

A SATISFICING MODEL OF CONSUMER BEHAVIOR

by

MARK JOSEPH RYAN

A DISSERTATION

Presented to the Department of Economics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

September 2011

DISSERTATION APPROVAL PAGE

Student: Mark Joseph Ryan

Title: A Satisficing Model of Consumer Behavior

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

Dr. Van Kolpin	Chairperson
Dr. Christopher J. Ellis	Member
Dr. Jeremy Piger	Member
Dr. Renee Irvin	Outside Member

and

Kimberly Andrews Espy	Vice President for Research & Innovation/Dean of the Graduate School
-----------------------	---

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded September 2011

© 2011 Mark Joseph Ryan

DISSERTATION ABSTRACT

Mark Joseph Ryan

Doctor of Philosophy

Department of Economics

September 2011

Title: A Satisficing Model of Consumer Behavior

Approved: _____
Dr. Van Kolpin

I develop a model in which a representative consumer selects an affordable consumption bundle, not as a single choice, but as the end result of a series of smaller, incremental purchase decisions. If the array of such incremental choices facing the consumer is sufficiently complex relative to the consumer's computational abilities, then the consumer may choose to employ a simplifying heuristic or rule-of-thumb to guide her behavior. I demonstrate the existence of a simple and well-defined example of such a strategy, based upon a satisficing decision rule. I further show that in the strategic setting defined by the interaction between consumers and firms that compete in prices, this satisficing strategy can form part of a Nash equilibrium, despite being *ex ante* only boundedly rational.

The use of this satisficing demand strategy fundamentally alters the nature of price competition between firms (relative to the standard Bertrand model), changing the shape of the firm best response functions. The use of a satisficing strategy alters the incentives of firms, and these altered firm incentives lead to pricing behavior which has the effect of rationalizing the satisficing consumption strategy, so that a truly novel class

of Nash equilibria in price-competing markets can be shown to exist under certain conditions.

We explore the nature of this new class of equilibria, and find that equilibrium prices may be higher than those which would be obtained in the standard Bertrand case. In general, demand curves for each distinct good will have a kinked shape, similar to those found in 1939 papers by both Sweezy and Hall & Hitch. The Nash equilibrium profile will involve the kink in each demand curve coinciding with the equilibrium price for the corresponding good. The equilibrium price vector will therefore be robust to “small” fluctuations in cost (since marginal revenue is discontinuous at the equilibrium price), and under certain conditions, we find that prices may be upwardly flexible but downwardly rigid. We make an argument that the main results of the paper generalize from a representative agent setting to one with a population of heterogeneous consumers.

CURRICULUM VITAE

NAME OF AUTHOR: Mark Joseph Ryan

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene
University of Tennessee, Knoxville

DEGREES AWARDED:

Doctor of Philosophy in Economics, 2011, University of Oregon
Master of Science, Economics, 2006, University of Oregon
Bachelor of Arts, Economics and Psychology, 2003, University of Tennessee

AREAS OF SPECIAL INTEREST:

Game theory
Microeconomic theory
Business Cycles

PROFESSIONAL EXPERIENCE:

Graduate Teaching Fellow, University of Oregon, 2004-2010

ACKNOWLEDGMENTS

I wish to thank the members of my doctoral committee for their indispensable help and guidance throughout this project. In particular, I am most grateful to Van Kolpin, my advisor. Your help, your advice, and (especially) your critiques were always insightful, practical, and elegant. I always enjoyed the discussions we had in your office about this dissertation, especially those discussions in which you and I did not fully agree on some particular point or topic. Those were always challenging and fun, and I will miss that part of this process very much.

I am also especially indebted to Chris Ellis. Your advice was extremely valuable in some early discussions which helped to clarify and focus the problem that I set out to investigate, as was your remarkable insight and economic intuition, which were always of great help to me throughout this project.

Finally, I wish to thank my wonderful wife, Leslie. You supported me throughout this process in every way possible. I cannot express how touched I am by and how grateful I am for all of the sacrifices you made along the way, and by all of the good things that you willingly and lovingly deferred while you and I were both working toward this goal. This would not have been possible without you.

For my wife, Leslie

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	1
II. THOUGHTS ON SATISFICING AND BOUNDEDLY RATIONAL EQUILIBRIUM	9
Microfoundations	9
Arbitrarily Inaccurate Information in Strategic Settings	17
Model Equilibrium Defined	20
The “Garbage In, Garbage Out Problem”	25
The Consumer Budget Problem as a Complicated Problem.....	28
III. THE CONSUMER SATISFICING DEMAND STRATEGY	47
The Optimizing Consumption Strategy as a Sequential Process	47
The Sequential Satisficing Process Outlined	52
Thoughts on the \mathfrak{R} Function.....	53
The \mathfrak{R} Hypothesis.....	55
Examples of the \mathfrak{R} Function.....	58
An Intuitive Sketch of the Satisficing Process Under \mathfrak{R}_5	79
Additional Comments on the \mathfrak{R} Function	82
The Path-Generating Process II	83
Beginning to Outline the Geometry of the Satisficing Demand Process.....	90
The Worth Constraints	97
Three Illustrative Cases.....	97

Chapter	Page
Worth Constraints in Action	99
IV. THE COMPLETE MARKET, WITH SATISFICING CONSUMERS	118
Deriving the Demand Curve	123
Individual Loci and Demand Curve.....	126
The Satisficing Market Equilibrium	148
The Kinked Demand Curves and Nash Equilibrium	154
Aggregation.....	159
Differential Rationality	175
V. IDENTIFICATION.....	180
Features Unique to Satisficing Demand Curves	182
Thoughts on Monetary Policy.....	190
Business Cycles	191
Expansionary/Inflationary Phase	195
Contractionary Phase	197
VI. CONCLUDING REMARKS.....	200
APPENDIX: SAMPLE MATHEMATICA CODE	203
REFERENCES CITED.....	229

LIST OF FIGURES

Figure	Page
1. Another Kind of Duality	30
2. Sequential Optimizing Consumption.....	50
3. One Generic Iteration of the Sequential Satisficing Process	54
4. Vector Fields and Consumption Trajectories	87
5. Identity Field (\mathbb{R}^2).....	88
6. Inverse Identity Field (\mathbb{R}^2).....	89
7. Budget Constraint and the Consumption Path.....	91
8. Varying the Vector Associated to the Origin.....	93
9. The Worth Constraint	100
10. The Worth Constraint and the Π^* Field	103
11. Complete System of Worth Constraints	103
12. Final consumption bundle.....	105
13. Worth Constraints Shaping Differing Π Fields	106
14. The system of worth constraints divides the consumption space	109
15. The same graph as above, with the Actual Consumption Path (ACP)	109
16. The Identity Field, with the inside worth constraint shown.....	112
17. The Identity Field, with the inside and outside worth constraints shown	113
18. The complete Π^* field.....	114
19. Same graph as above, with the inside worth constraints and peripheral streamlines removed	114

