = \min_{p_1 \in \Delta_{t_1}} [p_1([r]) + 2p_1(-[r])].$$

Suppose $\min_{p_1 \in \Delta_{t_1'}} p_1([r]) = 1/4$ and $\min_{p_1 \in \Delta_{t_1'}} p_1(-[r]) = 1/3$. Then

$$\mathbf{u}_{t_1}(t) = \min_{p_1 \in \Delta_{t_1}} [1 + p_1([r])] = 1 + \frac{1}{4}$$

$$\mathbf{u}_{t_1'}(b) = \min_{p_1 \in \Delta_{t_1'}} [1 + p_1(-[r])] = 1 + \frac{1}{3}.$$

Thus type t_1' of player 1 prefers playing strategy b, even though he weakly believes player 2 is playing r. Similarly, for type t_2' of player 2,

$$\mathbf{u}_{t_2'}(l) = \min_{p_2 \in \Delta_{t_2'}} [2 + p_2(-[b])]$$

and

$$\mathbf{u}_{t_2'}(r) = \min_{p_2 \in \Delta_{t_2'}} [1 + 3p_2(-[b])].$$

Then suppose that $\min_{p_2 \in \Delta_{t'_2}} p_2(-[b]) = 9/16$. Then

$$u_{t'_2}(l) = \frac{41}{16} < u_{t'_2}(r) = \frac{43}{16} .$$

So player 2 prefers playing strategy r despite weakly believing that player 1 plays b. Hence weak belief of actions is not sufficient to rule out non-Nash behavior.◇

The logical next step is to consider equilibrium in mixed strategies. There are two accepted interpretations of mixed strategy Nash Equilibrium. The traditional interpretation is that players actually mix according to the equilibrium strategies; the second interpretation holds that player 1's mixed strategy represents not his actual action but player 2's beliefs about what pure strategy player 1 is going to pick. Under the latter interpretation, Nash Equilibrium is usually stated as a pair of probability measures such that each pure strategy in the support of player 1's belief maximizes player 2's payoff and vice versa. The definition thus makes a simultaneous statement about strategies and beliefs: each player has a belief regarding the other player's set of optimal strategies (best responses) under which the event that the other player will play a non-optimal strategy is null, and the beliefs themselves constitute the equilibrium strategies.

As is well known, however, this equivalence between beliefs and strategies breaks down when players do not have expected utility preferences: when beliefs are not represented by a single probability measure, then a player's mixed strategy has no ready interpretation as the other player's belief over her possible pure strategies. The standard consistency requirement of Nash Equilibrium, where players expect other players to play best responses, now takes the form of specifying apppropriate supports for equilibrium beliefs, with the proviso that all strategies in the support of player i's belief be optimal for player j and vice versa. For example, Lo (1996) considers MEU preferences, and defines a notion of *Beliefs Equilibrium* as a pair (in the context of 2-player games) of closed and convex sets of probability measures such that the complement of each player's set of mixed-strategy best responses to his belief set is Savage-null from the point of view of the other player. Dow and Werlang (1994) define *Nash Equilibrium under Uncertainty* for players with CEU preferences as a pair of capacities such that all pure strategies in the DW-support of each player's capacity are optimal for the other player – that is, that the set of player 2's non-optimal strategies given player 1's belief is null in this sense, and vice versa.[7]

The issue is complicated by the fact that there seems to be very little agreement on whether the appropriate objects of choice are mixed or pure strategies. In Lo (1996), players have beliefs over opponents' mixed strategies, while Klibanoff (1996) treats players with MEU preferences as choosing mixed strategies, but as having beliefs over opponents' pure strategies. Mixed strategies as objects of choice are often justified on the grounds that uncertainty averse decision makers may have a strict incentive to randomize. Eichberger and Kelsey (1996) show that in the Anscombe-Aumann framework, a strict preference for randomization does exist, while it does not necessarily do so in the Savage framework.

I continue to consider players as playing pure strategies, and as regarding opponents as doing so. In the present context, the appropriate counterpart to an equilibrium in beliefs may be defined following Lo (1996) in terms of preferences.

Definition 14 $\{\succeq_{t_i}, \succeq_{t_j}\}$ *is a* Nash Equilibrium in preferences *if*

[7]This can be related to what Morris (1997) calls *Negative belief.* The negative belief operator for CEU preferences with capacity ν_ω is defined as follows:

$$B(E) = \{\omega \in \Omega \mid \nu_\omega(-E) = 0\}.$$

Negative belief can fail to be normal.

1. *there exists* $\Gamma_i \subseteq S_i$ *such that the event* $-F_i$, *where* $F_i = f_i^{-1}(\Gamma_i)$, *is null with respect to* $\succeq_{t_j}$, $j \neq i$; *and*

2. *for all* $s_i \in \Gamma_i$, $G_i(s_i) \succeq_{t_i} G_i(s_i') \; \forall \; s_i' \in S_i$.

In words, the preferences form a Nash Equilibrium if the event that each player is irrational is null from the point of view of the other player, and each player's action is optimal given her preferences.

Proposition 5 *Suppose for some* $t^* = (t_1^*, t_2^*)$ *that players are rational and that this is mutual belief, so that* $t^* \in R_1 \times R_2$, *and* $t^* \in B_*(R_1 \times R_2)$. *Let player 1's belief regarding player 2 be* $F_2 \subseteq T_2$, *so that* $t_1^* \in \hat{B}_1(F_2)$; *similarly, let* $t_2^* \in \hat{B}_2(F_1)$. *Suppose that* $t^* \in B_*(\hat{B}_1(F_2) \times \hat{B}_2(F_1))$, *i.e., that the beliefs are mutually believed. If the belief operator is normal, then* $\{\succeq_{t_1^*}, \succeq_{t_2^*}\}$ *is a Nash Equilibrium in preferences.*

Proof: We begin with two observations. First, mutual knowledge of beliefs implies that they are correct, so that if player 1 believes that player 2 believes F_1, then player 2 does indeed believe F_1, and vice versa. Secondly, if we let $\Gamma_i = f_i(F_i)$, then the first part of definition 14 is satisfied; we need only to prove that, for all $s_i \in \Gamma_i$, s_i is optimal for player i. So suppose $s_2' \in \Gamma_2$, and let $t_2' = f_2^{-1}(s_2')$. I show that s_2' is optimal for player 2. We know that at t_1^*, player 1 considers t_2' possible, and that she believes R_2 and $\hat{B}_2(F_1)$, so that she believes $R_2 \cap \hat{B}_2(F_1)$ (by [B3]). If the belief operator is normal, then every 'state' that (type t_1^* of) player 1 considers possible must be contained in the event player 1 believes; hence it follows that $t_2' \in R_2 \cap \hat{B}_2(F_1)$, so that s_2' must be optimal for player 2, given his beliefs about player 1. A similar argument from the point of view of player 2 establishes that all strategies in Γ_1 must be optimal for player 1 given her beliefs about player 2.◇

This result can be viewed as a general analogue to Theorem A in Aumann and Brandenburger (1995). Note that it imposes no explicit structure on preferences, and very little structure on beliefs. Moreover, it does not solve the problem of mixed versus pure strategies: if all definitions were in terms of mixtures, then the analogue of the proposition with preferences and belief defined over mixed strategies would go through.[8]

1.6 Conclusion

This chapter has offered a unified language for thinking about interactive beliefs in the context of non-expected utility maximisers. By exploring a minimal preference structure compatible with games from a subjective viewpoint, I provide a framework within which extensions of single-person decision theory to games can be compared and evaluated. By considering the relation of beliefs to preferences, I make it possible to identify the factors that drive notions of dominance and equilibrium. Some of the weaknesses of the existing extensions of non-expected utility theory to games are also identified.

[8]Proposition 6 in Lo (1996) is a special case that is very similar in spirit to this result, but employs mixed strategies.

Chapter 2

Price Dispersion in the Wholesale Market for Steel

2.1 Introduction

This chapter looks at the factors underlying price discrimination from an empirical perspective. I study the pricing decision of a durable commodity intermediary, a company which purchases bulk quantities of a homogeneous commodity at competitive prices on the spot market and sells inventories to customers at a mark-up over cost. The firm engages in minimal production processing, and makes profits by "breaking bulk," assuming substantial risk in the process.

Given the strong global competition in the steel market and the homogeneity of the product, we might expect to see few opportunities for arbitrage. However, if there is imperfect information about prices and the search costs are significant, the firm could make profits by arbitraging information. On the other hand, the firm is likely to have incomplete information on how knowledgeable a given customer is, and thus what the customer's reservation price might be. In addition, the firm is often quite exposed to fluctuations in the world price of steel, and may also be vulnerable to the risk of default on the part of buyers. In such an environment, it is not clear what would constitute an optimal pricing strategy.

In this chapter I analyze a rich dataset from a U.S. steel supplier ("Company S"), consisting of three years of daily data on sales, purchases and inventory holdings of steel plate. This is combined with information on the firm's customers. The findings can be summarized in the following *stylized facts*:

- Customers who make larger orders receive lower prices, suggesting the presence of quantity discounts;
- distance is negatively correlated with net prices, suggesting that the firm is pricing to take competition into account;
- there is considerable stickiness in the price charged to an individual customer within a period of a few weeks;
- during the period following the Russian currency crisis, purchase prices dropped quickly, but sale prices remained at their pre-crisis levels for several months;
- even controlling for variables such as quantity purchased and distance, individual customers appear to receive differential treatment from the firm.

Figures 2.1 and 2.2 display the time series of net price and unit cost for the firm's most heavily traded product (product P2) over a period of eighteen months each. There is clearly substantial dispersion in prices on a daily basis, and there are trend movements in costs, tracked imperfectly by prices. Figures 2.3 and 2.4 break the sample up into one year periods and display net prices, unit costs and net margins. It is interesting to note that, when costs rise around February 1998 (business day 150) and stay high for the next eight months, margins come down, but when costs fall subsequently around September 1998, net margins rise and stay high for the next several months. The average price level does not decline substantially until June of the following year, following a further sharp dip in costs. For the rest of the sample period, in fact, average margins remain above the levels seen in 1997 and through the first two quarters of 1998.

Table 2.1 provides a snapshot view of daily price dispersion for product P2 over the course of the sample. I indicate the number of times in the sample that each customer buys P2. I also indicate their distance in miles from Company S.[9] Transactions for each day are ordered by the net price, which is quoted in dollars per hundredweight of steel, and varies by over five dollars on one of the days (business day 580). Note the general decline in the net price level over the course of the thirty months that the table represents. The net profit in dollars per unit of the product sold also varies substantially, with some sales being more than twice as profitable per unit than others.[10] To get an idea of the magnitudes involved, consider business day 89. Sales 2 and 3 net a total revenue of $ 3,945 and $ 2,436 respectively. The net profit per unit between the two sales differs by over eleven dollars; at eighteen units, this translates into an additional net profit of over two hundred dollars for a single transaction. Even for smaller orders, the differences can be notable: on business day 580, sale 6 is netting an additional net profit of a hundred and thirty dollars over sale 5. There does appear to be some evidence of quantity discounts - larger orders are being sold at lower net prices in general, and hence lower net profit per unit, than smaller ones. Distance also seems to be a factor that influences prices, though at first glance its effects

[9]Company S is located in Connecticut.

[10]The careful reader might notice that the ordering of net profit per unit does not always match that of net price: this is due to the way Company S calculates margins and profits, which I discuss in more detail below.

are somewhat uneven.

Table 2.1: A Snapshot View of Daily Price Dispersion for Product P2

business day	sale by day	total number of purchases	units bought	net price ($/hundwt)	net profit per unit ($)	distance in miles
89	1	1	2	23.95	73.10	231.73
89	2	23	20	24.15	76.37	301.59
89	3	89	18	24.85	87.80	61.77
89	4	17	4	25.6	100.05	198.34
111	1	21	9	24.65	85.68	131.56
111	2	89	11	24.85	88.94	61.77
111	3	7	4	25.5	99.56	251.63
111	4	17	3	25.6	101.20	198.34
111	5	8	6	26.25	111.82	127.92
394	1	14	10	23.4	46.26	141.14
394	2	30	4	23.5	47.89	145.98
394	3	17	6	23.8	88.14	198.34
394	4	8	2	24	91.39	127.92
394	5	18	1	24	91.43	104.85
394	6	23	2	24.45	98.75	474.09
394	7	7	2	24.5	64.23	2
483	1	26	15	19.15	53.83	185.23
483	2	82	2	19.35	57.09	338.82
483	3	20	6	19.6	61.18	141.84
483	4	20	2	20.15	70.03	335.67
483	5	11	8	20.2	70.85	145.56
483	6	11	8	20.45	74.93	105.45
483	7	13	2	21.35	89.63	45.46
580	1	6	28	20.2	62.88	106.69
580	2	44	28	20.25	63.69	296.29
580	3	6	7	20.5	67.77	60.06
580	4	20	20	20.95	75.13	335.67
580	5	1	2	21.95	91.39	170.75
580	6	4	2	25.95	156.78	117.63

Table 2.2 presents some summary statistics for frequent buyers of product P2. The customer numbers refer to the rank of the customer in the sample, by the number of times they buy P2, which is indicated in the final column. For each customer, I indicate the average net price they receive over the course of the sample, as well as the average net percentage margin and the average net profit in dollars per transaction. I also include each customer's primary SIC code, which indicates their line of business. Most of the large customers are wholesalers and distributors of steel, and the rest are mostly construction

firms.[11] The mean net margin varies from thirty percent to as low as fifteen percent. Note that the most frequent customer for this product has one of the highest average net margins, while the next most frequent buyer has one of the lowest. Again, the table suggests the presence of quantity discounts, as well as lower prices in general to distributors of steel (2-digit SIC code 50). However, there is a surprising amount of dispersion here as well. Consider customers 2 and 13: at an average of twenty two units per transaction, customer 13 is bringing the firm an incremental profit of over three hundred dollars per transaction on average compared to customer 2, or about six thousand dollars over the course of the sample (averaging out at two thousand dollars a year).

[11] I discuss SIC codes and customer characteristics below in greater detail.

Table 2.2: Summary Statistics for Frequent Buyers, Product P2

Customer	Mean net price (\$/hwt)	Mean net margin	Mean net profit (\$)	Mean profit per unit (\$)	Mean units bought	Distance in miles	Primary SIC code	Number of purchases
1	22.44	28.84	818.46	80.53	10.236	61.77	505111	89
2	23.03	16.28	1106.18	50.55	22.898	338.82	179923	82
3	22.40	22.01	252.52	64.17	4.032	180.90	505106	62
4	22.96	21.60	557.41	64.97	9.341	296.30	505106	44
5	23.52	23.52	399.38	70.18	5.567	145.98	505106	30
6	22.78	23.12	663.35	65.79	12.654	185.23	179923	26
7	23.58	22.01	509.03	67.19	7.792	183.19	505106	24
8	25.58	30.11	141.60	93.47	1.542	281.08	505106	24
9	23.99	25.79	395.34	78.80	5	474.09	505106	23
10	22.88	17.93	820.68	55.23	18.261	301.60	344304	23
11	23.27	23.73	720.30	70.73	10	443.87	505106	21
12	24.44	21.10	416.62	66.95	5.857	131.56	505106	21
13	21.46	23.79	1452.50	65.02	22.65	335.67	179923	20
14	21.58	17.58	500.43	49.51	10.45	141.84	505106	20
15	21.91	24.23	501.41	68.79	7.4	424.70	505106	20
16	25.56	29.10	236.45	91.94	2.667	104.85	179103	18
17	24.38	24.19	555.92	75.11	7.588	159.24	505106	17
18	25.22	21.29	352.03	68.93	5.667	200.15	505106	17
19	24.75	27.27	249.13	83.10	3.118	147.36	505106	17
20	22.05	21.91	421.19	63.38	6.824	256.99	179103	17
21	24.00	26.56	289.88	80.14	3.706	198.34	505106	17
22	25.52	29.41	483.81	91.61	5.563	163.67	261102	16
23	23.31	21.99	489.75	67.06	7.133	226.60	505106	15
24	25.33	28.15	258.64	88.34	3.429	141.13	505106	14
25	23.59	19.81	416.07	61.52	6.786	158.80	505106	14
26	22.70	21.34	1286.37	63.38	19.571	246.42	179923	14
27	24.00	28.08	222.96	84.13	2.571	160.61	505113	14
28	22.97	28.91	245.33	82.04	2.923	45.46	331601	13
29	21.81	17.65	900.58	52.27	16.75	901.92	179923	12
30	23.19	17.13	1129.52	53.35	21.636	62.06	505106	11
31	22.23	17.60	1099.04	53.29	21.273	379.44	179103	11
32	22.68	24.69	363.51	70.85	5.273	105.45	505106	11
33	21.35	15.54	523.03	46.03	12.727	342.89	505106	11
34	21.75	22.52	200.79	63.34	3.273	145.56	505106	11
35	23.43	16.75	522.45	53.69	9.8	621.20	505106	10
36	23.99	28.50	312.07	85.36	3.8	20.53	349903	10

A growing theoretical literature attempts to deal with pricing decisions where buyers and sellers interact over several periods. Taylor (1998) studies competitive subscription markets, such as credit card markets, long-distance and wireless telephony and internet access. These markets are characterized by repeated consumer purchases, the presence of at least two firms, switching costs for consumers, and the fact that firms can distinguish between new and repeat buyers, enabling them to price discriminate between the two kinds. In his model, firms make "introductory offers" to new consumers, in order to induce them to switch suppliers.[12] He finds that although it is inefficient for consumers to switch when goods or services are homogeneous and changing suppliers is costly, in equilibrium firms price discriminate in this fashion and consumers engage in "supplier surfing".