Figure	Page
20. The budget constraint may be either more restrictive than the worth constraints.....	116
21. ...exactly as restrictive as the worth constraints.....	116
22. ...or less restrictive than the worth constraints.....	117
23. Deriving the Worth Constraint Locus in Price-Quantity Space.....	122
24. From Consumption Space To Demand Curve Loci.....	124
25. X^1 as a function of price (P_x).....	127
26. X^2 as a function of price (P_x).....	129
27. Both Loci Superimposed.....	129
28. Satisficing Demand Curve for good X (Axes Transposed)	130
29. The standard presentation of the satisficing demand curve	130
30. Firm X profit function.....	136
31. Firm X Best Response	136
32. The demand curve shown for good X, when P_y is equal to 20	137
33. The Bertrand equilibrium price vector, (9.39, 9.39)	141
34. Satisficing Demand Curve for X when $P_y = 9.39$	142
35. Satisficing Firm X profit function for $P_y = 9.39$	142
36. The equilibrium price vector, (30, 30)	143
37. Once again, observe that the kink price coincides with the equilibrium price of 30	144
38. Individual Demand Curves for Heterogeneous Consumers.....	163
39. Non-equilibrium Market Demand Curve with Heterogeneous Consumers.....	163

Figure	Page
40. Primary and Secondary Kinks	165
41. Consumer 1 Demand Curve for X	172
42. Consumer 1 Demand Curve for Y	172
43. Consumer 2 Demand Curve for X	173
44. Consumer 2 Demand Curve for Y	173
45. The aggregated market demand curves for good X	174
46. The aggregated market demand curves for good Y	174
47. CPI vs. PPI volatility.....	183
48. Initial Equilibrium at price vector (30, 30)	187
49. By decreasing μ from 10 to 9, the consumer can maximize utility	188

CHAPTER I

INTRODUCTION

There is a well-known economist joke:

A police officer is walking his beat one night, when he comes across an economist who is stooping beneath a lamp post, searching around on the ground.

“What are you doing?” asks the policeman.

The economist replies, “I lost my keys over there in that dark alley, and I’m trying to find them.”

The policeman, puzzled, then asks, “Well, if you lost your keys in the alley, then what are you doing looking all the way over here by the lamp post?”

The economist answers, “Because it’s easier to look over here!”

When modeling the behavior of actual human beings, economists have perhaps occasionally tended to make the mistake of the unfortunate economist from our joke. When modeling behavior, it is relatively easy to make the agents in a model behave as if they were very smart, and it is relatively easy to make agents behave as if they were very stupid. But it can be quite difficult to model agents as being only *reasonably* intelligent. Economists, therefore, have for the most part tended to build models for explaining behavior that are based on assumptions consistent with perfect rationality. This is despite the fact that we often have reason to believe that, in practice, realistic agents can be considerably less than fully rational. There is no doubt that this approach has proven tremendously fruitful. We are certainly able to make an impressive range of accurate predictions concerning general patterns of behavior, based on the economic models which assume full rationality of agents. Nonetheless, it would be unwise to pretend that we do not face some sort of a trade-off between the tractability of our models and the validity

of their results, even if we can find general agreement on the notion that, so far, the costs have been worth the benefits.

In the pages that follow, we shall attempt to build a model describing the behavior not of fully rational, but of merely “reasonably” rational agents. As we develop this model, it will no doubt be easy to convince the reader that the concepts and structures necessary to rigorously model these *boundedly rational* consumers are quite a good deal more cumbersome than those needed to model agents who are endowed with perfect rationality. We hope to make the case that the benefits of *this* modeling approach are also worth the costs, however. We believe that what we lose in easy tractability, we more than make up for with what is arguably a closer correspondence to actual human behavior. Along the way, we hope to illuminate a few novel insights into the behavior of markets and their participants.

Economists interested in studying complicated real-world situations will often find it necessary to translate these real situations into simpler models, in the hopes of boiling the problem down to only represent its most important features. These models may be concisely described as well-defined mathematical problems, and will often feature solution algorithms which guarantee finding the set of optimal solutions to the problem of interest.

However, economists always face the danger that, in transforming a real-world problem into its corresponding model, salient features of the real problem might be lost in translation. If this were to happen, economists who lose sight of the fact that the model and the real problem are not *literally* the same thing, may be lulled into a false sense of security by the fact that the model’s solution technique is often both perfect and unique.

Within the model, the logic justifying the use of a particular solution technique may often be totally unassailable. However, this does not necessarily mean that the solution algorithm in the model corresponds to the solution technique used by real agents, in their attempt to solve the real problem.

In other words, while no one denies the power of modeling, whenever we find *a* set of keys under the lamp post, we must be careful to make sure that we reserve an appropriate level of skepticism towards the assumption that we have found the *right* set of keys. Models built on the assumption of fully rational solutions to simple problems risk missing important consequences and implications, if the real world involves imperfectly rational solutions to complicated problems; even small deviations from the truly optimal solution can have potentially large consequences. To continue with the lost-key metaphor, the objective of this paper is to identify a second set of keys in the alley, and to begin to think about the process by which we will determine which set is the one we were looking for in the first place.

It is in this spirit that we now wish to revisit a well-known problem which cannot currently be said to be the subject of any active debate among economists, the consumer budget problem (CBP). The problem itself is familiar to first-year economics undergrads: a consumer has a certain amount of financial resources at her disposal, she has certain preferences over the types of goods that she is able to purchase with those resources, and she faces a vector of prices which she must pay in order to acquire those goods. Her challenge is to allocate her scarce financial resources in such a way as to derive the greatest possible satisfaction from her consumption choices. The standard solution to this problem is also well-known. The available wealth and prevailing prices determine the set

of bundles of goods that the consumer is able to afford, and the consumer chooses a bundle from within that affordable set which she most prefers.

It is clear that this *model* for describing the real-world struggle for consumers to get the most satisfaction from the wealth they have at their disposal does, in fact, abstract away from certain features of the problem which are present in the real problem. In particular, the obvious fact that we generally do not observe actual consumers directly choosing *final bundles* of goods, but rather, we tend to see them making a series of sequential decisions concerning what to purchase, and the aggregated outcomes of those decisions *end up* as the final bundle of goods that the consumer has chosen. The standard solution explicitly ignores this sequential and cumulative feature of the actual observed behavior on the part of consumers. This is not to say that the standard model is not extremely useful, nor to say that the standard solution is not deserving of its status as the fundamental construct of consumer theory and a workhorse of microeconomics in general. But, there may be reason to believe that the standard solution to the CBP is at least imperfect. Perhaps the correspondence between the real world and the standard solution breaks down in the absence of a fully rational consumer, or more specifically, in a setting where a boundedly rational consumer must make sequential purchase decisions over time.

In order to investigate this issue, and to therefore hopefully take our small step away from the lamp post and towards the alley, we will begin investigating the implications of *satisficing* behavior on the part of consumers. Herbert Simon (1956), who coined the term “satisficing,” was the first to argue in its favor as a plausible explanation of the behavior of reasonably rational individuals who face complicated problems.

Hogarth (1975) found evidence consistent with Simon's arguments when he tested subjects' abilities to estimate probability distributions of low and high complexity. But within the specific context of consumer behavior, this is a topic that does not yet appear to have been extensively investigated by economists. In a clever empirical study, Kapteyn, et al. (1979) used a series of surveys to attempt to measure consumer preferences directly, and then compared actual consumption behavior with the predictions of a utility-maximizing consumption plan (assuming the surveys measured the underlying preferences) and a satisficing consumption plan. They found that consumer behavior (with respect to the purchase of consumer durable goods) was not consistent with the utility maximization hypothesis, and, under their model, were able to reject the null hypothesis of utility maximization in favor of the satisficing alternative hypothesis.

Also, despite some limited work on making some operational assumptions for empirical models of satisficing behavior, there does not seem to have thus far been any attempt to develop a theoretical framework for deriving satisficing-based consumer demand, and investigating its implications and consequences. Doing precisely this is the primary aim of this paper.

The investigation is potentially an important one. As Abreu and Rubinstein (1988) found, incorporating implementation costs of complicated strategies into the payoff function can lead to large and discontinuous changes in the set of Nash equilibria. Their paper did not consider satisficing behavior explicitly. Nonetheless, as Simon initially argued, allowing for implementation costs of optimizing strategies (and recognizing the relative simplicity of satisficing strategies) can form a main line of supporting argument

for satisficing as a plausible description of actual behavior. As a result, we have at least some reason to believe that our investigation of satisficing consumer behavior will have important implications for the set of Nash equilibria in market games.

In the present paper, we will develop a model of consumer behavior which is based, not on the explicit optimization of a given objective function, but instead upon the serial application of a single satisficing rule of thumb. In the broadest sense, of course, this does not imply that the agents in our model do not maximize *something*. In principle, there will be *some* set of preferences, defined perhaps over strategies themselves instead of simply over consumption goods, and the behavior of our agents will be consistent with the maximization of that unspecified utility function. This is similar in flavor to optimal search models, for example, in which optimization of the objective function which includes search costs gives the result that, small deviations from the true maximum amount of utility obtainable through the direct consumption of goods can still be considered optimal, if those losses can be offset by the avoidance of the search costs necessary to identify the most favorable prices for each good. In this way, optimal search models give the impression that, relative to the direct utility from consumption only, there is a minimum acceptable level of satisfaction, and the consumer will select the first option she encounters which meets or exceeds this minimum level of direct utility.

We will emphasize the differences between our satisficing model and existing optimal search models in greater detail in a later section of this paper (Chapter V: Identification), but for now we will stress one main difference. Here, we will allow consumers to employ a true satisficing strategy, which unlike the consumers in optimal search models, is not predicated upon any underlying optimization process for a pre-

specified objective function. As a result, the behavior of our satisficing consumers will not be subgame perfect, and the possibility exists that our consumers might employ a strategy which is truly suboptimal for certain price vectors (this is never possible under the assumptions of optimal search models), while at the same time still being optimal for others. However, in our model, we will of course still specify a payoff function for our consumers. In the simplest version of this model, this payoff function will be equal to the direct utility function from consumption of goods. In this way, we still retain the ability to identify strategy profiles in which the consumers are playing a best response strategy to the strategies of the other players in the market game. As a result, even though our satisficing consumers are not explicitly acting in such a way as to guarantee the maximization of any particular objective function, it will still be possible to identify situations in which their behavior forms part of a true and complete Nash equilibrium in the market game.

Recognizing that consumer behavior based upon a true satisficing rule-of-thumb, rather than upon an optimal search process, implies that we are endowing the consumers in our model with a strictly lower degree of rationality than we are accustomed to granting them. As we explore the consequences of this in our model, we ought to have two questions in the back of our minds:

- 1) To what extent do the qualitative predictions and results of existing standard models still come through, even with the lower level of consumer rationality? (that is, to what extent are the standard results robust to the relaxation of full rationality assumptions?)

- 2) To what extent does the new model offer novel insight into the nature of consumer behavior, or of market outcomes? (to what extent does this model make plausible predictions not made by any single existing model?)

Specifically with respect to question 2, we ought to be very clear that our expectation as we develop this model is that any additional insights that we may find are certainly to be taken as complements to the profound insights developed through standard models based on true utility maximization. We certainly do not wish to suggest that these previous results are invalid, but perhaps they are incomplete. Our posture throughout the exploration of the implications of satisficing consumer behavior is basically summarized by the observation that, a well-defined model of consumer demand based upon true satisficing behavior is something that we have never really considered before. And now that it has occurred to us, it is something we, as a discipline, ought to begin to consider.

CHAPTER II

THOUGHTS ON SATISFICING AND BOUNDEDLY RATIONAL EQUILIBRIUM

Microfoundations

Any analysis which is based on the notion of satisficing behavior, rather than fully optimizing behavior, is going to require some defense against the charge that it rests on ad hoc assumptions, rather than on solid microeconomic foundations. Since satisficing, by definition, involves the selection of an alternative which is merely “good enough,” rather than the (a) best available option, the natural question to ask is, “If agents have an available option which is better than the one they have selected as ‘good enough,’ why wouldn’t they eventually wish to exercise that better option instead?” That is, why is it reasonable to model agents as measuring the satisfaction derived from any potential choice against some (possibly arbitrary) satisficing criterion, rather than against the best possible outcome (or equivalently, against every possible alternative outcome)?

There are at least three lines of argument which can be articulated to address these concerns. The first is that, under certain conditions, satisficing behavior can indeed converge to the same outcome which would have prevailed had the agent employed optimizing behavior instead. Particularly, it is possible to find examples in which individuals do not optimize, per se, at the outset of a particular task, but satisfice instead, and yet over time the process of satisficing converges, through some form of learning, to the outcome which would have prevailed had that individual optimized in the first place.

For example, Day (1967) gives an example of a firm which does not know its own profit function, and so satisfices with respect to increases in its own profit through alterations in its output level. He shows that, over time, a firm following this satisficing process (in the face of a constant demand curve) will eventually learn to supply the correct profit-maximizing quantity in its market, so that (eventually) there is no difference in the behavior of an optimizing firm and a satisficing one. Crain, et al. (1984) performed an empirical test of Day's model, and found support for Day's conclusions, using a sample of 107 Fortune 500 companies over a 10 year span.