Fudenberg and Tirole (1998) study monopoly pricing of two overlapping generations of a durable good, where the monopolist may continue to sell an "older" version of a good along with the new one. They show that the decision whether or not to do so, and the conditions that make it optimal to offer "upgrades" to the new good at less than the price of the new good to first-time buyers, depend on the sort of information the monopolist has about the past purchases of its consumers. Fudenberg and Tirole identify two opposite effects: a "reservation-utility" effect where former customers who already own the old generation must be offered a reasonable price to acquire the upgraded version, and a "ratchet" effect whereby former customers have signalled a higher willingness to pay than nonpatrons and can thus be charged a high second-period price.

Both these approaches shed some light on pricing in the wholesale market for steel. As I discuss below, the presence of information asymmetries may imply that there are substantial costs to buyers of switching suppliers, so that it might be possible for the seller to sustain long-term price dispersion despite the presence of competition while at the same time running the risk that some buyers might switch suppliers. Although I have no data on transactions that did not take place, conversations with officials at the firm suggest

[12] Chen (1997) studies the same problem in a two-period duopoly model in which consumer switching costs are distributed uniformly.

that long-term buyers do sometimes take their business elsewhere. On the other hand, the ratchet effect discussed by Fudenberg and Tirole (1998) may be in operation as well: as shown above, there appear to be considerable differences in the pricing to different customers over time, with some paying significantly higher prices than others.

The chapter is organized as follows. Section 2 describes the institutional background and market in which the firm operates. Section 3 describes the data and presents some summary statistics. A more detailed look at the sources of price dispersion is carried out in Section 4, and Section 5 concludes.

2.2 The Steel Market

World production of steel in 1999 was of the order of about 750 million tons. The largest steel-producing countries in the world are China, Japan and the U.S., accounting for about 100 million tons each, with Russia, Germany and South Korea producing 40 to 50 million tons each. According to the American Iron and Steel Institute, 1998 saw the largest U.S. steel imports in history, and 2000 saw the second-highest (37.8 million tons), up 5.9 percent from 1999, with the main sources of imports being the Ukraine, China, India, Turkey, Taiwan and Russia. Table 2.3 lists the the total mill shipments and imports of industrial steel in 1999 and 2000. Table 2.4 lists the major sources of imported steel by region in 1999 and 2000.

The majority of sales by the Company S are to buyers who are classified as wholesalers of steel. Though the firm itself is listed primarily as an importer of steel, it may be considered part of the Steel Service Center (SSC) industry, since it performs a number of

Table 2.3: Mill Shipments (million net tons)

	Total Mill Shipments		Imports	
Category	2000	1999	2000	1999
Carbon	101.5	98.7	32.3	30.9
Alloy	5.4	5.4	4.4	3.7
Stainless & Heat Resisting	2.1	2.1	1.2	1.1
Total	109.0	106.2	37.9	35.7

Source: American Iron and Steel Institute

Table 2.4: Imports by Country of Origin (million net tons)

Country or Region	2000	1999
Canada	5.2	5.1
Mexico	3.3	3.5
Other Western Hemisphere	4.9	5.2
EU	7.1	6.6
Other Europe (including Russia)	6.1	4.4
Asia	9.7	9.2
Oceania	0.9	1.1
Africa	0.7	0.6
Total	37.9	35.7

Source: American Iron and Steel Institute

the functions that would qualify it as such. The Steel Service Center Institute (SSCI) website describes an SSC as "an operation that buys finished steel, often processes it in some way and then sells it in a slightly different form." Service centers thus act as intermediaries between steel producers and finished product producers; often, there is more than one layer of intermediaries between steel producers and final input users.

The SSC industry is a $ 50 billion per year industry. It ships about 27 million tons of steel annually, and employs over 88,000 people. A recent report by the SSCI (available online at http://www.ssci.org/exec_summary.adp) estimates that in 1997, almost a quarter of the total shipments of steel in the U.S. occurred via SSCs rather than through direct sales from manufacturers. There are over 5,000 SSC locations throughout the U.S., with locations in every state. Table 2.5 lists the locations of the top 100 metal service centers. The greatest concentration is in the Great Lakes region (New York, Pennsylvania, Ohio, Indiana, Illinois, etc.) – where the country's metals production and consumption are also concentrated – and the next biggest concentration of SSCs is in the Southeast (Kentucky, West Virginia, North Carolina, South Carolina, Georgia, etc.). Table 2.6 shows the distribution of imports by receiving customs districts. All service centers locate near major highways, and they tend to locate close to each other as well.

Table 2.5: Locations of Top 100 Metal Service Centers in the U.S. and Canada, 1997, percentage of sites

Location	HQ	Centers
Midwest	47	72
Mid-Atlantic	13	35
West	12	30
Northeast	9	27
Southeast	7	44
Southwest	7	36
Plains	1	15
Rockies	0	20
Canada	4	14
Mexico	0	5

Source: Steel Service Center Institute

Table 2.6: Imports by Customs District (million net tons)

District	2000	1999	1998
Atlantic Coast	8.2	8.2	8.0
Gulf Coast - Mexican Border	13.4	11.3	16.2
Pacific Coast	6.9	6.7	6.5
Great Lakes - Canadian Border	8.9	8.9	10.4
Off Shore	0.5	0.6	0.4
Total	37.9	35.7	41.3

Source: American Iron and Steel Institute

SSCs' customers can be divided into three broad groups: fabricators, manufacturers and the construction industry. Their suppliers are domestic and foreign steel mills. The market is highly competitive: the C4 concentration ratio is 14.37, the C10 ratio 26.13 and the C100 ratio 49.63. Table 2.7 lists the major markets to which steel was shipped in 1999 and 2000; note that SSCs themselves form the largest market.

Table 2.7: Mill Shipments to Major Markets (million net tons)

Category	2000	1999	1998
Service Centers and Distributors	30.1	28.1	27.7
Automotive	16.1	16.7	15.8
Construction & Contractors' Products	20.3	18.4	15.3
Machinery, excluding agricultural	3.8	3.9	4.4

Source: American Iron and Steel Institute

SSCs engage not only in distribution of steel, but also in a number of pre-production processing activities. The major activities, along with the SSCI's figures for value added per ton, are listed in Table 2.8.[13]

Table 2.8: Value Added by Steel Service Centers

Item	$ per ton
Storage and shipment	5-15
Slitting	15-45
Cold rolling	150-275
Cutting to length	15-45
Edge rolling	2-8
Plate cutting	15-35
Pickling	14-45
Annealing	20-85
Blanking	85-175
Stamping	75-150
Tempering	35-120
Galvanizing	100-125

Source: Steel Service Center Institute

Firms may be more or less specialized, as Table 2.8 indicates. Company S engages mostly in storage and shipment of plate and pipe, and also in some leveling. Its major product is a particular grade of steel plate, in which it often advertises 'inventory specials'. The firm buys from other intermediaries, and sells both to intermediaries and to end-users such as the construction industry. Table 2.9 reports the SSCI's "cost of possession" worksheet – that is, the cost increase that would be incurred by firms which choose to carry their own metals inventory rather going through an SSC – if the SSCI is to be believed, as

[13] For a detailed description of each of these activities, see the SSCI's website, at http://www.ssci.org/faq_processing.adp.

much as 50 percent!

Table 2.9: Cost of Possession Worksheet

Item	% increase
Inventory Housing	6.8
Inventory Capital Costs	20.0
Equipment Costs	2.5
Scrap Costs	8.5
Manpower Costs	12.3
Total	50.1

Source: Steel Service Center Institute: http://www.ssci.org/cost_possession.adp

The market for steel plate and coil in the U.S operates basically as a dealer or "telephone" market, with SSCs intermediating a substantial share of the trade. Interestingly, there is no commodity market for steel, even though, as argued by Hall and Rust (2002), steel is an example of a standardized commodity that would appear to be an ideal candidate for trading over an exchange. There are no centrally posted prices, and no central market place at which steel is traded. Buyers search for prices by making phone calls to producers or intermediaries in their area who deal in the product they wish to buy. Typically, there is a substantial amount of negotiation involved, so that an initial quote may differ significantly from the final sale price. Hall and Rust (2002) discuss the entry of a potential "market maker," e-STEEL.com, in some detail, and point out that it currently acts as a computerized extension of the existing telephone market. It does not post publicly observable bid/ask prices, but instead appears to act as an intermediary, allowing buyers and sellers to post the quantity, location and price of goods for sale or purchase, with the posted prices often being "first offers" in a negotiation. It does not post historical price data either, which are actually quite hard to come by in any detail: a few trade publications (such as *Purchasing* magazine and *American Metal Market*) report an average "market price" for some categories of steel and some products (including the benchmark product in Group P), but this is compiled with a substantial lag and is indicative at best rather than representative.

2.3 Data

The data set consists of daily data on purchases and sales between July 1997 and June 2000, for a total of 716 business days of data. For every transaction, I observe the quantity (both in units and in weight) of steel bought or sold, the sales price, the material cost, the shipping costs if Company S delivers[14] and the identity of the buyer or supplier. The last of these, identity, is the unique feature of the data set, for it enables me to track individual buyers over time. Accordingly, I have data on the buyers in the sample, regarding their location, line of business and firm size.

2.3.1 Sales data

The company buys and sells over 3,000 products, in the form of steel coil and plate. In the interests of clarity, I focus initially on a particular product, which I will call product P2. This is one of the products in a class of products ("Group P") which form a sort of industry benchmark. These products tend to be sold in fixed dimensions, with little or no production processing by the firm. Chan (2001a) shows that, even within Group P, there is a substantial amount of price variation across products.[15] Products in Group P are widely

[14]Delivery is carried out through an independent transportation firm located next to Company S. In about twenty percent of the transactions, the buyer picks up the steel at Company S.

[15]With one qualification: prices across products in Group P are usually correlated *by customer*. That is, the by-weight prices quoted to a particular customer for different products in the group are usually the same if the purchases are made on the same day (or, more generally, within a few weeks, following the within-product stickiness in price).

produced and traded; product P2 is one of the highest-volume products the firm deals in, with 1195 sales taking place over the sample period.

Prices for each product are in dollars per hundredweight of steel;[16] percentage margins are computed over the cost figure supplied by the firm. The cost figure deserves some explanation. Cost figures, which are recorded separately for each transaction, are indexed by the location from which the firm brings into its plant the steel it ships to that particular customer, and are calculated on a first-in-first-out basis. The recorded cost of the product is thus actually the historical cost at which the first unit in the inventory was bought. As such, the material cost figure over which the margin is computed for each transaction can vary in the course of a single day. For instance, consider the ordering of net profits on business day 394 in Table 2.1; the anomalous item is a transaction for which the quoted material cost is higher than for some of the other transactions on the same day (there are other transactions with the same quoted cost and still lower prices).[17] This raises two issues. The first is that simply considering the sale price when quantifying dispersion may be somewhat misleading, inasmuch as the margin over recorded cost may be considerably higher or lower relative to another transaction at the same price depending on the cost figure used. However, over the course of the sample this appears to smooth out: going back to Figures 2.1 and 2.2, the distribution of margins mirrors the distribution of net prices quite closely, and the results do not vary much when margins are computed over a daily average cost figure (the difference between the recorded cost and the daily average cost is over a dollar in less than ten percent of the transactions). Moreover, as I argue below,

[16] A unit of product P2 weighs about 1600 pounds, or 16 hundredweight; at \$ 23 per unit, this translates into roughly \$ 300 per ton.

[17] On a few occasions, the firm sends the product out for some pre-sale processing; the processing cost is included in the recorded cost, and I am able to identify and discard these transactions.

the margin figure itself is likely to be somewhat unreliable, and therefore net price is the relevant variable to consider in attempting to quantify the extent to which buyers could do better.

The second issue relates to the way the firm's incentive scheme is structured. Sales are carried out by a number of sales agents, each serving a particular geographical area. Individual agents tend to interact with particular customers on an extended basis, and do not compete with each other in making sales to individual customers. Commissions are calculated based on the profit over the recorded material cost. As noted above, the cost is based on the price at which past purchases were made. Company S buys relatively infrequently, making only 120 purchases during the sample period.[18] As a result, the historical cost can differ substantially from the current spot price, so that if sales agents make decisions with respect to the replacement cost - whether current spot price or the expected future price - rather than the historical cost, then the margin figure is likely to be subject to significant measurement error. The dispersion in costs for the reasons outlined above only exacerbates this problem; different agents may be pricing above different cost figures on the same day, although in reality the shadow price of steel might be quite different from either or all of these cost figures. Nevertheless, the extent of short-run dispersion in prices is sufficient to suggest that there is more at work than simply different perceptions of the right cost figure. Since the firm itself considers the recorded cost when setting prices and computing profits, I would argue that the cost figure is a relevant variable in studying the firm's pricing behavior, even if not the correct variable for the firm to be using in providing

[18] Hall and Rust (2000) investigate the firm's inventory decision in detail.

incentives.[19]

The dispersion in costs actually offers an interesting insight. A number of the unusually high-cost transactions result in margins that are much lower than average (typically between six and ten percent, and sometimes as low as three percent, with the average at eleven percent). Moreover, most of the lower margin transactions are to a small subset of customers. This appears to be partly due to the fact that prices are sticky in the short run, so that if the firm is being forced to fill an order at higher cost than usual, it absorbs the difference rather than passing it on to the customer. While some of this is an artifact of the firm's accounting procedures, it is also possible that it may be absorbing some of the rise in costs in order to retain customers. Discussions with firm officials suggest that customer "loyalty" has been on the decline in the past few years, so that the firm is being forced to work on smaller margins and larger quantities, possibly indicating that it is absorbing a fair amount of risk on account of the purchasing cycles, without necessarily being compensated for all of it. Still, as Hall and Rust (2002) observe, there appear to be incentives for both sides to maintain relationships, and informal conversations also suggest a reluctance on the part of executives to disappoint an existing customer.

As noted above, the market is basically a "phone market," where individual customers negotiate with their assigned agent over a price. Prices are not posted publicly – at most, the firm's web site indicates if there are products on which it is offering 'specials', but prospective customers must call the company for a quote. Prices to different customers on a day can vary by as much as seven or eight dollars per hundredweight of steel, so that the

[19]See Chan (2001b) for an investigation of the costs to the firm of providing incentives based on historical cost rather than replacement cost.

percentage margin on any given date can vary from 2 to 30 percent. The highest margin observed in the sample is 87 percent. Table 2.10 summarizes the sales data for product P2: prices are quoted in dollars per hundredweight, revenues and profits in dollars and margins as percentages over unit cost. Figure 2.5 indicates the distribution of margins over the course of the sample. Figure 2.6 plots the distribution of units sold for all sales of less than forty units (all but two of the sales); most sales are quite small, with the mean sale being ten units and the median sale being only six units.

Table 2.10: Summary Statistics for Sales of Product P2

Variable	Mean	Std. Deviation	Minimum	Maximum
units	9.09	8.86	1	107
distance	315.46	333.91	2	2427.341
gross price	24.26	2.33	18.5	31.5
net price	23.29	2.32	18	29.65
unit cost	18.97	2.59	15.65	26.2
gross revenue	3526.78	3344.38	339.05	32948.86
net revenue	3379.45	3208.88	308.82	31637.89
gross profit	711.78	649.88	34.32	4719.47
net profit	564.45	523.33	15.84	3543.42
gross margin	28.92	10.66	1.88	92.99
net margin	23.72	10.29	1.87	87

Since we only observe transactions that actually take place, it is not possible to isolate factors which might determine a customer's decision to buy from those which account for how much a given customer is willing to pay, and how much the firm calculates it can charge them and still complete a sale. For instance, if Company S has local monopoly power, as seems likely, it could price differentially on the basis of location, on the theory that customers in the surrounding areas are less able or likely to search for or buy from competing firms; in such a case the customer might both be more likely to buy from the firm, and may be willing to pay a higher price instead of having to incur search costs.

2.3.2 Customer information

During the period under consideration, 157 customers appear in the sample. Of these, thirty six customers buy ten times or more, fifteen buy twenty times or more and five buy over thirty times. For each customer, I have data on a number of variables, including the firm's location, primary and secondary SIC code, number of employees, credit rating, and in some cases, sales volume.[20] I make use of the information on the firm's line of business as indicated by the SIC code, and I use the number of employees as a proxy for firm size. I do not use sales volume as a proxy for firm size, as it is unavailable for a number of customers. The credit rating is available in three categories, roughly indicating 'very good', 'good' and 'bad'; most of the buyers for whom I have information fall into the first two categories, leading me to suspect that the rating does not actually convey much information. Moreover, this credit rating variable is compiled by Reference USA, not by Company S, so I do not believe it is an important determinant of the pricing decision.[21]

Buyers are located mostly in Connecticut, Massachusetts, New Jersey, New York and Pennsylvania. The average distance of the buyers in the sample is about 300 miles: Figure 2.7 summarizes the distribution of distances. There are buyers located as far away as Illinois, Oklahoma and Texas, but they account for less than half a percent of total sales, and the median distance is about 200 miles. Figure 2.8 plots the distribution of sales by distance; most of the sales occur to customers in Connecticut, Maine, Massachusetts, New Hampshire, New Jersey, New York and Pennsylvania.

[20] The data were acquired from Reference USA, a web site, and supplemented by information from the Yellow Pages.

[21] For another view, see Chan (2001b) for an investigation of customer characteristics and bargaining power in this market.

As mentioned earlier, the biggest group of customers is under the two-digit SIC code of 50 (Durable Goods, Wholesale). Table 2.11 summarizes the major SIC code groups for buyers of product P2. Of these, Steel Distributors & Warehouses account for eleven of the fifteen most frequent customers and sixty percent of the sales in the sample period. The next biggest category is Construction & Specialty Contractors, comprising thirty seven customers, of whom thirty are classified as Steel Fabricators. This group includes three of the fifteen most frequent customers and accounts for a quarter of the sales. The remaining customers are mostly classed as Primary Metal Industries, and Fabricated Metal Products.

Table 2.11: Major SIC code groups, Product P2

SIC Code	Number of customers
Steel Dist & Warehouses (505106)	64
Steel Fabricators(179103)	30
Primary Metal Industries (33)	8
Fabricated Metal Products (34)	17

Buyers are mostly small, both by employee size and by sales volume. The one buyer in the 500-plus employee category is actually a division of a major corporation. The distribution of customer sizes is summarized in Table 2.12.

Table 2.12: Customer Firm Size by Number of Employees

Employee range	Number of firms
1 - 4	4
5 - 9	19
10 - 19	29
20 - 49	47
50 - 99	29
100 - 249	19
250 - 499	4
500 - 999	1

2.4 Sources of Price Dispersion

Our preliminary look at the data in section 1 suggested that the firm offers quantity discounts, and that prices are influenced by location and other customer characteristics as well. Here I analyze the effect of these variables in greater detail. Figure 2.9 plots the average net price, average per unit profit and average net margin respectively against the mean number of units bought for the 36 most frequent customers listed in Table 2.2 (who account collectively for 819 of the 1195 observations). Each observation is plotted using the customer's rank by frequency of purchase. Notice that while the relative positions of most customers are quite consistent (including customers 3 and 8 in the top ten customers), there are notable exceptions (such as customers 1, 2 and 10). The relationship is broadly negative, and quite pronounced for the case of the per unit net profit and net margin. Figure 2.10 plots these quantities against the distance from Company S for the same subset of customers, and the resulting graphs look quite similar; again, the relative positions of customers change between the net price and net margin graphs, due in part to the firm's accounting procedures but also possibly indicating a degree of customer-specific pricing.

I estimate an ordinary least squares specification with net price (in dollars per hundredweight) to each customer as the dependent variable. Independent variables include a constant, per unit cost, quantity purchased (in hundredweight) and a time trend, where each time dummy corresponds to a period of about 60 business days. Results are displayed in Table 2.13. Not surprisingly, this picks up a large part of the variation, since prices track costs over the course of the whole sample. Notice that the coefficient on t6 is actually larger than that on t5, despite the fact that this period is *after* the Russian crisis; it is not until t7 that prices begin to fall, and they seem to "overcorrect" themselves between business day 480 and 540 (period t9) before going back up following a slight rise in costs. Quantity is significant; roughly, buying an extra unit would mean a discount of 50 cents per hundredweight bought, or \$ 8 less per unit, so that a customer who bought the average of ten units would pay \$ 80 less per transaction than if he were charged the price paid by a customer who bought just one unit.

Table 2.13: OLS with Time Dummies

DEPENDENT VARIABLE: NET PRICE			
Variable	Estimate	Std. Error	t-value
constant	15.91	0.52	30.6
unit cost	0.339	0.031	10.86
hundredweight	-0.0033329	0.0002274	-14.66
t1	2.323	0.236	9.86
t2	2.569	0.203	12.64
t3	2.315	0.225	10.26
t4	2.362	0.237	9.95
t5	2.371	0.223	10.63
t6	2.898	0.171	16.87
t7	1.293	0.179	7.19
t8	0.278	0.159	1.74
t9	-0.4866	0.157	-3.094
t10	0.2033	0.17	1.19
t11	0.132	0.169	0.78
Number of observations		1195	
Adjusted R^2		0.77	
F(13,1181)		303.97	

Table 2.14 adds customer characteristics to this regression, in the form of distance in miles from Company S, firm size as proxied for by the number of employees, and line of business as indicated by the SIC code. Distance is significant; every extra 100 miles would result in a discount of 7 cents per hundredweight off the net price, or \$ 1.10 off per unit, so that a buyer located at the median distance of 200 miles would pay on average \$ 2.20 less per unit, or about \$ 20 less per transaction if they bought ten units. Firms with less than ten employees seem to do well, as do firms with between 20 and 100 employees. As

suggested above, steel distributors also do better, as do makers of fabricated metal products (SIC code 34).

Table 2.14: OLS with Customer Characteristics and Time Dummies

DEPENDENT VARIABLE: NET PRICE			
Variable	Estimate	Std. Error	t-value
constant	16.719	0.540	30.93
unit cost	0.3444	0.0301	11.14
hundredweight	-0.0032468	0.0002649	-12.26
distance	-0.000679	0.00016	-4.12
$1 \leq$ emp≤ 4	-0.546	0.265	-2.06
$5 \leq$ emp≤ 9	-0.457	0.222	-2.06
$10 \leq$ emp≤ 19	-0.203	0.212	-0.96
$20 \leq$ emp≤ 49	-0.627	0.202	-3.11
$50 \leq$ emp≤ 99	-0.712	0.204	-3.5
$100 \leq$ emp≤ 249	-0.332	0.205	-1.62
Dist/Warehouses	-0.314	0.13	-2.41
Fabricators	-0.173	0.155	-1.12
Primary Metal	-0.212	0.234	-0.91
Fab Metal Products	-0.419	0.179	-2.33
Contractors	-0.286	0.170	-1.69
t1	2.323	0.231	10.06
t2	2.529	0.199	12.65
t3	2.268	0.221	10.25
t4	2.381	0.233	10.22
t5	2.381	0.218	10.92
t6	2.908	0.169	17.25
t7	1.315	0.176	7.48
t8	0.295	0.157	1.87
t9	-0.499	0.155	-3.22
t10	0.272	0.167	1.63
t11	0.161	0.166	0.96
Number of observations		1195	
Adjusted R^2		0.78	
F(25,1169)		169.36	

In Table 2.15 I add customer dummies for the top ten most frequent buyers; customers 3 and 10 do significantly better than the others, and customer 8 appears to be paying over a dollar per hundredweight more than the others in the sample, despite being a distributor; interestingly, customer 2 does not appear to do significantly better once quantity and distance are accounted for. Distributors pay less here as well, though other SIC code groups do not seem to be at a significant advantage; also, only firms with between 20 and 100 employees seem to have a significant price advantage over firms of other sizes. There are two possible reasons why distributors might pay lower prices in general. The first is that they tend to buy larger quantities on average, accounting for a large part of the sample, so that they might have an incentive to undertake search; if Company S expects them to be better informed regarding the market price of steel, it is likely to charge them less than buyers who are considered less well-informed. The second possibility is that distributors, by the nature of their business, tend to have more elastic demands than, say, construction companies or steel fabricators, leading to lower reservation prices and hence lower sale prices.

Table 2.15: OLS with Distance and Customer Characteristics, Product P2

DEPENDENT VARIABLE: NET PRICE			
Variable	Estimate	Std. Err.	t
constant	16.719	0.532	31.42
unit cost	0.3433	0.0304	11.26
hundredweight	-0.0030876	0.000268	-11.52
distance	-0.0008348	0.000172	-4.83
$1 \leq$ emp≤ 4	-0.4376	0.3479	-1.26
$5 \leq$ emp≤ 9	-0.4057	0.2206	-1.84
$10 \leq$ emp≤ 19	-0.2387	0.22005	-1.085
$20 \leq$ emp≤ 49	-0.4807	0.2031	-2.37
$50 \leq$ emp≤ 99	-0.6396	0.2029	-3.15
$100 \leq$ emp≤ 249	-0.2696	0.2283	-1.18
Dist/Warehouses	-0.3335	0.1367	-2.44
Fabricators	-0.1900	0.1565	-1.21
Primary Metal	-0.2240628	0.2309	-0.97
Fab Metal Prod	-0.2306	0.1949	-1.18
Contractors	-0.1155	0.2089	-0.55
t1	2.266	0.223	9.86
t2	2.485	0.197	12.63
t3	2.262	0.219	10.35
t4	2.378	0.230	10.33
t5	2.353	0.215	10.94
t6	2.8720	0.166	17.28
t7	1.328	0.174	7.64
t8	0.240	0.155	1.56
t9	-0.498	0.153	-3.26
t10	0.248	0.164	1.51
t11	0.134	0.164	0.81
Cust1	0.066	0.202	0.32
Cust2	-0.240	0.229	-1.05
Cust3	-0.624	0.159	-3.90
Cust4	-0.020	0.186	-0.10
Cust5	-0.144	0.213	-0.68
Cust6	-0.239	0.409	-0.58
Cust7	-0.072	0.233	-0.30
Cust8	1.167	0.244	4.78
Cust9	0.469	0.242	1.94
Cust10	-0.618	0.288	-2.14
Number of observations			1195
Adjusted R^2			0.787
F(35,1159)			127.34

Table 2.16 reports the results for a regression of net margin on the same subset of variables; the results look quite similar, but notice that the fit is significantly worse.

In Table 2.17 I report the results for a regression of net price on the same subset plus additional dummies for nearby states, in which most of the customers are located. The

Table 2.16: OLS with Distance and Customer Characteristics, Product P2

DEPENDENT VARIABLE: NET MARGIN			
Variable	Estimate	Std. Err.	t
constant	106.099	3.004	35.32
unit cost	-4.284	0.172	-24.89
hundredweight	-0.0169658	0.001535	-11.21
distance	-0.0042675	0.000975	-4.38
$1 \leq$ emp≤ 4	-2.517	1.964	-1.28
$5 \leq$ emp≤ 9	-2.082	1.245	-1.67
$10 \leq$ emp≤ 19	-1.555	1.242	-1.25
$20 \leq$ emp≤ 49	-2.723	1.147	-2.38
$50 \leq$ emp≤ 99	-3.543	1.146	-3.10
$100 \leq$ emp≤ 249	-1.519	1.289	-1.18
Dist/Warehouses	-1.777	0.772	-2.30
Fabricators	-1.190	0.884	-1.35
Primary Metal	-1.261	1.303	-0.99
Fab Metal Prod	-1.178	1.10	-1.07
Contractors	-0.980	1.179	-0.83
t1	9.853	1.298	7.59
t2	10.808	1.111	9.73
t3	10.281	1.235	8.32
t4	10.658	1.30	8.2
t5	10.617	1.214	8.75
t6	13.586	0.938	14.48
t7	5.462	0.982	5.56
t8	1.573	0.877	1.79
t9	-3.145	0.862	-3.65
t10	0.983	0.925	1.06
t11	0.385	0.924	0.42
Cust1	0.644	1.143	0.56
Cust2	-0.717	1.297	-0.55
Cust3	-3.45	0.902	-3.83
Cust4	-0.420	1.050	-0.4
Cust5	-0.473	1.203	-0.39
Cust6	-0.065	2.31	-0.03
Cust7	-0.378	1.318	-0.29
Cust8	6.295	1.378	4.57
Cust9	2.228	1.365	1.63
Cust10	-3.344	1.628	-2.06
Number of observations			1195
Adjusted R^2			0.655
F(34,899)			65.88

distance variable gets subsumed by the state dummies; customers in Massachussetts and New Hampshire appear to be paying more, while customers in New Jersey pay less, which is consistent with there being considerably more competition in New Jersey than there is to the north. Also, Company S appears to have substantial local market power, with

Table 2.17: OLS with Distance and Location Dummies, Product P2

DEPENDENT VARIABLE: NET PRICE

Variable	Estimate	Std. Err.	t
constant	16.397	0.55	29.82
unit cost	0.324	0.03	10.84
hundredweight	-0.003014	0.000267	-11.3
distance	-0.0002828	0.0002306	-1.23
$1 \leq$ emp≤ 4	-0.176	0.355	-0.49
$5 \leq$ emp≤ 9	-0.171	0.234	-0.73
$10 \leq$ emp≤ 19	-0.0006	0.229	-0.002
$20 \leq$ emp≤ 49	-0.337	0.208	-1.62
$50 \leq$ emp≤ 99	-0.435	0.217	-2.001
$100 \leq$ emp≤ 249	-0.136	0.236	-0.58
Dist/Warehouses	-0.177	0.139	-1.28
Fabricators	0.034	0.155	0.22
Primary Metal	-0.685	0.237	-2.89
Fab Metal Prod	-0.174	0.202	-0.86
Contractors	0.002	0.216	0.008
t1	2.263	0.224	10.09
t2	2.526	0.192	13.17
t3	2.309	0.213	10.83
t4	2.420	0.225	10.78
t5	2.408	0.21	11.47
t6	2.879	0.162	17.82
t7	1.307	0.169	7.72
t8	0.193	0.151	1.27
t9	-0.535	0.148	-3.61
t10	0.241	0.160	1.50
t11	0.152	0.161	0.95
Cust1	-0.729	0.268	-2.72
Cust2	0.068	0.529	0.13
Cust3	-0.90	0.212	-4.24
Cust4	0.41	0.262	1.57
Cust5	-0.235	0.218	-1.07
Cust6	0.482	0.432	1.11
Cust7	-0.348	0.268	-1.29
Cust8	1.141	0.274	4.17
Cust9	0.858	0.292	2.94
Cust10	-0.477	0.292	-1.63
conn	1.112	0.233	4.78
delaware	-0.195	0.512	-0.38
maine	-0.286	0.222	-1.29
mass	0.383	0.18	2.13
newhamp	0.541	0.215	2.52
newjersey	-0.538	0.204	-2.62
newyork	0.15	0.174	0.86
penn	0.167	0.171	0.98
Number of observations			1195
Adjusted R^2			0.8
F(43,1151)			112.17

customers located in Connecticut paying a premium (although half the sales in Connecticut are accounted for by just one buyer). Distributors do not appear to do any better than other groups once location is accounted for; it is likely that between the state dummies and individual customer dummies, most of the larger customers – who are distributors – are accounted for. Firm size mostly does not matter either, while individual customer dummies are still significant for some of the customers.

What is interesting is that twenty percent of the dispersion in prices (and almost thirty five percent of the variation in margins) is still unaccounted for, and that there appear to be significant idiosyncratic effects in pricing to individual buyers that are hard to capture fully in summary variables such as firm size and line of business. To explore this further, I now look briefly at some of the other high-volume products in Group P. Table 2.18 provides summary statistics on all buyers who make more than 20 purchases of any of the major Group P products (products P1, P2, P3, P5, P7, P10, P12, P15). I rank customers by the number of times they buy these Group P products (rankings for product P2 are given in parentheses for those customers who buy P2). I also indicate the mean quantity bought in hundredweight, since different products in the group have different weights, which makes comparison by number of units unreliable.

Consider customers 1 and 2 in the table. Customer 2 pays an average price of $2.50 per hundredweight more than customer 1; at an average of 101 hundredweight per transaction, this translates into an extra $ 250 per transaction, or an extra $ 57,600 over the course of the sample period. Similarly, customer 5 (customer 1 in our original sample of P2) is "overpaying" relative to customer 1 by $ 2.70 per hundredweight on average, which

Table 2.18: Summary Statistics for Frequent Buyers, Group P Products

Customer	Mean net margin	Mean net price	Mean units (Group P)	Mean hundwt (Group P)	distance	SIC code	number of purchases
1 (3)	12.47	20.66	2.97	123.91	180.9	505106	598
2 (16)	16.11	23.06	1.62	101.82	104.85	179103	236
3 (21)	12.65	20.8	1.96	88.14	198.34	505106	229
4 (5)	15.05	21.45	2.89	92.70	145.98	505106	135
5 (1)	26.89	23.36	9.64	146.39	61.77	505111	132
6 (7)	16.71	22.53	5.69	131.06	183.19	505106	128
7	11.39	21.96	5.13	216.56	251.63	344403	127
8 (4)	18.19	23.17	8.74	147.20	296.3	505106	98
9 (2)	15.74	23.24	23.15	378.26	338.81	179923	87
10 (59)	9.84	21.86	1.15	79.76	9.7	331204	68
11 (90)	11.26	22.15	1.07	83.49	19.36	354498	54
12 (6)	23.12	22.78	12.65	206.76	185.23	179923	26
13 (8)	30.11	25.58	1.54	25.19	281.08	505106	24
14 (9)	25.79	23.99	5	81.70	474.09	505106	23
15 (10)	17.93	22.88	18.26	298.38	301.6	344304	23
16 (50)	21.39	20.52	3.36	109.89	0	179103	22
17	10.87	21.63	10	653.40	126.16	509315	22
18 (94)	13.2	20.7	1.29	63.01	287.79	505106	21
19 (11)	23.73	23.27	10	163.40	443.87	505106	21
20 (12)	21.1	24.44	5.86	95.71	131.56	505106	21

totals to an extra \$ 52,000 over three years.

These magnitudes are surprisingly large, especially over the long-term. One is led to wonder why Company S is able to sustain the dispersion in prices, particularly when the more frequent customers buy sufficient amounts to warrant investing in search. One possibility is that, despite the apparent homogeneity of the product, buyers are paying for more than simply the steel itself. Hall and Rust (2002) point out that in the steel market, timely delivery and payment and quality assurances are of primary importance; from conversations with steel executives, they observe that many sellers of steel do not deliver the goods when promised and that the steel can be of lower quality than promised. Some fraction of the units in an order can exhibit non-uniformities or defects, so that buyers may be disinclined to shop around for fear of encountering "lemons" problems. Company S itself stresses the importance of just-in-time delivery and stock availability.