Secondly, there is a modeling consideration. There is arguably an attitude within economics which holds that, since we believe agents are generally rational, goal-oriented utility maximizers, then any microeconomic description of their behavior should be consistent with our standard concepts of optimizing behavior. Any model based upon full rationality is therefore well-supported on solid and acceptable microfoundations, while any model featuring less-than-full rationality is thought of as not compatible with standard results of microeconomics, and must therefore rest upon shaky foundations or ad hoc assumptions.

I wish to take issue with this view. From my perspective, if our aim is to understand the way that markets actually function, then it is important to recognize that our main goal should be to describe behavior, rather than to prescribe behavior. More specifically, it is important to conduct our analyses of markets in a way that is consistent with that goal. In situations where we have good reason to suspect that consumers do not act in a manner consistent with full optimization, then we ought to develop models in which the underlying micro decision process is itself distinct from standard optimization.

That is, we should reserve the term “ad hoc” only for cases in which it is clear that we, as modelers, are employing some sort of clear trick or gimmick which was never intended to have any compelling correspondence to any aspect of the real world. If, on the other hand, we imagine a model which is intended to bear a correspondence to reality in its underlying mechanics and structure, then I argue that it is not appropriate to refer to these structures as resting on ad hoc assumptions. Of course, such a model will still be either true or false, correct or incorrect, etc. But it will not be, in the most complete sense of the term, *ad hoc*, to the extent that it attempts to build a description of the actual decision process.

As an example, consider two similar models, and the extent to which the assumptions of each can be thought of as truly ad hoc.

The Calvo pricing mechanism is perhaps the most well-known example of an ad hoc assumption; there is a population of firms in an economy, and at any point in time, each firm has a certain probability of being allowed to reset its price to that which would be optimal under then-current conditions (otherwise, each firm must retain its previously-set price). This implies that, on average, a certain proportion of the total mass of firms is able to adjust their price at any point in time, while the remainder must keep their price fixed. Despite the usefulness of this construct in matching actual macro level data, clearly this is an ad hoc assumption. We do not seriously believe in the price fairy.

And yet, in another setting, an idea which is very similar to the Calvo mechanism might not be so obviously ad hoc. Consider, for example, the mechanism by which technological progress occurs, as modeled by Eaton and Kortum (1999). In their setup, firms produce output by processing intermediate goods into final goods, but have the

option to invest in research and development which can potentially increase the efficiency with which they consumer intermediate goods. Research and development yields improved production technology in the following way. Each worker who is occupied in R and D will produce a fixed number of “ideas” in a given period of time. For the population of ideas which is produced by the collection of R and D workers at any point in time, there is a fixed probability that any one idea will be a “good” idea, which actually leads to improved productive efficiency. Otherwise, the ideas are bad, and are discarded. So, if a firm gets lucky, and experiences a favorable draw from the population of possible ideas, then they get to use that idea to improve their technology. But otherwise, they are forced to continue using their previously existing technology, in the same way that “unlucky” firms in the Calvo model must continue to charge their existing price.

Despite the fact that one model features a “price fairy” and the other features a “good idea fairy,” I contend that not both of these modeling structures are deserving of the label ad hoc. In fact, as one who has from time to time “produced” extremely long successions of ideas without finding any of them to be favored by the “good idea fairy,” I can say that this structure described by Eaton and Kortum, in addition to being a quite elegant and creative means of formalizing a model of technological progress, also bears a striking resemblance to reality. It is this close correspondence with reality that is, in my view, the reason that this production model does not deserve the label of “ad hoc-ery.”

Similarly, for the model which we develop in the subsequent pages, we very much do have in mind that the mathematical processes and structures are more or less isomorphic to the actual decision process employed by reasonably rational and realistic

consumers. As such, while we should remain quite open to the possibility that this model is incorrect, I do not feel as though it is appropriate to criticize the model on the grounds that descriptions of satisficing behavior are inherently ad hoc. They are merely different from standard models of consumer optimization.

Finally, and most importantly, the distinction between static and strategic settings, and the resulting implications for the concept of microfoundations, needs to be emphasized. In static problems, where agents face an unchanging objective function, the argument that we ought to expect those agents to make the best choice they possibly can (i.e., optimize) can be quite compelling. But it is important that we remain mindful of the important distinctions between static problems and strategic ones.

In a strategic setting, it is of fundamental importance that, instead of each individual actor having a single, unchanging objective function, we have interaction between the *actions* of one player and the *objective* function of others. That is, in a game, the action choice of any single player potentially impacts, not just the payoff of that player, and not just the payoff of other players, but also the *relationship* between other player's actions and those player's own payoffs.

It is precisely this extra layer of interaction that makes game theory a distinct field, and it is precisely this extra layer of interaction that requires game theory to employ its own unique array of tools for analyzing problems of this type. Indeed, it is also precisely this extra layer of interaction that makes normal static analysis, and normal static notions of what constitutes solid microfoundations, inadequate. Since choices of one player effect more than just the payoff of that player, there is the possibility of strategic interaction amongst players which increases/reinforces the losses incurred by

satisficing rather than optimizing. But it is equally *possible* that this type of interaction can have the opposite effect; the strategic consequences of satisficing might attenuate or overturn completely the losses we would expect to find based upon a static concept of rationality. In particular, since ex ante irrational behavior on the part of one player has the potential to change the incentives, and therefore the behavior, of other players, it is possible that strategies which may seem less than fully rational in a static setting may end up being rationalized by the behavior of other players in a game-theoretic setting. There are, potentially, important implications and outcomes which static notions of optimization are not capable of identifying.

It is for this reason that we wish to call special attention and emphasis to a fact that is widely agreed upon in economics in general: in strategic settings, the proper concept of microeconomic foundations of behavior is that of Nash equilibrium or its common refinements. That is, when we wish to assess the rationality of a particular strategy or of a particular strategy, it is meaningless to attempt to do so outside of the *context* of the game itself. Strategic behavior is rational *relative* to what else is happening in the game itself, and the Nash equilibrium concept allows us to identify important strategy profiles in which all players are acting rationally, *relative* to each other.

This point requires elaboration, of course. In a market game, it is certainly possible to demonstrate the existence of a class of Nash equilibrium which is nonetheless *clearly* an ad hoc description of consumer behavior. In this setting, a consumption strategy is simply any mapping from the space of all possible price vectors onto the space of all affordable consumption bundles. Strictly speaking, any consumers complete strategy set is the set of all possible such mappings. By construction, we can demonstrate

that a particular strategy forms part of a Nash equilibrium, even though there is no possible way we would ever accept that strategy as a reasonable explanation of actual consumer behavior.