Conversations with officials at the firm also indicate that the nature of Company S's business exposes it to significant risk, not just in terms of the aggregate risk of fluctuating world prices for steel, but also because there can be substantial credit risks involved in selling to some customers. In fact, customer 8 in the P2 sample, though apparently overpaying, appears to have cost the firm about $ 20,000 through a failure to pay his bills. The firm claims that its pricing and delivery policies take individual credit histories into account; in the case of some customers, for instance, steel is not shipped until payment is received, and customers who are poorer credit risks tend to be charged higher prices. Customers actually have an incentive to make large orders, as they can make use of inventory to secure loans; on the other hand, Company S's inventory is unsecured, and it can often be exposed to the risk its customers assume as a result.

The observed inefficiencies in the market are almost certainly due in part to the old-fashioned nature of the steel market. Hall and Rust (2002) study the wholesale fish market, where the emphasis on timely delivery and payment and quality assurance is similar (not surprisingly, given the nature of the commodity). They point out that gofish.com, the fish market's equivalent of a centralized web site, differs from e-STEEL.com in one crucial way: transparency of transactions. Price information is publicly available to all market participants; in addition gofish.com has set up a grading system where buyers have an opportunity to grade suppliers on every transaction. This system acts as a screening device and a powerful incentive to ensure delivery and quality as promised. There is also a conflict resolution department to resolve disputes. On the other side, gofish.com has ties to Seafax, a major provider of credit information on firms in the wholesale fish business; this enables sellers to look up the credit history of potential buyers before agreeing to a sale. In addition, buyers on gofish.com acquire a line of credit with gofish.com; this involves a credit check. If the buyer does not pay, gofish.com covers the debt, and the seller knows how much gofish.com will cover in advance of the transaction's taking place. By contrast, e-STEEL.com's response to buyers who don't pay is to give the seller the option of restricting which potential buyers can view the seller's postings. It does not perform credit checks on members and does not provide a payment/delivery history of members either. As a result, for firms operating in this market, it becomes quite important to know the identity of the buyer or seller one is dealing with. Hall and Rust (2002) also note that there appears to be considerable resistance on the part of executives in the steel market to the possibility of centralized market-making. Some of this is clearly due to the fact that a successful market maker would probably drive a significant number of SSCs out of the market, leaving lower profits to the survivors as a result of the reduction in bid/ask spreads.

2.5 Conclusion

This chapter explores a rich data set on transactions in the wholesale market for steel, with the objective of investigating the existence and extent of dispersion in prices. I find considerable dispersion in prices, both in the short term and over the long run. Apart from quantity discounts and local market power, the identities of the buyers affect the pricing decision of the seller, and there also seems to be a premium to information. Discussions with steel executives indicate that information asymmetries might be contributing to the observed dispersion in prices. The analysis suggests that there are substantial inefficiencies in the market for steel, and that buyers might benefit considerably from increased transparency in the market as a whole and lower barriers to search.

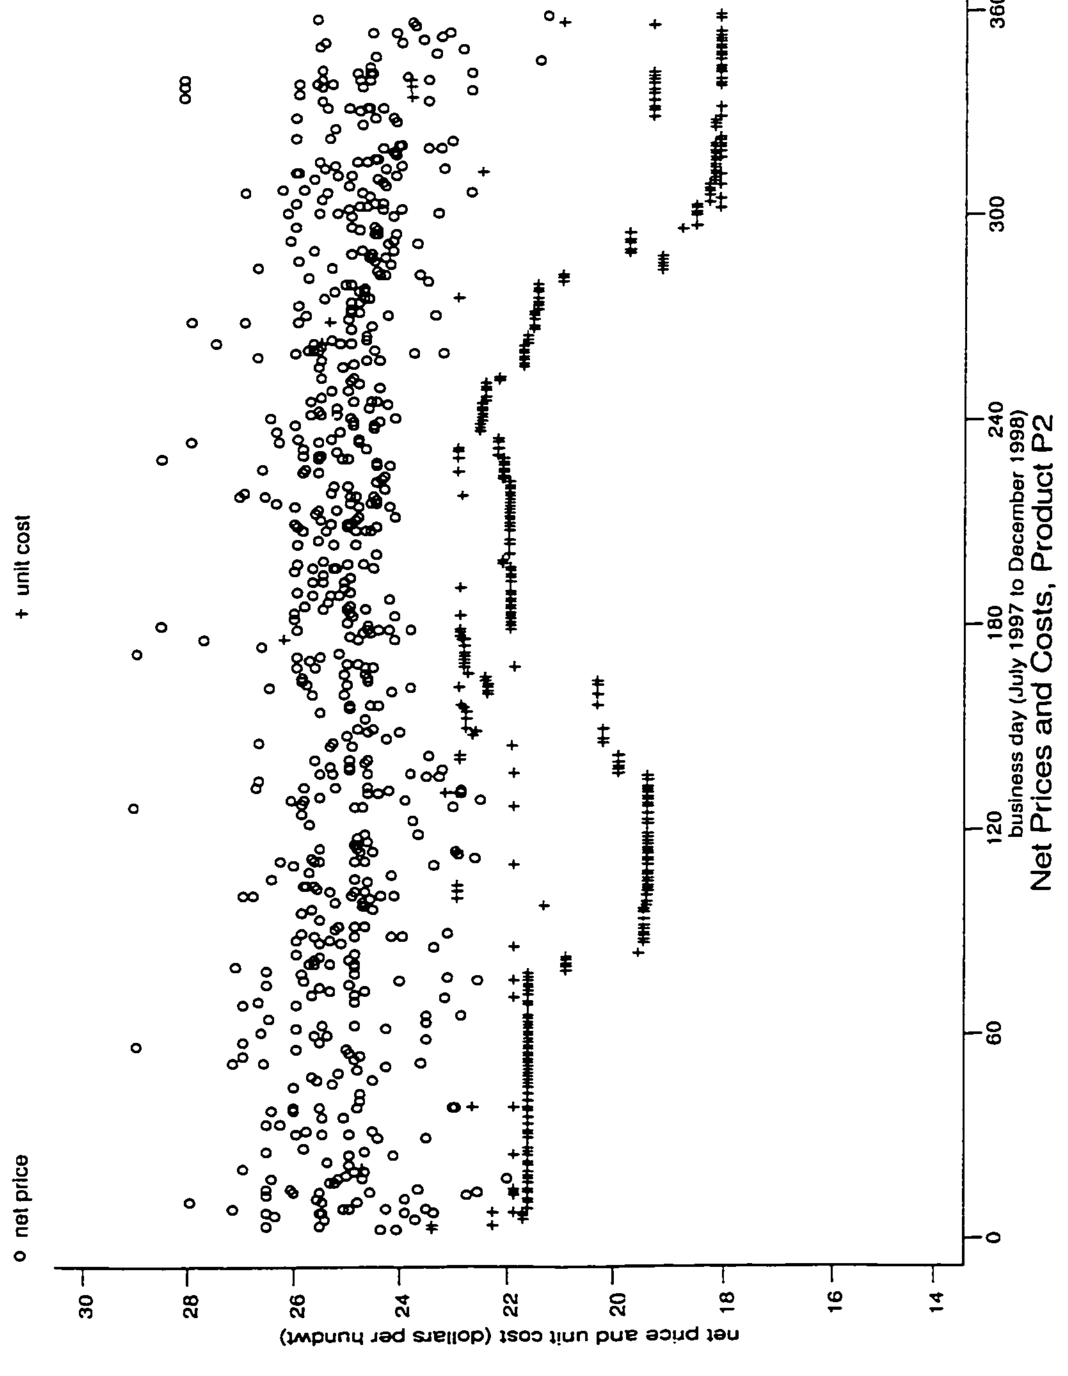

Figure 2.1: Prices and Costs for Product P2, July 1997 to December 1998

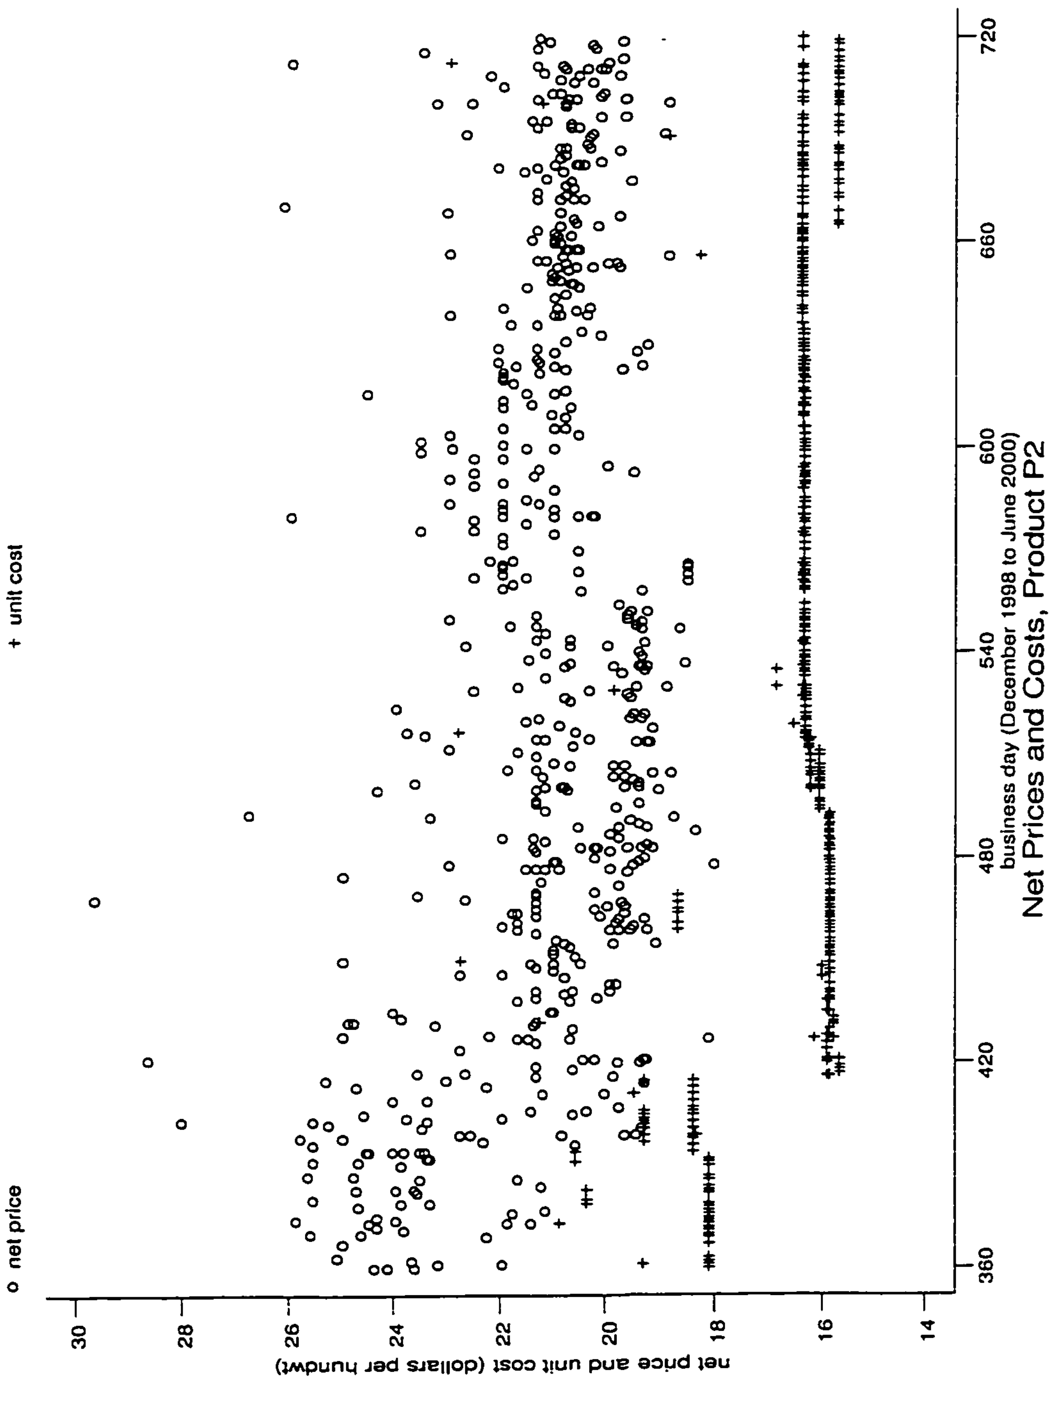

Figure 2.2: Prices and Costs for Product P2, December 1998 to June 2000

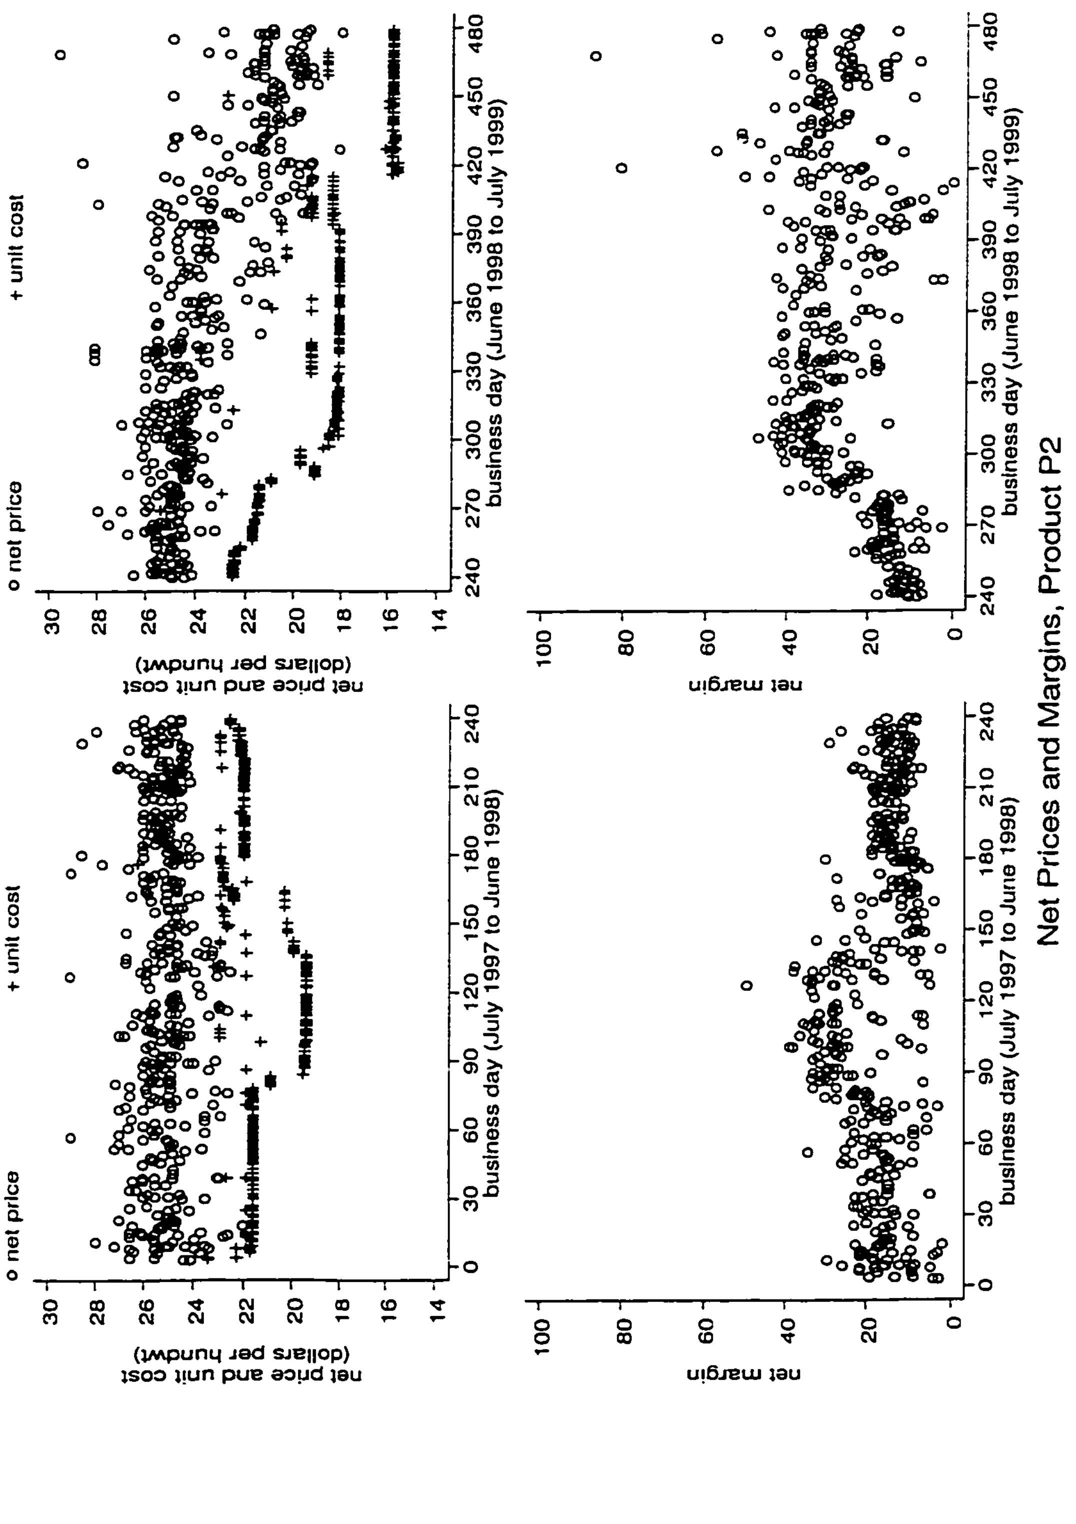

Figure 2.3: Prices and Margins for Product P2, July 1997 to July 1999

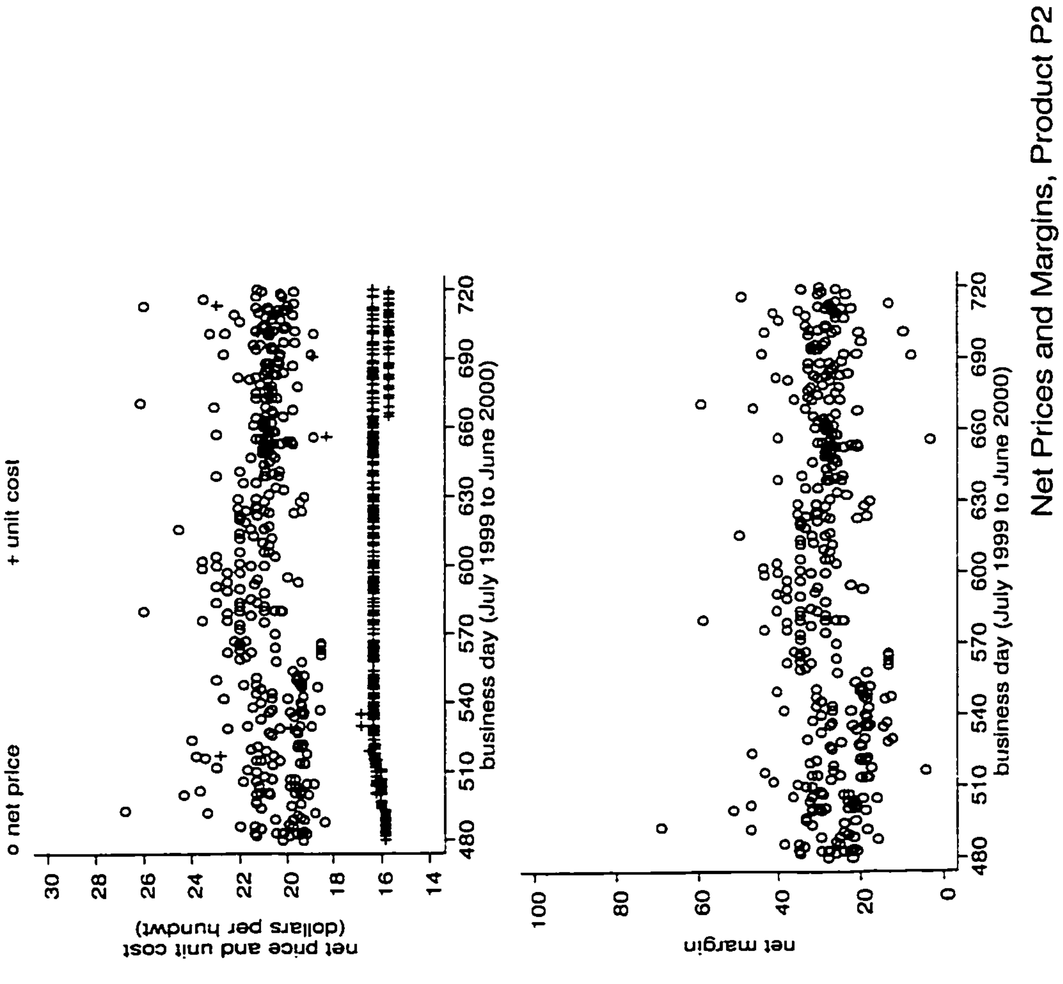

Figure 2.4: Prices and Margins for Product P2, July 1999 to June 2000

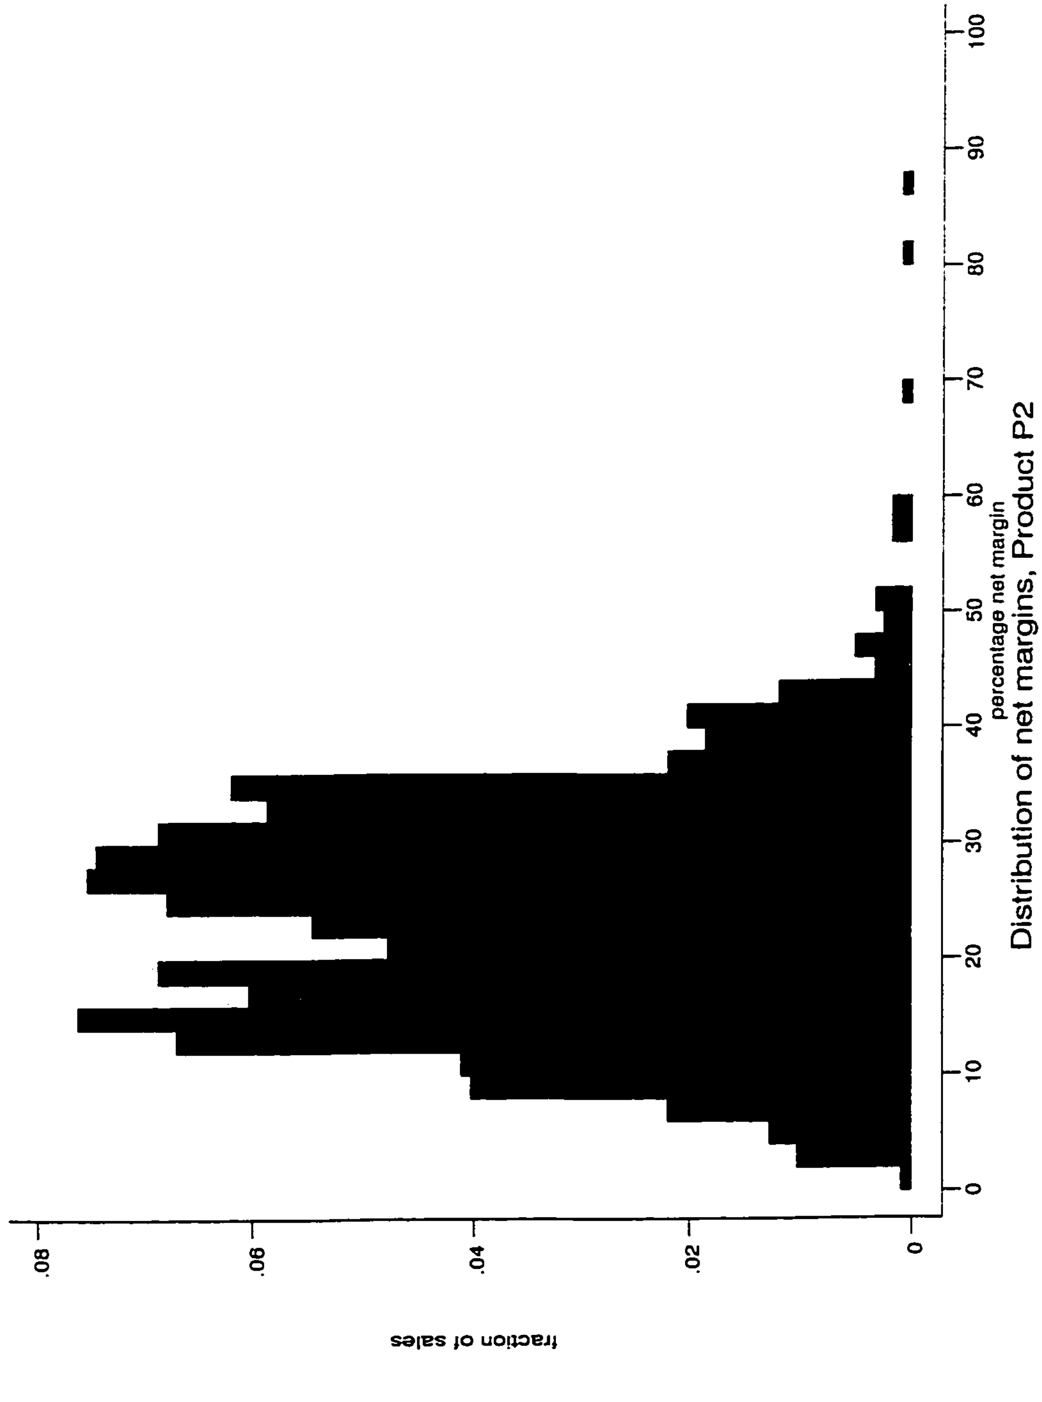

Figure 2.5: Distribution of Net Margins for Product P2

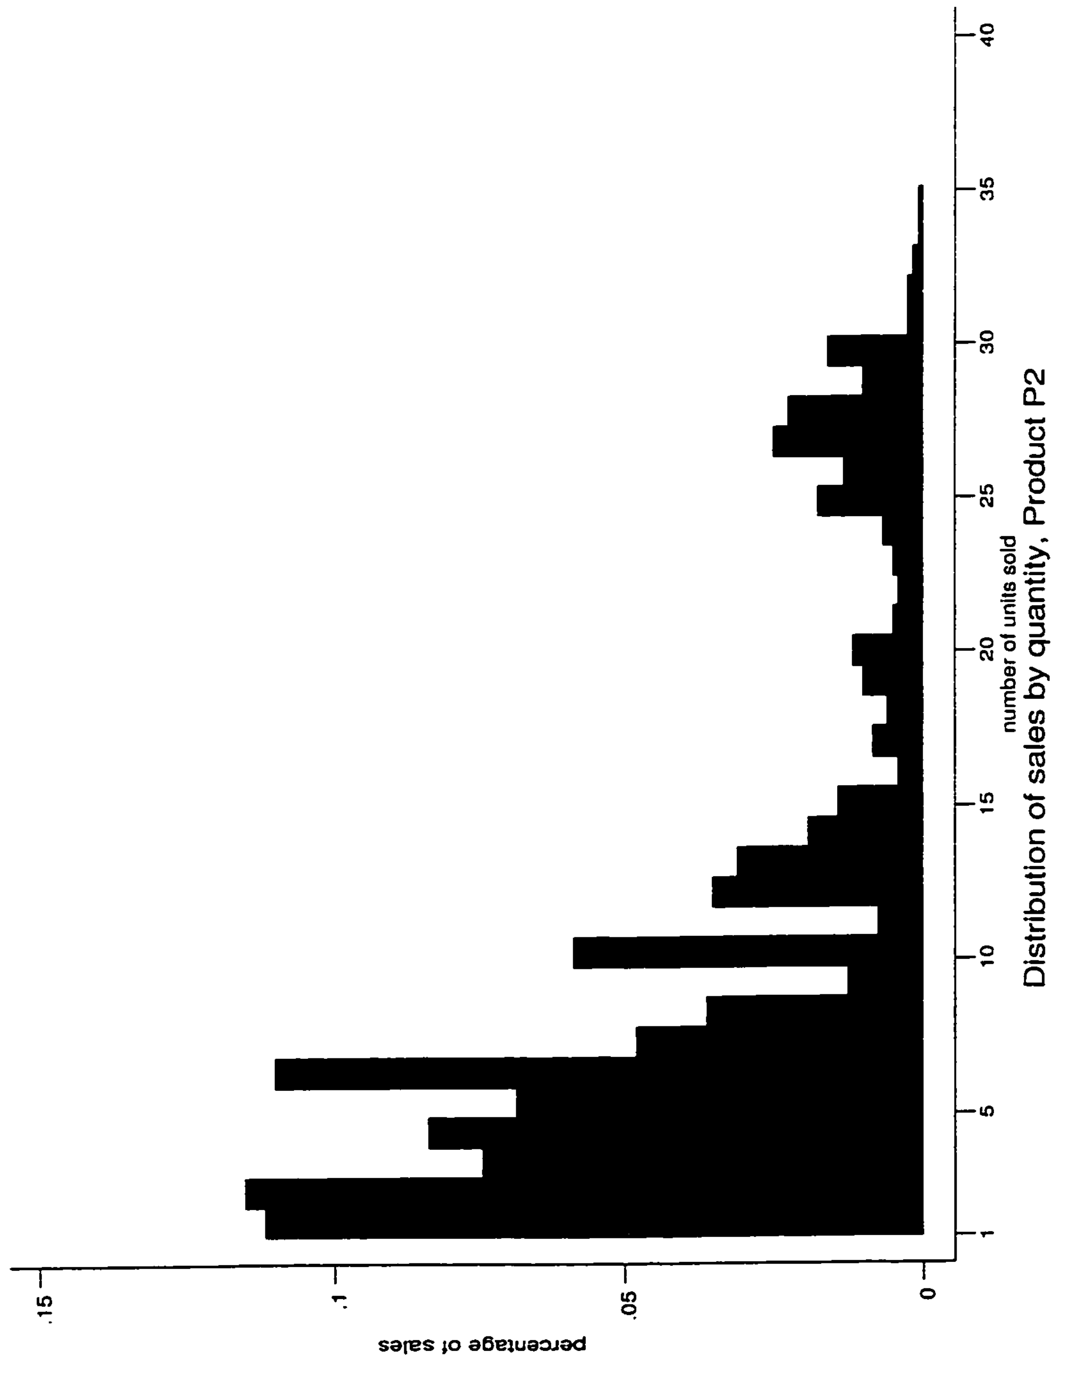

Figure 2.6: Distribution of Sales by Quantity, Product P2

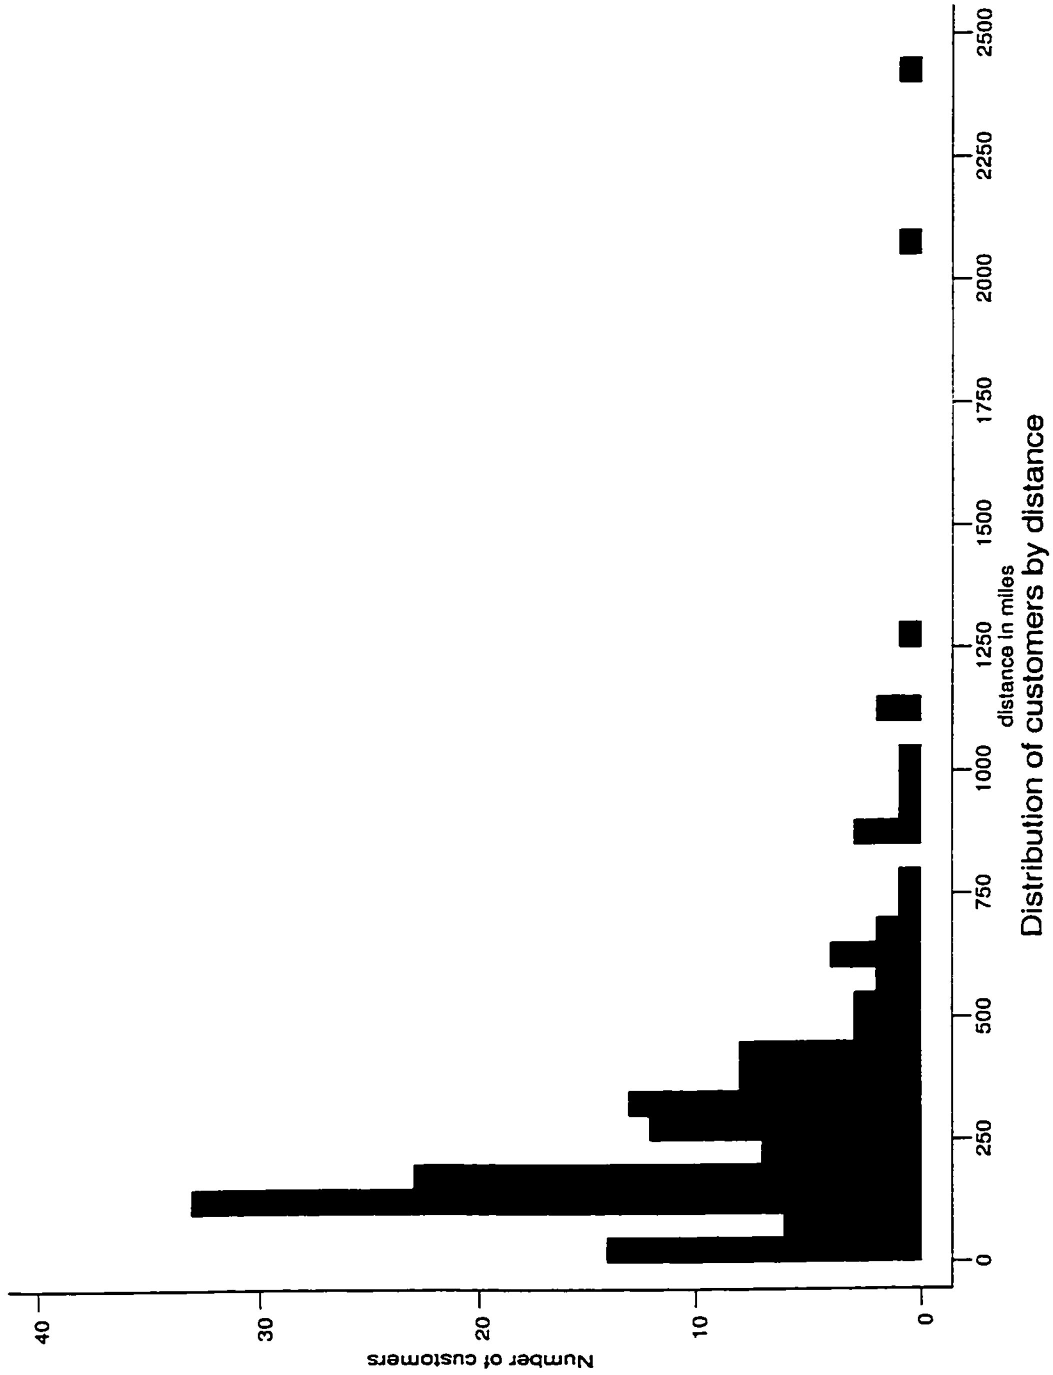

Figure 2.7: Distribution of Customers by Distance, Product P2

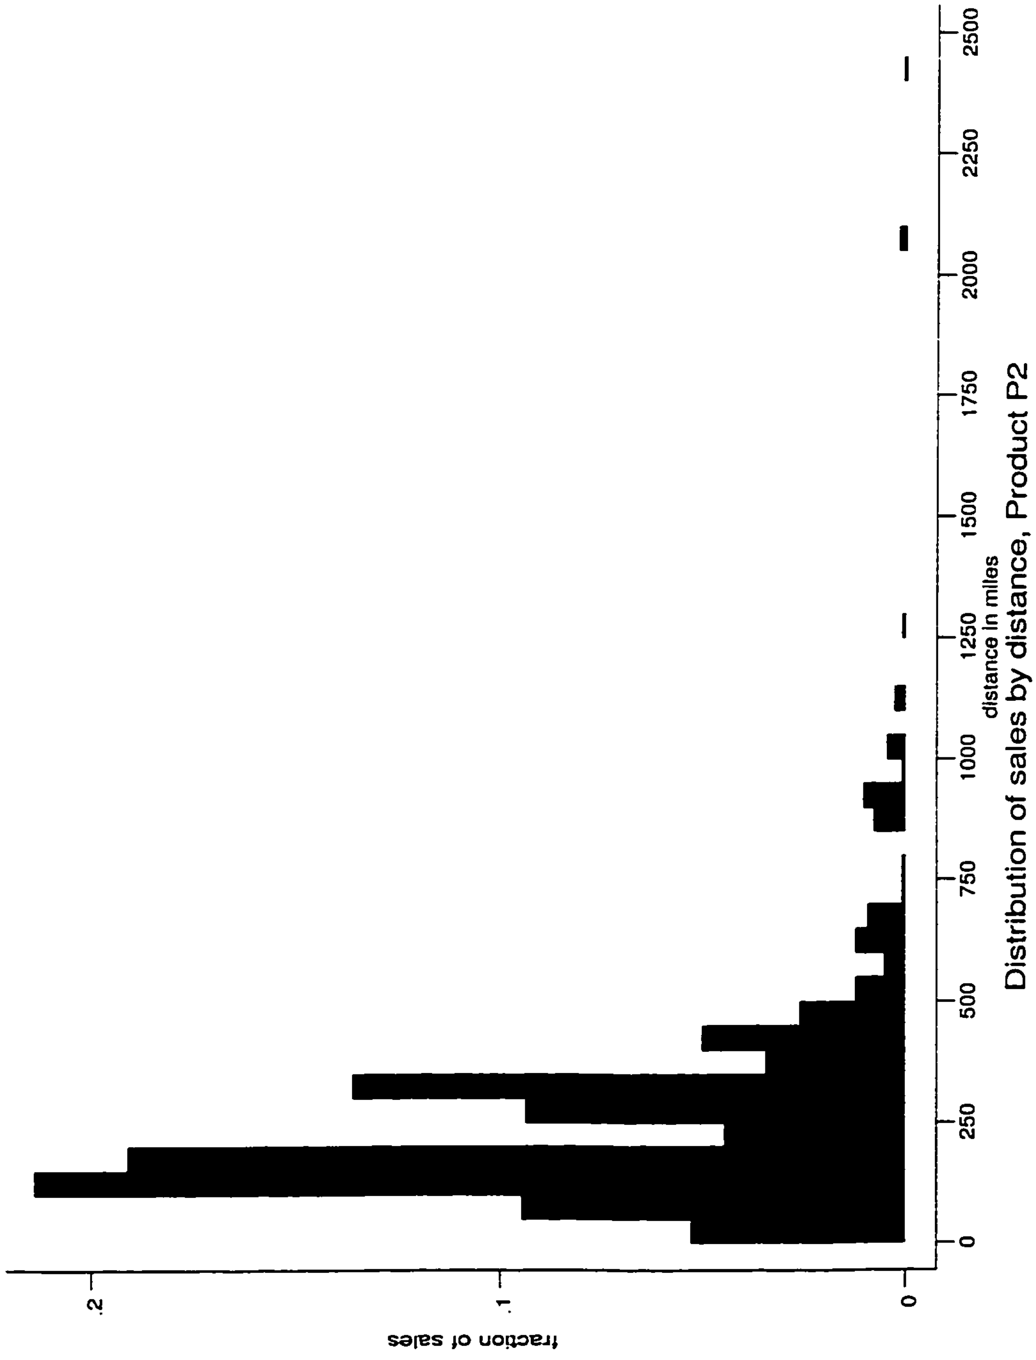

Figure 2.8: Distribution of Sales by Distance, Product P2

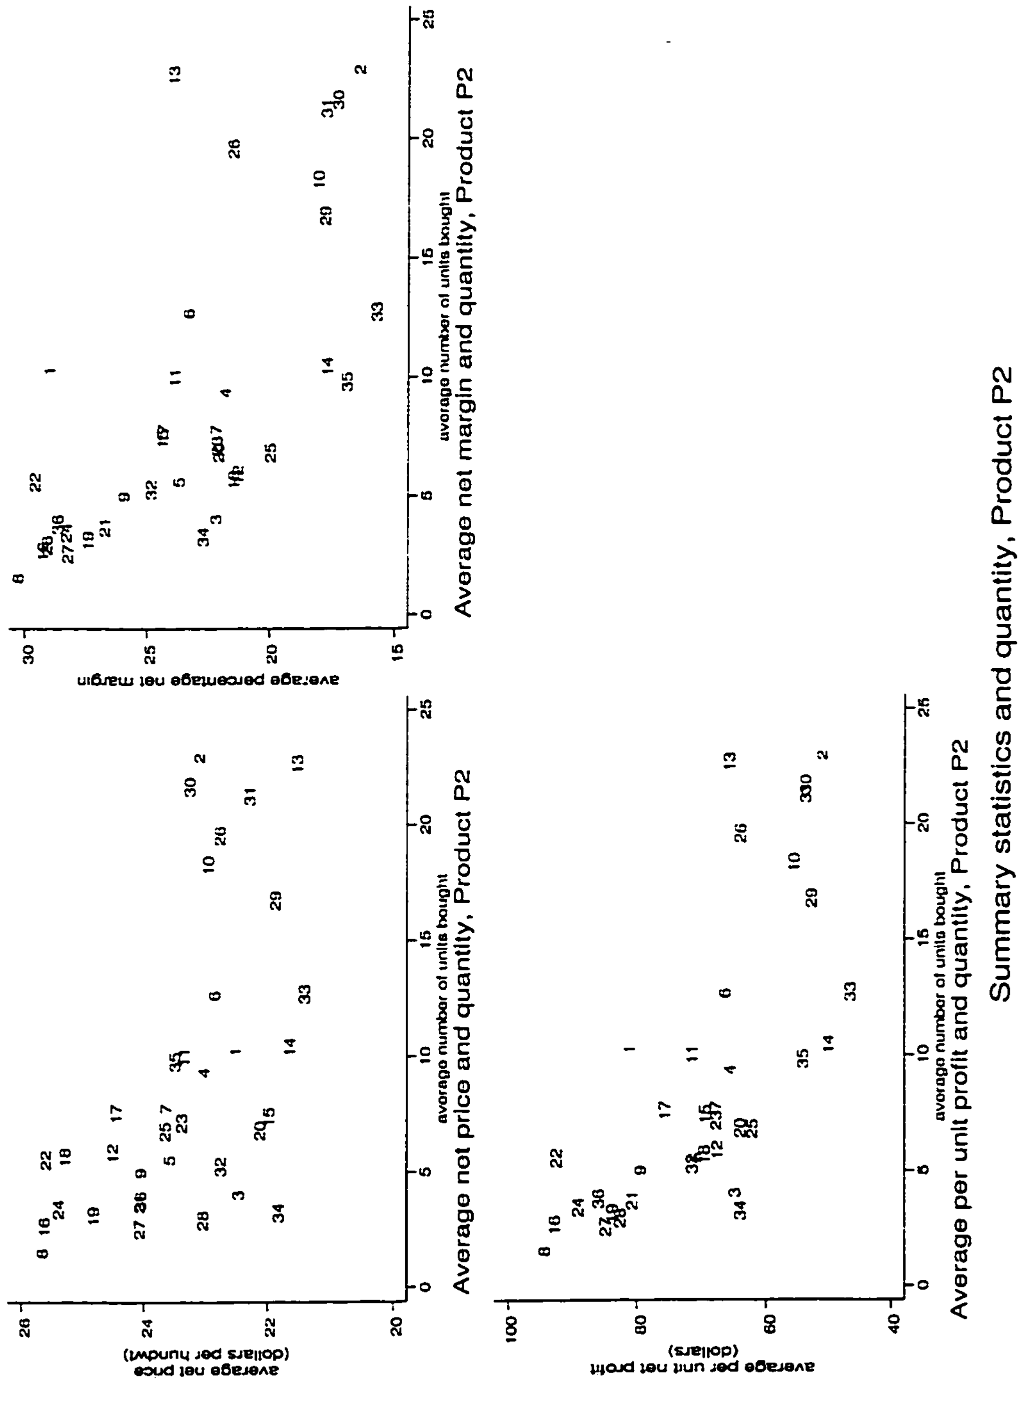

Figure 2.9: Summary Statistics for Frequent Buyers by Average Quantity Purchased

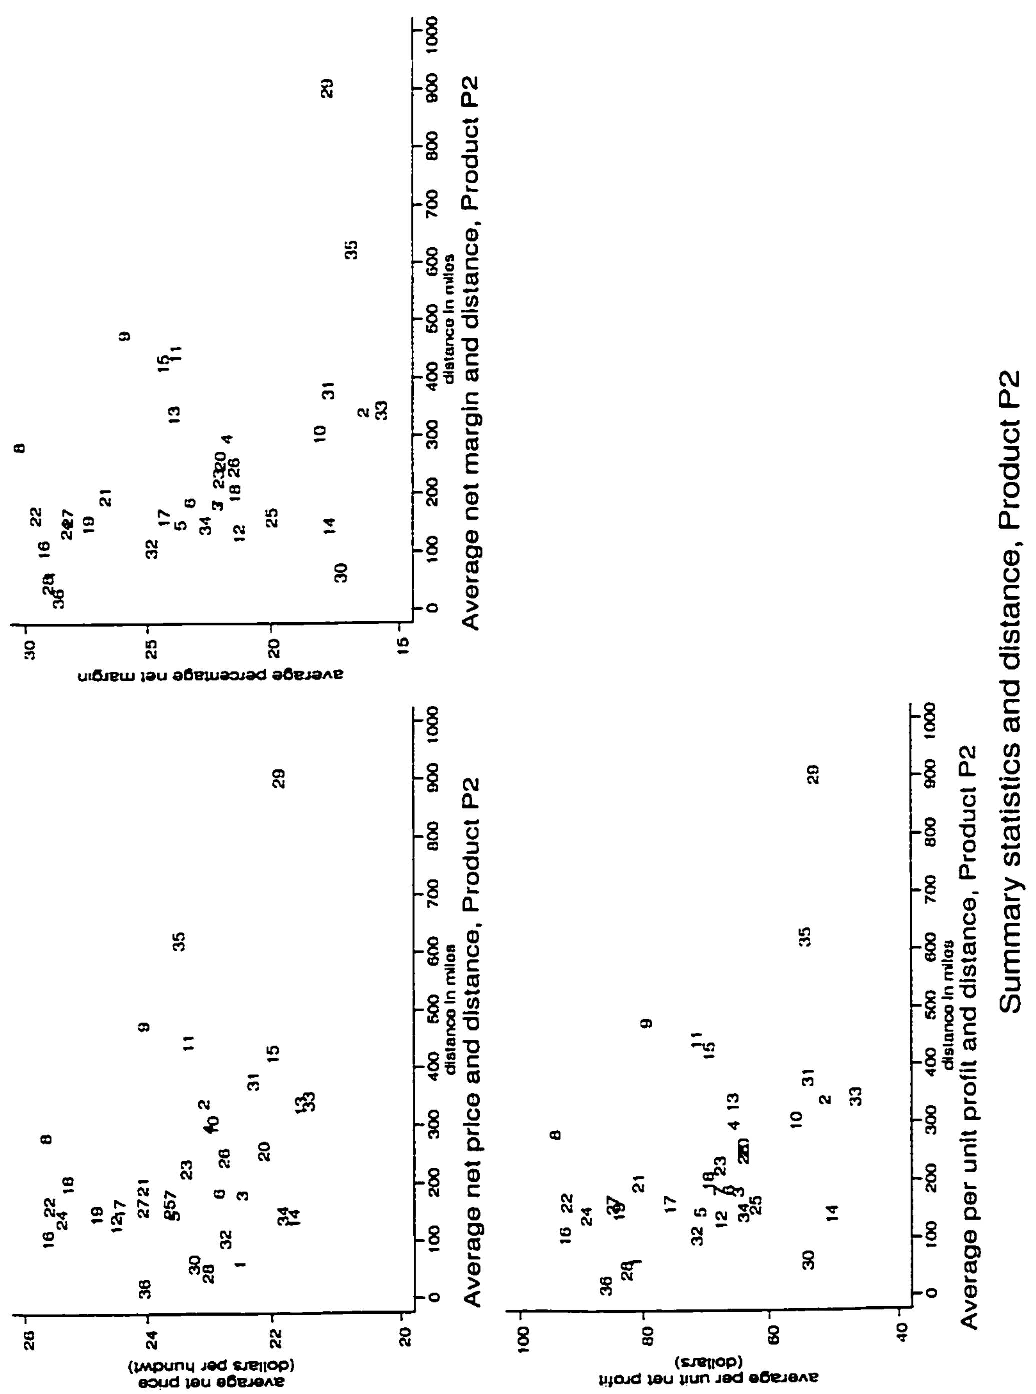

Figure 2.10: Summary Statistics for Frequent Buyers by Distance from Company S

Chapter 3

Motivation, Perception and Morale

3.1 Introduction

Can employers motivate workers through higher pay? Why do firms dislike cutting wages? How do workers' perceptions influence their reactions to the employer's actions? In this chapter I study these and related questions from a perspective which brings together issues of work motivation, incentives and the interaction between an individual and the environment.

Bewley (1999) explores some of these issues in his book on wage rigidity. He finds that firms do not like to cut pay for fear of causing poor morale among employees, and he also finds little support among managers for the "shirking" theory of effort. Pay equity is viewed as important, particularly in the primary sector, where jobs are full-time and long-term, and there is communication between employees. This also makes the pay of new hires in the primary sector downwardly inflexible, as opposed to the secondary sector, where lack of communication often allows firms to sidestep the need for internal pay equity.

Good morale, according to Bewley, can mean a number of things: a willingness to cooperate with company objectives, a sense of common purpose consistent with the firm's goals, enthusiasm for the job, moral behavior and mutual trust. Business people value good morale because it reduces labor turnover, makes recruiting good workers easier, and increases the productivity of the workforce. This increase in productivity, Bewley argues, arises not so much because employees work harder at assigned tasks that are monitored, but because