Define $D(\mathbf{P})$ as the usual, rational and optimizing demand function, which is consistent with the utility maximization algorithm. Let \mathbf{P}_1 be any arbitrary price vector for which each firm's price is greater than or equal to its own marginal cost of production. Whatever the particulars of the vector \mathbf{P}_1 , there exists a Nash equilibrium in which firms jointly price at the levels specified by \mathbf{P}_1 . This is true because, technically, the consumer always has the option to play the following strategy:

[If the actual price vector in the market is \mathbf{P}_1 , then buy the consumption bundle given by $D(\mathbf{P}_1)$.
If the actual price vector in the market is anything other than \mathbf{P}_1 , then purchase nothing (the consumption bundle given by the zero vector).

This strategy profile is clearly a Nash equilibrium (firms all earn less profit at any other price, and $D(\mathbf{P}_1)$ is utility-maximizing under the price vector \mathbf{P}_1 , by the definition of D , so the consumer experiences less utility from any other affordable bundle. Nonetheless, this is clearly a nonsensical description of consumer behavior.

This “silly” example of market Nash equilibrium underscores the appeal of the restricting our attention in consumer theory to fully rational demand functions, as well as reinforces the justification for the standard presentations of Bertrand competition, in which consumers are assumed, either explicitly or implicitly, to formulate their demand

functions on the basis of the fully optimizing process D. In making this assumption, analysis of equilibria in the Bertrand game is restricted to cases of subgame perfection only, and we obtain the well-known result that Bertrand markets have a unique *subgame perfect* equilibrium. Again, the usefulness of subgame perfection here is, in part, that it allows us to dispense with the nonsensical Nash equilibria of the type described above.

However, subgame perfection is perhaps a double-edged sword here, at least to a certain extent. From our perspective, it is important to emphasize that, subgame perfection is appealing *because* it refines away certain nonsensical equilibria which are not in any way compelling descriptions of actual behavior. Equally worthy of emphasis is the fact that this is *not* the same thing as saying that those undesirable equilibria are nonsensical *because they are not subgame perfect*. To be clear, we claim that whether or not a particular equilibrium is a reasonable description of behavior is something that ought to be judged on a case-by-case basis, and not with the blanket test of whether or not it is consistent with subgame perfection.

We intend to demonstrate that there does exist a class of Nash equilibrium in market games which is not subgame perfect, but which nonetheless offers a reasonable and compelling description of consumer behavior. We will therefore argue that, by restricting attention to cases of subgame perfection, as is currently customary, we are improperly censoring our ability to describe the full set of reasonable outcomes in markets where firms compete in prices.

As this pertains to our discussion of what is and what is not properly considered an “ad hoc” description of behavior, we argue that (examples of nonsensical equilibria notwithstanding) any Nash equilibrium which is based upon reasonable descriptions of

consumer behavior, in which the strategies and payoff functions for all players are continuous and have a natural correspondence to what we believe to be realistic behavior, and which are in some sense stable, and have a meaningful basin of attraction within the relevant strategy spaces, cannot be considered ad hoc in the usual sense of the term.

Arbitrarily Inaccurate Information in Strategic Settings

To further elaborate on this idea of stable Nash equilibrium as the most appropriate notion of microeconomic foundations, we now want to take our discussion of bounded rationality one step further. We now wish to demonstrate that, by imposing a small amount of additional structure to our analysis of strategic behavior, it is possible to show that even strikingly irrational behavior may potentially persist indefinitely as part of an equilibrium. In doing so, we hope to further strengthen our case that consumer behavior which is based upon the type of less-than-fully rational rule of thumb. In particular, we would like to begin to imagine what sort of conclusions might reasonably be drawn about equilibrium in a strategic setting with a less restrictive set of assumptions considering player information and rationality. Rather than assuming that all players know each of the relevant facets of the game they are playing, we will allow that each player has some internal concept of that game, which we will call a model. Each player's model may, in principle, be either correct or incorrect with respect to any of the defining characteristics of the game itself, and misconceptions about the nature of the game on the part of one player may affect the incentives and behavior of any of the other players.

To keep this as general as possible, we will assume that each player's model has all the features that we normally ascribe to an entire game (set of players, strategy sets, payoff functions, etc.).

In strategic settings, the actions of one player affect the payoff functions of other players. It is, in principle, plausible to suggest that a player's internalized conception of a game is more directly responsible for determining their behavior than the objectively true game is. Players will make decisions based upon what they *believe* to be the nature of the game they play, even if those beliefs are mistaken.

Since the behavior of one affects the payoffs of others, the possibility that any individual player may make decisions which are based upon incorrect beliefs may have real consequences in strategic settings. Here, we begin a rough sketch of what it might look like if we were to attempt to develop a conceptual framework for finding an equilibrium which allowed players to vary widely in their beliefs about the structure of the game itself, as well as to choose a strategy from amongst their available strategy set. We will refer to the equilibrium concept we are groping for here as "model equilibrium."

Model equilibrium is basically the idea of Nash equilibrium, with relaxation of the assumption that players actually understand the structure of the game. In model equilibrium, all players perceive that they are doing the best they can, *given* their own (possibly incorrect) perception of the parameters of the game.

The most general formulation of this idea probably looks something like the following:

Γ : the true underlying game, the parameters of which are possibly not fully known to the players. It is fundamentally defined by the following standard functions and sets:

I : the set of players

S : the collection of individual strategy sets available to each of the players in I

σ : the strategy profile describing the behavior of all players

P : the payoff function, which maps from the set of strategy profiles to the set of payoff vectors

Σ : extensive form game structure (redundant, since this is actually implied by S , but we will list it explicitly here)

N : the set of Nash equilibria (also implied by S and P . Furthermore, we can redefine N in terms of any common Nash equilibrium refinements, as desired)

If we relax the assumption that the players possess full and common knowledge of the above underlying game parameters, then we can begin to describe a new layer of information which is likely to be important in determining how actual play of the game will unfold.

We may hypothesize a second layer of structure on top of the fundamental or “true” game structure. If the true underlying game Γ describes the manner in which strategy profiles map to payoff vectors, then the each individual’s game model Γ_i describes what player i believes would be

In addition, the manner in which the play of the true game Γ progresses will be determined in part by the set of the individual consumers' internalizations, or models, of that true game Γ .

Γ_i : the individual model for each player i . For each $i \in I$, Γ_i is that player's internal conception of the game. Γ_i will in turn be composed of I_i , S_i , σ_i , P_i , and Σ_i in a manner similar to Γ above.

Γ_i need not necessarily be equal to Γ . If Γ_i and Γ are in fact equal, then we say that player i *knows* the game. Otherwise, player i simply *models* the game.

Γ is the game that the players are actually playing, while Γ_i is the game each player i *thinks* they are playing (Or, at a minimum, that each player i *acts* as if they are playing Γ_i . It is entirely possible that the player i knows that Γ_i is not the true game, or possibly even that player i does know the full game Γ , but chooses to use Γ_i as a simplifying heuristic to guide her play of Γ anyway).

H_i^t : the history of events, as of time " t ", which have been observed by player i .

$\Lambda_i(H_i^t, \Gamma_i)$: the Γ_i updating function ($\Lambda_i: (H_i^t, \Gamma_i^t) \rightarrow \Gamma_i^{t+1}$).

Model Equilibrium Defined

There are at least two different senses in which we can describe a strategic situation such as this one as being in equilibrium.

The first, which we may provisionally refer to as "weak model equilibrium" occurs when each player's beliefs about the nature of the game induce a sequence of

actual play which never falsifies any player's model Γ_i . That is, given each player's updating function Λ_i , we have

$$\text{for each player } i \in \mathbf{I}, \text{ we have } \Lambda_i(H_i^t, \Gamma_i) = \Gamma_i$$

In other words, if no realized history of play is ever sufficient to make any player alter her understanding of the game (as defined by her model Γ_i), then the vector of all models is stable, even if the actual action profile in each period that the game is played varies from period to period.

We may therefore define what we will provisionally call a “strong model equilibrium.” A game is in strong model equilibrium if and only if for each player $i \in \mathbf{I}$, player i is playing a strategy that is a best response to σ_i in Γ_i , and $\Lambda_i(H_i^t, \Gamma_i) = \Gamma_i$.

That is, if all models are stable, and each player, based upon her own model, has no perceived incentive to alter her own behavior, then the game is in strong model equilibrium.

The extent to which the Γ_i 's converge to Γ (if at all) depends in large measure upon the nature of the Λ_i functions. The more sensitive the Λ_i functions are to falsifying events in the sequence of observed events H_i^t , the more quickly and more accurately Γ_i will approach Γ .

As one example of how this might work in practice, consider the game of chess. Recalling Zermelo's theorem, for any 2 player, zero-sum, finite game of perfect information, it must be the case that exactly one of the following is true: Either the first player has a strategy which can guarantee that she wins 100% of the time, OR the second player has such a strategy, OR each player has a strategy which can force no worse than a tie.

It is clear that the game of chess satisfies the hypotheses of the theorem, so it should be clear that so long as black wins some of the time and white wins some of the time, then it must be the case that at least one of the players is not playing a best-response strategy. Consequently, Nash equilibrium does not describe the actual play of the game. Instead, as we outlined in the introduction, it is likely that each player forms a model of the game and uses that model as a heuristic to guide her choice of moves as actual play unfolds. In this context, one way to describe what such a model might look like is to imagine some sort of preference ranking of potential board states. In the limit (as the players become “fully rational”) this preference ranking is consistent with backward induction, but short of common knowledge of full rationality, the ranking is, in general, distinct from the backward induction ranking.

There is a widely held maxim amongst chess players that the only way to get better at chess is to regularly play against players who are better than you. The logically equivalent statement being that, if one only plays players who are equal to or worse than oneself, then one will not improve at the game. This second version of the statement is perhaps the more direct statement of the concept of model equilibrium. If we imagine two novice players who play only against each other, then we would not expect either player to improve quickly, because neither player is capable of conspicuously demonstrating the flaws in the other’s understanding of the game. Both players may go on for quite a long time playing the game at a low level of ability, despite the fact that there are, in the abstract, many different opportunities for either player to improve their game. That is, both novice players have several profitable deviations from their current chess strategy/model, but they are not able to easily identify those areas for improvement.

So, despite the fact that such play cannot be supported by the concept of Nash equilibrium, it can, in principle be supported by the concept of model equilibrium. That is, play which is sub-optimal, relative to the object idea of full-rationality, may nonetheless persist indefinitely if it only takes place within a narrow enough context so that its shortcomings are not exposed by the superior player.

That is, two novice players, i and j , may find themselves in a situation where the sequences H_i and H_j never realize sufficient number of events that would case Λ_i and Λ_j to update either player's game model Γ_i or Γ_j . In other words, the two sub-optimal models (and the sequences of play which may result from them) may in principle persist indefinitely under certain circumstances. That is, the models and their resulting histories of play may be part of some relevant equilibrium, even if that equilibrium concept can be shown clearly *not* to be that of Nash equilibrium.

In addition, if we consider an alternative situation, where two players of differing skill levels play repeatedly, we would expect the more skilled player to win the majority of the time. Through the series of losses, the less-skilled player will slowly be made to become aware of the flaws in her game model Γ_i^t . Over time, Λ_i will eventually furnish her with a new and better game model Γ_i^{t+a} (where $a > 0$). The regardless of the actual sequence of Γ_i^t 's, we can assume that there will be some sort of selection process by which "worse" Γ_i^t 's are eventually selected out, and "better" Γ_i^t 's are eventually selected in, since "worse" game models are overturned by experience more easily and more frequently than "better" game models. The end result is that, over time, player i tends to have a "better" game model than she started with, assuming that she consistently plays against more skilled players than herself. Of course, in this context, saying that a given

player develops a “better” game model over time is precisely what it means to say that the player herself improves at the game of chess. The larger point, though, is that poor models, or poor understandings of the actual nature of a game, may persist indefinitely. It is not until something specific happens which *shows* the player that (or how) her understanding of the game may be improved that these incorrect beliefs begin to be improved.

This is an attempt to develop a conceptual framework for thinking about strategic situations in which relevant actors are not necessarily perfectly aware of the parameters of the strategic interaction among players. It is probable that there are wide variations in the histories which may be supported in equilibrium, and the specific sequence of outcomes depends to a large extent upon the nature of the updating function Λ_i . To the extent that these functions are “scientific,” which we can loosely take to mean very good at identifying observed events which are incompatible with the existing model, and both “valid” and “creative” by which we mean good at forming alternative models Γ which are consistent with all previously observed events, and in the simplest and most plausible way, then we would expect the model of the game to converge to the true game more quickly. But, to the extent that the Λ functions are irrational, animistic, superstitious, or biased towards emphasizing only those observations which reinforce the model, and towards understating or dismissing those observations which are incompatible with the model, then we would expect a much slower (or nonexistent) convergence of the model to the true game. In the same way that that some cultures have embraced science, and have advanced their knowledge and technology considerable enormously throughout history, other cultures have reached an incorrect, but stable, understanding of the world

which does not seem to undergo any revision or correction, despite being wrong. (examples abound) In any case, the preceding theory has far too many degrees of freedom to be considered predictive, but hopefully nonetheless serves as a convincing conceptual illustration of how irrational behavior may still persist indefinitely as part of an equilibrium, and therefore be worthy of consideration and study, even if we know that it is not truly optimal behavior.