workers do the right thing even when they are not being monitored, make suggestions for improvements, set themselves extra tasks, help each other and share information with each other and with superiors. Not surprisingly, good morale is considered especially important for productivity in jobs where it is difficult to monitor performance, where good performance requires creativity and where workers must deal with customers.

Bewley (1999) outlines a model in which effort depends on morale through a two-stage maximization process. In the first stage, the mind unconsciously chooses the "mood" so that conscious choices maximize an unconscious profit function; the conscious person makes a "decision" about how much effort to exert, and the actual action taken depends on the unconsciously chosen mood. The impact of morale on mood is to vary the subset of possible moods available to the unconscious maximizer; improvements in morale make the set of available moods larger, so that an improvement in mood (weakly) increases effort, and vice versa. A cut in pay reduces morale by worsening the mood, which accounts for firms' reluctance to reduce wages. Morale also manifests itself as an internalization by the workers of the firm's objectives, so that workers are motivated to exert effort in order to raise the firm's profits.

George and Brief (1996) argue that feelings are central to obtaining a richer understanding of work motivation. They suggest that workers strive to become many things, with the self as a worker being just one of many possible selves potentially in need of motivational attention. Feelings help to determine which possible self is focused on motivationally at any point in time and when shifts in this attention take place. They posit that moods, with their origins in person-context interactions, impact both distal motivation (which determines a worker's choice of task) and proximal work motivation (which determines how much effort a worker will expend on a task already selected).

A number of motivational approaches emphasize the beliefs that people have about their performance on a task. Included in this category are theories of goal-setting, self-efficacy and expectancy. Benabou and Tirole (1999) consider a model partly motivated by self-efficacy considerations. In their model, the agent does not have full information about his own productivity, but extracts information from the principal's action. The principal can thus affect the agent's self-confidence, modeled as the agent's estimate of his likelihood of success, through his choice of policy, and in equilibrium higher rewards are bad news, being correlated with lower self-esteem.