The following sections of the larger paper which deal with the derivation and the implications of the satisficing demand strategy are not precisely a concrete example of this equilibrium concept, although we will return to this concept more explicitly as we discuss a hypothetical business cycle in the final section. Instead, the following sections are inspired by the spirit of this model equilibrium concept. What follows is an attempt to explore the implications for economic theory when we relax the usual assumption that agents are fully rational, or that they have full knowledge of the game in which they are operating

The “Garbage in, Garbage out Problem”

In principle, this concept of model equilibrium can be used to justify or support literally any possible outcome in the face of any underlying game. For any arbitrary action profile α , there is a strategy profile σ^α in which each player’s strategy is to play the specific action from the given action profile in all circumstances. For any strategy profile so constructed, there is *some* game Γ^{σ^α} for which that profile constitutes a Nash equilibrium (for example, construct a game for which each player’s payoff under the strategy profile is 1, but each player’s payoff would be 0 if they were to employ an

different strategy). Now if each player's model Γ_i is the game Γ^{σ^α} which has been so constructed, then, regardless of the true underlying game Γ , the arbitrary action profile α is supported by a model equilibrium. Once again, the generality of this line of argument implies both that, if we start with a given outcome, we can support that outcome as a model equilibrium of literally any game, and if we instead start with a given game, then we can support literally outcome as a model equilibrium of that game.

In light of this, we might ask ourselves whether this concept is at all useful. A model that predicts literally anything can in some sense be said to predict nothing. However, we should recognize that we can raise the same objection about any analysis based on the concept of preferences or utility maximization. Literally any type of behavior is consistent with *some* set of preferences, so we may use the concept of rational agents pursuing their own preference maximization to also support an absurdly large set of possible outcomes to nearly any problem. That is, the idea of using preferences as a fundamental basis for the motivation or behavior of agents also, in principle, suffers from the “garbage in, garbage out problem.”

Insert whatever wildly implausible or otherwise seemingly irrational action you like here, and that behavior can be justified by the simple assumption that the relevant agent's preferences were such that her behavior was, in fact, completely rational and utility-maximizing.

Of course, the GIGO problem does not derail our attempts to use preferences or utility functions to motivate very useful models or to draw insightful conclusions about agent behavior. This is true because, while we recognize on some level that strange preferences can certainly lead to strange outcomes, we tend to focus on what we believe

to be realistic description of preferences, and we build our models on what we believe to be plausible and well-behaved utility functions. Here, we will attempt to do the same thing with our Γ_i 's and our Λ_i 's. If we can discover cases of particular complicated games where player modeling likely plays a non-trivial role in determining the outcome of the game, then we can show cases where the concept of model equilibrium adds to our understanding of strategic behavior by relaxing the assumptions of full knowledge or full rationality.

As in formal logic, where an argument may be *valid* if the reasoning and inferences are each themselves valid and correct, we may here have in infinitude of model-profile-strategy-profile combinations which are model equilibria, so long as the combined effect of all Γ_i 's and the individual strategies are such that no player would choose to change either their strategy or their model. An argument's validity does not depend on its premises actually being *true*, however, just as the existence of a model equilibrium does not depend upon the collection of models (Γ_i 's) being plausible and realistic internal representations of the actual game Γ . That is to say, that a valid argument may yield a conclusion which is in fact false, if one or more of the premises of that argument are themselves false. Model equilibrium may provide an unrealistic description of agent behavior and of the strategic outcome if one or more of the models in that equilibrium are themselves unrealistic or implausible.

In formal logic, an argument is *sound* if it is valid, while at the same time it is the case that all of the argument's premises are true.

Returning to the issue of the relative usefulness of defining a concept such as model equilibrium, the *mathematical* question is whether or not these model equilibria

exist. The *scientific* question is which, if any, of these equilibria form a compelling basis for describing the behavior of actual agents. The first question is loosely analogous to the logical question of validity, while the second is loosely analogous to the logical question of soundness. Insofar as economics is a science, we are mainly concerned with questions of the second type. From this perspective, we do not have to accept or reject the merit of the model equilibrium concept as an all-or-nothing proposition. We may reject some instances of model equilibrium while retaining others, using our judgment and empirical evidence to suggest which should be kept and which should be discarded.

If we have reason to believe that the only plausible assumptions concerning the players' rationality and degree of knowledge of the underlying game is that of complete rationality and complete knowledge of the game, then the only *plausible* model equilibria describing that game are the set of appropriately defined Nash equilibria. However, if there are reasonable deviations from full rationality/full knowledge, then there may be additional compelling equilibria for a specific game beyond the set of refined Nash equilibria.

It will later be our contention that, despite not being consistent with full rationality (*ex ante*), the satisficing consumer demand strategy can be thought of as a *plausible* means of describing the boundedly rational consumption behavior of actual human beings.

The Consumer Budget Problem as a Complicated Problem

We now return to our explicit examination of the consumption choice as an example of a complex decision, which by nature of its size and sequential nature, is

arguably sufficiently difficult a problem to tackle at once that it strains the cognitive resources of a realistic consumer. In order to explore the question concerning possible consequences of the sequential aspects of consumer decision-making, and of bounded consumer rationality, we shall attempt to re-cast the familiar consumer budget problem in a different context. Rather than specifying an n -dimensional consumption space, and defining a feasible consumption region as all n -tuples within that space, on or below the budget hyperplane corresponding to the consumer's available wealth and the prevailing price vector¹, we will reformat this problem into an equivalent, but richer structure which incorporates a sequential feature to the process of selecting a consumption bundle. We may begin by attempting to describe the consumer decision-making process as a tree.

In principle, we can easily describe the activities of the consumer using a decision tree. The task of a consumer then becomes to make a series of small decisions which *aggregate* to the best possible terminal node (consumption bundle), rather than to select directly the best bundle from among the set of affordable bundles. The differences in the two formulations (a budget tree vs. a feasible consumption set) are subtle, but before we begin to explore them, we ought to consider the similarities (fig. 1).

In principle, if we wish to describe the set of possible decisions that a consumer could possibly make concerning how to spend her income, there is no fundamental difference between using the construct of a decision tree, and using the construct of a budget set. One may be more suitable (or less suitable) for illustrating certain aspects of the decision-making process than another, but whatever conclusions one draws about

¹ This is consistent with the well-known formulation and solution algorithm of the Consumer Budget Problem, as seen in standard economics texts at every level, as in, for example, Krugman and Wells at the principles level, Varian at the intermediate micro level, and Mas-Colell, et. al. at the graduate level.

consumer behavior from using one construct ought to be equally as valid under the other as well.