Wiley (1997) details the results of several years of motivation surveys, where respondents ranked the top factors that motivate them in their jobs, including good wages, full appreciation for work done, job security, promotion and growth and interesting work. The results indicate that pay is important to people, not just because it enables them to meet material needs but because of the psychological boost in self-esteem that public recognition associated with monetary compensation affords. Rewards such as wages and praise that reflect ability may also lead to greater intrinsic motivation; this suggests that it is not necessarily the reward itself that determines how people respond, but rather the type of feedback implied by the reward. Thus, for instance, extrinsic rewards can increase intrinsic motivation if they are perceived as providing information about competence.

Motivation is a complex, dynamic phenomenon, rather than a static state of mind. As the discussion above would suggest, perception seems to be important: people react in different ways to the same thing depending on the kinds of inferences they make from the observed actions of others. In addition, how people feel about their environment, the behavior of others and their own performance will influence how they react in any given context. In this chapter, therefore, I use a cognitive approach that takes feelings into account, to build a model of motivation in which the worker cares intrinsically about the employer's perception of him, and where his self-image influences his motivation to exert effort.

The chapter is organized as follows. In section 2, I conduct a necessarily brief review of approaches to motivation in the industrial psychology and organizational behavior literature. In section 3, I present a simple example of the feedback from rewards to morale. In section 4, I look at approaches to issues of motivation in the economics literature. Section 5 concludes.

3.2 Theories and Evidence

There is a vast literature on motivation in the fields of industrial psychology and organizational behavior. Wiley (1997) lists three assumptions as guiding contemporary research on work motivation:

- Motivation is inferred from a systematic analysis of how personal, task and environmental characteristics influence behavior and job performance.
- Motivation is not a fixed trait, but refers to a dynamic internal state resulting from

the influence of personal and situational factors.

- Motivation affects behavior, rather than performance; initiatives designed to enhance job performance by increasing employee motivation may not be successful if there is a weak link between job performance and the employee's efforts.

Locke and Latham (1990) observe that motivation theory in psychology has undergone several changes during the past century, with broad, all-encompassing, noncognitive theories being replaced for the most part by narrower, "limited-domain" theories, most of which include key cognitive elements. Early all-encompassing theories include psychoanalysis and behaviorism; the former is not viable as a theory of motivation because, among other reasons, there is no evidence that people possess instincts as postulated by Freud in the sense of inborn drives leading them to pursue preset ends in the absence of learning. Behaviorism, with its emphasis on stimulus-response conditioning, cannot yield an adequate explanation of behaviour either: the so-called reinforcers that make subsequent similar responses to a stimulus more probable do not change behavior unless people want or value them, are aware of the connection between the response and the reinforcers, and believe they can make the required responses.

Another broad theory, formulated by Hull, postulated that all behaviors are based in four primary drives: hunger, thirst, sex, and the avoidance of pain. According to this view, these drives provide the energy for behavior, whereas the associative bonds that develop between drive stimuli and behaviors through the process of drive reduction provide the direction for behavior. When the physiological need deficit that gives rise to the drive state is reduced, activity ceases. Again, this theory fails to explain a whole range of human behaviors: some actions (such as exploration) cannot be traced to physiological need deficits, drive reduction does not always lead to reduced activity, and drive increases such as brain self-stimulation can even be rewarding.

Modern approaches to motivation proceed from the assumption that behavior cannot be explained without reference to consciousness, and seek to account for a more limited set of phenomena rather than attempting to explain all human action (Locke and Latham 1990). These theories are generally cognitive and situationally specific. For instance, there are theories concerned with motivational phenomena in the realm of work (Hackman and Oldham 1980); equity theory (Adams 1963) is concerned with the effects of a single value as it operates in one specific (social) domain. A broader theory is expectancy theory (Vroom 1964), sometimes known as VIE theory (for valence-instrumentality-expectancy) to reflect its major components. Expectancy is defined as an action-outcome estimate, usually in the form of a probability; instrumentality refers to the relationships between an initial outcome like performance and secondary outcomes like pay or promotions, and valence refers to the value that one places on the receipt of these secondary outcomes. The theory suggests that the individual calculates the "expected values" of alternative actions in choosing among them; the measurements used to predict and explain action are based on the individual's perception of each specific situation.

There is some disagreement in the literature as to what constitutes motivation. Some simple notions refer to motivation in terms of someone wanting (and committed to trying) to do something, or "an inner desire to make an effort." Some motivation theorists have argued that intrinsic motivation is that which comes from the task itself, whereas extrinsic motivation comes from sources outside the task.

Deci and Ryan (1985) argue that intrinsic motivation is based on an innate need for competence and self-determination. According to them, it energizes a wide variety of behaviors and psychological processes for which the primary rewards are the experiences of effectance and autonomy. Intrinsic needs differ from primary drives in that they are not based in tissue deficits and they do not operate cyclically; but like drives, they are innate to the human organism and function as an important energizer of behavior. The needs for competence and self-determination motivate an ongoing process of seeking and attempting to conquer optimal challenges. When people are free from the intrusion of drives and emotions, they seek situations that interest them and require the use of their creativity and resourcefulness; the emotions of enjoyment and autonomy represent the rewards for intrinsically motivated behavior.

Deci and Ryan dub their approach "cognitive evaluation theory" and claim that it describes the effects of events that initiate or regulate behavior on motivation and motivationally relevant processes. A central concept in the theory is that of the *perceived locus of causality*. The hypothesis is that intrinsically motivated behavior has an internal perceived locus of causality: the person does it for internal rewards such as interest and mastery, while extrinsically motivated behavior has an external perceived locus of causality, and is engaged in for the purpose of gaining an external reward or to comply with an external constraint. With an external reward or constraint, an instrumentality develops such that the activity becomes a means to an end rather than an end in itself. According to them, the effects of external intervention can be summarized in the following propositions:

- External events relevant to the initiation or regulation of behavior will affect a person's intrinsic motivation to the extent that they influence the perceived locus of causality for that behavior. Events that promote a more external perceived locus of causality will undermine intrinsic motivation.

- External events will affect a person's intrinsic motivation for an optimally challenging activity to the extent that they influence the person's perceived competence, within the context of some self-determination. Events that promote greater perceived competence will enhance intrinsic motivation, and vice versa.

- Events relevant to behavior have three aspects - *informational*, *controlling* and *amotivating*. The first provides effectance-relevant feedback in the context of choice; the second pressures people to act in certain ways, and the third signifies that effectance cannot be attained. It is the relative salience of the three aspects to a person that effects changes in perceived causality and perceived competence and that alters the person's intrinsic motivation.

These conclusions are based on studies that attempt to measure intrinsic motivation by allowing the experimental subject a free choice of activities, and observing how much time he or she chooses to spend on them. Some of the studies also look at the quality of performance or of outcomes as indicators of intrinsic motivation, arguing that because intrinsic motivation has been associated with greater creativity and spontaneity, the presence of those characteristics can signify intrinsic motivation. A number of the studies find that giving individuals extrinsic rewards such as money for performing a task leads to lower intrinsic motivation after the rewards are withdrawn than would have existed had the rewards not been offered at all. Deci and Ryan argue that money, if its controlling rather than its competency aspect is emphasized, lowers intrinsic motivation because it undermines the individual's sense of choice and self-determination.

Locke and Latham (1990) argue that there are a number of serious problems with this conceptualization of intrinsic motivation and its application to the effects of rewards. According to them, virtually no studies using a behavioral criterion have verified Deci and Ryan's interpretation of the reward effect by actually measuring the hypothesized mediating variables and showing that the reward effect works through variables such as feelings of competence and self-determination. Most studies have been interpreted by making inferences from their experimental design and manipulations rather than from an understanding of how the participants actually experience them. Moreover, this conceptualization of intrinsic motivation does not distinguish between liking an activity for its own sake and liking it because it makes one feel competent. Locke and Latham argue that a more logically defensible classification of types of motivation might be as follows:

- *Intrinsic motivation* is involved when the pleasure is derived from the task activity itself, as in the case of an interesting task;
- *achievement motivation* is operative when the pleasure comes from performing well in relation to a standard or goal;
- *extrinsic motivation* is aroused when the pleasure comes from outcomes to which task performance leads.

Since the first two are different phenomena, their causes are likely to be different; Deci and Ryan's measure of intrinsic motivation – that is, time spent on an activity in the absence of pressure or external constraints – seems to be more relevant to the intrinsic motivation, as defined above, than to achievement motivation, which is what feelings of competence would be related to. The latter would be revealed more clearly when performance is undertaken in the presence of standards. In Locke and Latham's view, the so-called free choice behavior measure itself is probably not an adequate measure of intrinsic interest or motivation because time can be spent on an activity for many reasons beside interest. Moreover, they point out that if intrinsic motivation is largely wiped out, as Deci and Ryan claim, by such factors as salient incentives and rewards, competition, imposed goals, standards and deadlines, pressure, anxiety, self-doubt, conflict, instrumental task consequences, feelings of obligation to others, appraisals of performance by others, negative feedback, surveillance, ego involvement, and so on, "...then it is doubtful that it has much application to real life. It seems incongruous that the need for self-determination and competence are considered to be, on the one hand, the wellsprings of all human motivation and at the same time so fragile that their effects can be negated by the most common of life's exigencies."

According to Mitchell (1997), goals have emerged as the major underlying explanatory mechanism for motivational processes; goal and goal attainment discrepancies can be seen as major determinants of the arousal and direction of behavior, and goal importance and goal commitment are important for intensity and performance. The following ideas are key supplements to the theory of goal-driven motivation:

- Humans seek pleasure and avoid pain: this idea is most clearly reflected in the expectancy model discussed above, which suggests that people choose actions based on their perceptions of the consequences of those actions.

- Humans want to have a positive self-view, and to be viewed positively by others. Much of the work on memory and social cognition suggests that people construct, revise and interpret events in ways that reflect positively on themselves.

- Humans are constantly involved in the process of social comparison, and interpretations of fairness and appropriateness can have a powerful impact on actions.

- Humans seek and prefer settings and activities where we have control and mastery. These are integral parts of building competence and self-efficacy, which have been shown to have a powerful effect on motivation and motivational processes.

However, while there is general consensus among motivation researchers that motivation revolves around those psychological processes involved with the arousal, direction, intensity and persistence of voluntary actions that are goal directed, differences in the definition of motivation continue to persist and evolve.

3.3 A Simple Example

There is a large literature in economics on contingent rewards, effort and performance. Economic theory predicts that when pay is contingent on effort, high effort is accompanied by higher wages; any motivation for workers to exert effort, which is costly, comes solely from the fact that more effort means higher wages. When effort is unobservable and there is uncertainty, the efficient contract for a risk-averse agent lies between the two extremes of full insurance (with no incentives) and a pure incentive contract (with no insurance).

In this chapter, I am interested in the feedback from rewards to motivation. One way to conceptualize the problem is in terms of the distinction between "motivation" and "morale," with the first term referring to the agent's attitude toward work, and the second reflecting the effect on the agent's utility of his perception of the principal's actions. According to Bewley (1999), good morale has three components: identification with the organization or internalization of its objectives, good moods, and trust and mutual affinity among members of the organization. Low pay can hurt morale because workers associate pay with self-worth and recognition of their value to the company. I take the agent's "intrinsic" motivation as given, and study the effect of the principal's choice of reward on the agent's morale. The agent cares intrinsically about the principal's perception of his motivation, and in equilibrium the principal has an incentive to reward an agent who has high motivation, or risk lowering the agent's morale and hence the effort exerted.

To make this concrete, consider the following example. There are two players, an agent and a principal. The agent can be of two types, $a_L = 0$ and $a_H = 1$, corresponding to his motivation. The agent chooses effort $e \in [0, 1]$. The principal does not know what type the agent is, but knows that the agent may be a high type with prior probability p_H, and a low type with the complementary probability. She receives a signal $\sigma \in [0, 1]$ about the agent's motivation, with conditional densities denoted by $g_H(\sigma)$ and $g_L(\sigma)$. I assume that the conditional densities satisfy the MLRP.

The timing of the game is as follows. The principal receives a signal about the agent's type and forms a posterior belief denoted by $\theta(a_H|\sigma)$, which indicates the weight she places on the agent being a high type. She sets wage $w \in [0, 1]$ to maximize expected profit. The agent observes the wage and makes an inference about the principal's 'meta-type,' or belief about the agent's type. Let $\hat{\theta}(a_H|w)$ denote the agent's belief about the principal's type. I take the principal-agent relationship as given, so that a type i agent solves

$$\max_e u_i(e,w) = a_i e - \frac{1}{2}e^2 - \frac{1}{2}(a_i - \hat{\theta}(a_H|w))^+ e$$

where $(a_i - \hat{\theta}(a_H|w))^+ = \max\{0, a_i - \hat{\theta}(a_H|w)\}$. The idea is that a motivated agent has positive utility for effort, but that he also cares about whether the principal considers him to be motivated. If he assigns a low probability to the principal placing a high weight on him being motivated, then this lowers his morale, which I model as a scaling down of his utility of effort. An unmotivated agent, on the other hand, has no positive utility for effort, and does not care about the principal's beliefs or intentions. Note that positive motivation could arise for number of reasons: the agent may enjoy the task (so that the motivation might be *intrinsic* in Locke and Latham's sense) or the agent may be trying to achieve certain goals or to perform well in relation to a standard (*achievement* motivation), or the agent may be willing to exert high effort out of altruism towards the principal. In all these cases, the principal benefits from the agent's attitude toward the task; intuitively, therefore, the agent cares about sharing the gains from his effort, and will exert higher effort if he believes the principal to be rewarding him for his intention to do so.

I denote a type i agent's optimal choice of effort by e_i^*. The principal's problem is then

$$\max_w \theta(a_H|\sigma)e_H^*(\hat{\theta}(a_H|w)) + (1 - \theta(a_H|\sigma))e_L^*(\hat{\theta}(a_H|w)) - w$$

The principal must trade off the chance that the agent will simply "take the wage and run" against the likelihood of upsetting a highly motivated agent by offering too low a wage. For a high wage to elicit high effort, it is important that the agent believe that the principal is rewarding him for his intrinsic motivation. I analyze the perfect Bayesian equilibria of this two-stage signalling game. Clearly, there is a continuum of pooling equilibria, indexed by the wage level, in which the principal sets a constant wage regardless of her belief about the agent's motivation level. I am interested in separating equilibria, where the principal successfully signals her perception of the agent's motivation, and the agent responds by conditioning his effort on the wage.

Proposition 1 *In a separating equilibrium:*

1. *the low type agent always exerts zero effort;*

2. *the high type agent exerts higher effort, the higher the wage offered; and*

3. *the principal offers a higher wage, the higher the agent's expected motivation level.*

Proof. (1) is obvious. For (3), note that the high type solves

$$\max_e \; e - \frac{1}{2}e^2 - \frac{1}{2}(1 - \hat{\theta}(a_H|w))^{+}e$$

Unless the last term binds, this yields $e_H^*(\hat{\theta}(a_H|w)) = \frac{1}{2}(1+\hat{\theta}(a_H|w))$; if the last term binds, then we simply have $e_H^* = 1$.

The principal's problem reduces to

$$\max_w \; \theta(a_H|\sigma)\; e_H^*(\hat{\theta}(a_H|w)) - w$$

which yields

$$\theta(a_H|\sigma)\; e_H^{*\prime}(\hat{\theta}(a_H|w^*))\; \hat{\theta}'(a_H|w^*) = 1$$

In a separating equilibrium, the agent can infer the principal's belief from the wage, so that $\hat{\theta}(a_H|w) = w^{-1}(\theta(a_H|\sigma))$ and $\hat{\theta}'(a_H|w) = \frac{1}{w'(\theta(a_H))}$. This implies that the firm's FOC is

$$\theta(a_H|\sigma)e_H^{*\prime}(\hat{\theta}(a_H|w^*))\frac{1}{w^{*\prime}(\theta(a_H|\sigma))} = 1$$

Now $e_H^{*\prime}(\hat{\theta}(a_H|w)) = \frac{1}{2}$, so that the firm's FOC becomes

$$\theta(a_H|\sigma)\frac{1}{2} = w^{*\prime}(\theta(a_H|\sigma))$$

which yields $w^*(\theta(a_H|\sigma)) = \frac{\theta^2(a_H|\sigma)}{4}$, so that the higher the probability the principal places on the agent being a high type, the higher the wage she offers.

For (2), observe that since beliefs are correct in equilibrium, it must be the case that $\hat{\theta}(a_H|w) = \theta(a_H|\sigma) = 2\sqrt{w^*(\theta(a_H|\sigma))}$, so that the high type agent exerts effort $e_H^*(\hat{\theta}(a_H|w)) = \frac{1}{2} + \sqrt{w}$. ◇

Notice that the wages here are not contingent on effort, but are simply offered as a fixed payment. Clearly, this is a good strategy for the principal only if she believes the agent to be a highly motivated type; in this context, the equilibrium may be viewed as one in which the principal trusts the agent, and the agent reacts favorably to that trust. Note that trust arises endogenously in the separating equilibrium but that I take the agent's initial intrinsic motivation as given, so that a low type agent will always exert zero effort. A different approach is that of Casadesus-Masanell (1999), who studies trust in a similar setting. He examines the circumstances within which it is in the agent's best interest to develop intrinsic motivation, which is defined as an "endogenously-determined enticement to work based on preferences, without the intervention of external punishments and rewards." Building on a distinction between the "object self" and the "acting self" that is common in the sociology and psychology literature, he studies situations where it is rational for the agent to develop a preference for adherence to norms and ethical standards, or to develop altruism. He finds that, under plausible conditions, offering a larger fixed payment and lower-powered incentives can lead to larger total surplus.

Rabin (1993) considers the role of fairness in human interaction, on the premise that people care not merely about what other people do but also about *why* they do it. He develops the notion of a "fairness equilibrium," and uses it to argue that one cannot fully capture realistic behavior by invoking pure altruism. In his model, people extract information about the motives of others from the actions they take, so that a player can infer from another player's choice of strategy his reasons for playing that strategy, and this knowledge affects both his payoffs and his choice of strategy. Rabin applies this idea to a labor model similar to the "gift-exchange" framework in Akerlof (1982). The worker chooses whether to exert high or low effort and the firm simultaneously chooses a benefit level for the worker. Each player's dominant strategy is to choose the minimum possible level of his action variable, so that the unique Nash equilibrium is the "nasty" one in which the players feel negatively towards each other. However, there can also be situations in which each player is "kind" to the other, and where each player believes the other player is being kind, so that there exists a fairness equilibrium in which the worker exerts high effort and the firm pays a positive benefit to the worker.

Falk, Fehr and Fischbacher (2000) provide experimental evidence regarding the importance of fairness. According to them, both models that consider only fairness *outcomes* (such as Fehr and Schmidt, 1999) and those that consider mainly fairness *intentions* (such as Rabin, 1993) perform less well than those that consider both outcomes and intentions. They find evidence that the attribution of fairness intentions is important both in the domain of negatively reciprocal behavior (a hostile response to an action that leads to an unfair outcome or is driven by an unfair intention) and in the domain of positively reciprocal behavior (a kind response to a fair outcome or fair intentions). When the experimental design rules out the attribution of fairness intentions, reciprocal responses are substantially weaker. When fairness attributions are ruled out, a significant proportion of experimental subjects display no reciprocal behavior; but when fairness attributions can be made, all subjects display some reciprocity.

The example in this chapter can be interpreted as incorporating concerns about reciprocity. The high type worker intends to exert positive effort, and cares about being rewarded for his motivation. In a separating equilibrium, the employer signals through higher wages (which are costly to her) that she believes the worker will exert high effort, and the worker responds with higher effort than he would exert if he believed the employer did not trust him.

3.4 The Economics of Motivation

Economists have tended to approach motivation indirectly, usually in the context of attempting to explain apparently anomalous relationships between rewards and effort. A number of papers propose that "efficiency wages," above the market-clearing wage, will be paid to workers to induce higher effort by those workers. A well-known example is the paper by Akerlof (1982) on gift exchange, a view of social exchange emphasized in sociology and anthropology. In the basic application of this idea to the labor market, some firms would willingly pay workers in excess of the market-clearing wage, and in return would expect workers to supply more effort than they would if equivalent jobs could be readily obtained (as is the case if wages are just at market-clearing).

Akerlof (1984) outlines four other paradigms within which such wage-setting behavior arises. These can be broadly summarized as follows:

- The dual labor market hypothesis: primary sector jobs have stability, low quit rates, good working conditions, acquisition of skills, a promotion ladder and good pay, while secondary sector jobs have high quit rates, harsh discipline, little chance of promotion, low acquisition of skills, and poor pay. The pay differential between the two sectors can be seen as the difference between wages in excess of market clearing and wages at market clearing.

- The Weberian theory of organization: organizations are modelled as bureaucracies with a hierarchical structure, a well-specified division of labor and impersonal discipline which arises out of the personal loyalty of the employees to the goals of the bureaucracy. According to Akerlof, the most essential feature of the gift exchange model is the importance of employee loyalty to the operation of the firm; thus the evidence supporting the Weberian description of bureaucratic organizations serves as evidence for the importance of this essential feature of the gift exchange model as well.

- The work groups paradigm: modern Weberian theory incorporates the theory of work groups into its description of internal labor markets. The theory builds on the observed discrepancy between formal and actual authority in many different work situations and the complex equilibria which have been shown to exist in which official work rules are partially enforced, coexisting with informal customs and some individual deviance from both the official work rules and informal norms.

- Equity theory: as mentioned earlier, this builds on the work of Adams (1963), in order to address the question whether, *ceteris paribus*, workers with greater pay produce greater output. Many variants of Adams' original experiment have been conducted; not all reproduce his original result that "overpaid" workers will produce more, but the evidence appears strongest for the withdrawal of services by workers who are led to believe that they are underpaid.

Bewley's surveys provide anecdotal evidence that at least some of these factors, most notably the dual labor market and concerns about equity, play a part in determining firms' wage-setting behavior. He finds a consistent reluctance among primary sector employers to hire job-seekers who are deemed to be overqualified for the task at hand. Not only do employers expect to lose such employees once labor market conditions improve, but they also fear that under-utilized employees would be unhappy and may cause trouble with other employees and supervisors. In the secondary sector, overqualification appears to be less of a concern, although this is partly explained by the fact that hiring and training costs are often too low to make it worthwhile to pay high enough wages to induce employees to stay. Moreover, given the nature of the jobs, employers seem to expect lower motivation and higher turnover, even relying on it to keep costs down in some cases. Employers' attitudes to morale also tend to differ between the sectors. Bewley points out that morale is hurt by threats, and that though companies fire some workers, it is thought to be bad business practice to have people work in a negative atmosphere. However, precisely this management style is often used with low-level and low-paid labor doing short term jobs that are easily monitored, most of which are in the secondary sector.

Solow (1979) and Akerlof and Yellen (1990) describe models where pay equity is important in determining the effort exerted by workers. In Solow's model, pay rates have a positive impact on productivity through their impact on morale. In Akerlof and Yellen's model, there are two types of labor, and a lower variance in compensation between the groups leads to more harmonious relations within the firm. If the effort exerted by lower-paid workers depends on the wage gap, the firm chooses wages for higher-paid employees to clear the market, and then maximizes profits in determining the number of lower- paid employees to hire and the wages they are paid. The wage differential between the two types of labor is "compressed" relative to the perfectly competitive equilibrium.

Bewley argues that flattening, relative to marginal products, tends to occur within one job or labor classification, not across jobs, and that it is normally not true that workers in low-level jobs are paid more than their marginal products, while those in high-level jobs are paid less. Such disparities would induce employers to dismiss some low-level employees and hire more high-level ones, an implication not supported by the evidence.

Recently, there have been a number of papers that build on Deci and Ryan's conceptualization of intrinsic motivation and their claim that raising extrinsic incentives can reduce intrinsic motivation. Kreps (1997) writes of the "stylized fact" that the imposition of extrinsic incentives can sometimes be counterproductive, because it may destroy the workers' intrinsic motivation, leading to lower levels of quality- weighted effort and lower net profits for employer. Eichenberger and Frey (1997) write about intrinsic motivation being "crowded out" by extrinsic motivation in some situations. They draw on Deci and Ryan's arguments to suggest that when extrinsic intervention is perceived as "controlling" then it works to shift the agent's focus from an intrinsic desire to exert effort to externally-based reasons for working.

Benabou and Tirole (1999) study situations where an agent who has imperfect knowledge of his own ability extracts information about his ability from the actions of a better-informed principal. They consider a model in which the agent can choose whether or not to exert effort ($e \in \{0, 1\}$), and where the success of the project depends on the agent's ability, which he does not know but about which he receives a random signal. The principal knows the agent's ability; she uses the choice of bonus to signal to the agent her knowledge of the agent's likelihood of success. In equilibrium, the principal offers a high-ability agent a lower contingent bonus than a low-ability agent, since the higher the agent's ability, the larger his expected payoff from exerting effort and the lower the inducement he must be offered in order to choose $e = 1$. They interpret the equilibrium as being one in which the principal signals with a low incentive scheme that she trusts the agent. By conceptualizing intrinsic motivation in terms of the agent's *self-confidence*, or estimate of his own ability, they also argue that rewards, or extrinsic motivation, have a limited impact on the agent's current performance; further, rewards reduce the agent's intrinsic motivation to undertake similar tasks in the future, since they provide bad news in equilibrium.

Notice that in this formulation, the "motivation" that induces the agent to exert effort is actually the extrinsic reward he gets in case of the project's being successful; it corresponds to neither the pleasure the agent derives from the task itself (intrinsic motivation as defined in section 2), nor to the pleasure the agent gets from performing well in relation to a goal or standard (achievement motivation). Nor is "trust" used in the common sense of the word: although the principal's behavior is observationally equivalent to trust, she offers a lower contingent payment not because she trusts the agent to "do the right thing" but because she knows that the agent's expectation of his own ability is high enough to induce him to exert effort without needing to be supplemented by a higher bonus.

Lazear (1999) studies the effect of incentive pay on the employees of a large autoglass company. He finds that when the company switched incentive schemes from hourly wages to a piece rate scheme, average levels of output per worker rose by 44%; of this, half can be attributed to the increase in production by the average worker due to incentive effects, and some of the rest results from an increased ability to hire more productive workers and possibly from a reduction in quits from the highest ability workers. The firm shared the gains in productivity with its workforce, with a given worker receiving about a 10% increase in pay, and the variance of output increased, which Lazear attributes to the increased incentive for more ambitious workers to differentiate themselves in the presence of piece rates. He argues that his evidence implies that the choice of compensation method has important incentive effects, and that it unambiguously refutes the claims of psychologists and sociologists that monetizing incentives may reduce output.

Wiley (1997) reports on evidence from surveys administered in 1946, 1980, 1986 and 1992. Respondents ranked a number of factors that motivated them. She finds that although the relative rankings of the top factors (as reported above) has changed over time to reflect prevailing economic conditions, both high wages and full appreciation for work done are consistently important to people. She argues that in addition to the psychological boost from high rewards, employees may also consider good wages to be solid feedback concerning their work, and that an effective compensation scheme is critical to motivating employees.

3.5 Conclusion

In this chapter I explore some of the ways in which motivation depends on one's interaction with the environment and one's perception of the actions and intentions of others. Feelings, which are often neglected in the standard analysis of work motivation, are important in determining motivational focus and direction, and workers' perceptions of their employer's intentions can affect their morale and their attitude towards work. In the example sketched in section 3, the principal was able to signal to the motivated agent through a high wage that she considered the agent to be a highly motivated type, and that she was rewarding the agent for his intention to exert effort. Evidence indicates that monetary rewards are only one of the things that people care about, and that the psychological effects of monetary rewards are important in their own right. A richer treatment of the inner lives of economic agents is therefore important to understanding a number of issues not explained fully by standard economic theory.

Bibliography

Adams, J.S. (1963), "Wage Inequities, Productivity and Work Quality," *Industrial Relations*, 3, 9-16.

Akerlof, G. (1982), "Labor Contracts as Partial Gift Exchange," *Quarterly Journal of Economics* 97, 543-569.

Akerlof, G. (1984), "Gift Exchange and Efficiency-Wage Theory: Four Views," *American Economic Review*, 74, 79-83.

Akerlof, G. and J.L. Yellen (1998), "Fairness and Unemployment," *American Economic Review, Papers and Proceedings*, 78, 44-49.

Akerlof, G. and J. L. Yellen (1990), "The fair wage-effort hypothesis and unemployment," *Quarterly Journal of Economics*, 105, 255-283.

Anscombe, F.J. and R.J. Aumann (1963), "A definition of subjective probability," *The Annals of Mathematical Statistics* 34, 199-205.

Aumann, R.J. (1976), "Agreeing to Disagree," *Annals of Statistics* 4, 1236-1239.

Aumann, R.J. and A. Brandenburger (1995), "Epistemic

Benabou, R. and J. Tirole (1999), "Self-Confidence and Social Interactions," working paper, Princeton University.

Bernheim, D. (1984), "Rationalizable Strategic Behavior", *Econometrica* 52, 1007-1028.

Bewley, T. F. (1999), *Why Wages Don't Fall During A Recession.* Cambridge, MA: Harvard University Press.

Blume, L., A. Brandenburger and E. Dekel (1991a), "Lexicographic Probabilities and Choice Under Uncertainty," *Econometrica* 59, 61-80.

Blume, L., A. Brandenburger and E. Dekel (1991b), "Lexicographic Probabilities and Equilibrium Refinements," *Econometrica* 59, 81-98.

Börgers, T. (1993), "Pure Strategy Dominance," *Econometrica* 61, 423-430.

Brandenburger, A. (1992), "Lexicographic Probabilities and Iterated Admissibility," in P.S. Dasgupta *et al.*, eds. *The Economic Analysis of Markets and Games.* Cambridge, MA: MIT Press.

Brandenburger, A. and H.J. Keisler (2000), "Epistemic Conditions for Iterated Admissibility," working paper.

Casadesus-Masanell, R. (1999), "Trust in Agency," working paper, Harvard Business School.

Chan, H. (2001a). "Analysis of Variance in Steel Sale Price," manuscript, Yale University.

Chan, H. (2001b). "An Empirical Bargaining Model of Steel Sales," manuscript, Yale University.

Chen, Y.M. (1997). "Paying Customers to Switch," *Journal of Economics and Management Strategy*, 6, 877-897.

Deci, E. and R.M. Ryan (1985), *Intrinsic Motivation and Self-Determination in Human Behavior.* New York: Plenum Press.

Dow, J. and S. Werlang (1994), "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory* 64, 305-324.

Eichberger, J. and D. Kelsey (1996), "Uncertainty Aversion and Preference for Randomisation," *Journal of Economic Theory*, 71, 31-43.

Eichberger, J. and D. Kelsey (1998), "Non-additive Beliefs and Strategic Equilibria," *Games and Economic Behavior*, 30, 183-215.

Eichenberger, R. and B.S. Frey, (1997) "On the relationship between intrinsic and extrinsic work motivation," *International Journal of Industrial Organization*, 15, 427-439.

Ellsberg, D. (1961), "Risk, Ambiguity, and the Savage axioms", *Quarterly Journal of Economics* 75, 643-669.

Epstein, L. (1997), " Preference, Rationalizability and Equilibrium," *Journal of Economic Theory*, 73, 1-29.

Falk, A., E. Fehr and U. Fischbacher (2000), "Testing Theories of Fairness - Intentions Matter," working paper, University of Zurich.

Fehr, E. and K. Schmidt (1999): "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114, 817-868.

Fudenberg, D. and J. Tirole (1998). "Upgrades, Tradeins and Buybacks," *RAND Journal of Economics*, 29 (2), 235-258.

George, J.M. and A.P. Brief (1996), "Motivational Agendas in the Workplace: The Effects of Feelings on Focus of Attention and Work Motivation," *Research in Organizational Behavior*, 18, 75-109.

Gilboa, I. (1987), "Expected Utility Theory with Purely Subjective Non-additive Probabilities," *Journal of Mathematical Economics*, 16, 65-88.

Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141-153.

Hackman, J.R. and G.R. Oldham (1980) *Work Redesign*. Reading, MA: Addison-Wesley.

Hall, G. and J. Rust (2000). "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries," *Carnegie- Rochester Conference Series on Public Policy.*

Hall, G. and J. Rust (2002). "Middle Men versus Market Makers: A Theory of Competitive Exchange," working paper, Yale University.

Klibanoff, P. (1996), "Uncertainty, Decision and Normal Form Games," working paper, Northwestern University.

Kreps, D. "Intrinsic motivation and extrinsic incentives," *American Economic Review*, 87, 359-364.

Lazear, E. (1999), "Performance Pay and Productivity," working paper, Stanford University.

Lo, K.C. (1996), "Equilibrium in Beliefs under Uncertainty," *Journal of Economic Theory* 71, 443-484.

Locke, E.A. and G.P. Latham (1990), *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice Hall.

Marinacci, M. (2000), "Ambiguous Games," *Games and Economic Behavior*, 31, 191-219 .

Mitchell, T.R. (1997), "Matching Motivational Strategies With Organizational Contexts," *Research in Organizational Behavior*, 19, 57-149.

Monderer, D. and D. Samet (1989), "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170-190.

Morris, S. (1996), "The Logic of Belief and Belief Change: A Decision Theoretic Approach," *Journal of Economic Theory* 69, 1-23.

Morris, S. (1997), "Alternative Definitions of Knowledge," in M.O.L. Bacharach *et al.* eds., *Epistemic Logic and the Theory of Games and Decisions*. Dordrecht, The Netherlands: Kluwer Academic.

Mukerji, S. (1997), "Understanding the Nonadditive Probability Decision Model," *Economic Theory* 9, 23-46.

Mukerji, S. and H.S. Shin (1997), "Equilibrium Departures from Common Knowledge in Games with Non-Additive Expected Utility," mimeo, University of Southampton.

Pearce, D. (1984), "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* 52, 1029-1050.

Rabin, M. (1993), "Incorporating Fairness into Game Theory and Economics," *American Economic Review*, 83, 1281-1302.

Ryan, M.J. (2000), "What do Uncertainty-Averse Decision Makers Believe?," mimeo, Australian National University.

Samuelson, L. (1992), "Dominated Strategies and Common Knowledge," *Journal of Economic Theory* 73, 1-29.

Savage, L. (1954), *The Foundations of Statistics*, John Wiley and Sons.

Schmeidler, D. (1989), "Subjective Probability and Expected Utility Without Additivity," *Econometrica* 57, 571-587.

Stahl, D.O. (1995), "Lexicographic Rationalizability and Iterated Admissibility," *Economics Letters* 47, 155-159.

Tan, T. and S. Werlang (1988), "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory* 45, 370-391.

Taylor, C. (1998). "Supplier Surfing: Price Discrimination in Markets with Repeat Purchases," working paper, University of Texas.

Vroom, V.H. (1964), *Work and motivation*. New York: John Wiley.

Wiley, C. (1997), "What Motivates Employees According to Over 40 years of Motivation Surveys," *International Journal of Manpower*, 18, 263-280.