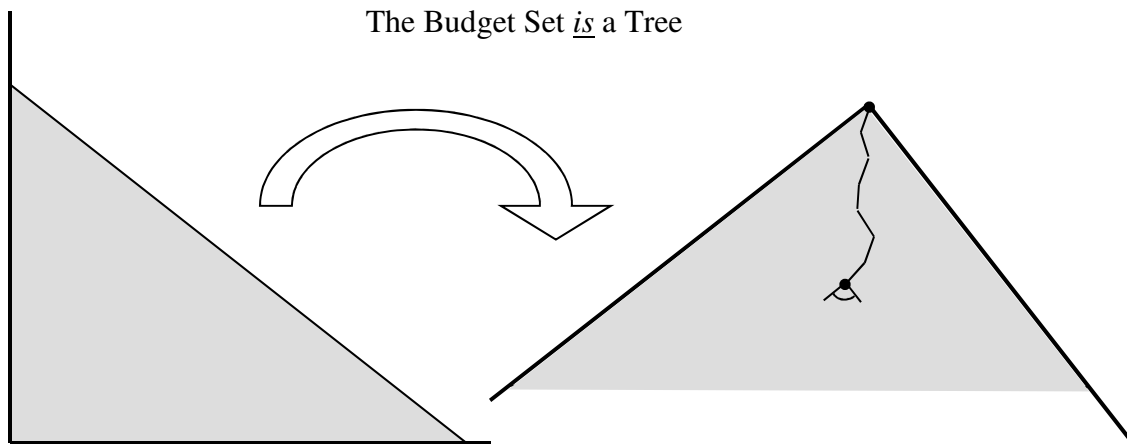


Figure 1. Another Kind of Duality: We can think of the budget set itself as a sort of tree, with its root at the origin, and the budget hyperplane as the set of all terminal nodes. Each point in the interior of the budget set is an intermediate node, and the set of action choices available at each intermediate node can be thought of as the set of all (arbitrarily small) vectors in \mathbb{R}^L_+ . The set of all possible histories in the tree is the set of all non-decreasing paths from the origin to the budget line.

Drawing on this equivalence, we can begin to imagine what a boundedly rational consumer's behavior might look like. It may be hard to imagine that a consumer would not be able to identify a most-preferred object from among a set of objects (in this case, that a consumer might be unable to identify at least one most-preferred point on the budget hyperplane), even in the case where set of candidate objects (the budget set) is very large. However, it is, perhaps, more plausible to think that a boundedly rational consumer might have trouble solving a large decision tree by backward induction. In order to do this, a consumer must know, at the time that she makes any individual purchase, whether or not that particular purchase is a part of the optimal final

consumption bundle. In other words, if the consumer is truly behaving in a manner consistent with fully rational utility maximization, then by the time she decides whether or not to make her very first individual purchase in a particular period – no matter how small that purchase may be – then she must have already solved her entire budget problem, and have already identified the precise complete combination of items and quantities that she will purchase during that period.

This implies that a truly rational consumer who solves her budget tree by backward induction (which is equivalent to employing the standard utility maximization algorithm, and finding the point of tangency of the budget hyperplane and the highest attainable indifference curve) would know on the first day of the month everything that she planned on buying throughout the course of that month. For example, such a consumer would never choose, on the first day of the month to buy, say, a ham sandwich, unless she had already decided whether or not she was going to pay for, say, an oil change for her car in the third week of the month. In optimizing, she would have already compared every margin of consumption against every other margin of consumption, and decided only to purchase a particular sequence of goods which yielded the truly maximum level of satisfaction. In this paper, we hypothesize that this is not a valid description of the way actual, realistic (and boundedly rational) consumers behave. Instead, we will posit that a more accurate description of consumer behavior would be something along the lines of:

A reasonably rational consumer gets paid on the first of the month. At that time, she decides whether or not she wants to purchase a ham sandwich for lunch that

day. She makes this decision, not on the basis of an exhaustive comparison of all her affordable consumption bundles, but instead on some simplifying heuristic which *roughly* compares the costs and benefits of that sandwich. She buys the sandwich if she estimates, at the time of purchase, that the benefits outweigh the costs. When she gets to the third week of the month, she will buy the oil change if she estimates, at *that* time, that the benefits outweigh the costs, and if she still has money left over after all of the purchases she has made using the same heuristic over the previous three weeks. If she doesn't still have enough money, she will put off buying the oil change until the next month, even if, in hindsight, she wishes she had bought less of some other good(s) so that she could afford the oil change in the current month.

Clearly, if consumers behave in a way which is inconsistent with full rationality, then it will be very possible that they end up making decisions which lead to a strictly less-than-optimal level of utility. However, it is not *necessarily* the case that these less-than-fully rational consumers will be any worse off than they would have been had they behaved rationally. The most important result of this paper is to demonstrate that, even if consumers do not behave in a manner consistent with full rationality/backward induction, it is still possible for that consumer to achieve the highest attainable level of utility from her available wealth, given her preferences and the price vector that she faces. That is, "full rationality," in the usual sense, on the part of consumers is not a necessary condition for a Nash equilibrium in an overall market game.

To see why, consider the way in which we economists typically use the term “utility maximization.” Although we rarely have the need or the occasion to explicitly recognize the difference, it should be clear that, at times, economists use this term in at least two distinct ways:

1. “Utility maximization” can refer to the *utility maximization algorithm*, or the process/strategy which guarantees that, *for any possible* price vector, the consumer will allocate her wealth in such a way as to realize the highest possible satisfaction.
2. “Utility maximization” may also refer to the idea that the consumer’s actions/strategy constitute a best response to a particular strategy profile of other actors (such as firms and possibly other consumers) whose decisions and behavior have an effect upon the original consumer’s payoff. Usually, this means that the consumer’s strategy is such that the consumer is realizing the highest possible satisfaction, *given some particular* price vector.

Obviously, these two uses are not, strictly speaking, equivalent to each other. When we are careful, we certainly realize this. But if we aren’t careful, we run the risk of punning on these two different meanings of the same term, and conflating the idea of global optimization, or of subgame perfection, with the weaker idea of a best response to one specific set of strategies. When we are mindful of the difference, we immediately recognize that condition 2 is all that is necessary for a particular consumer strategy to

constitute part of a Nash equilibrium, although condition 1 is necessary if we wish to require that Nash equilibrium to be subgame perfect.

Our primary goal in this paper will be to demonstrate that there are consumer strategies which are both distinct from the utility maximization algorithm of sense 1, but which can nonetheless satisfy the utility maximization requirement of sense 2, when played against certain firm strategies. Moreover, we will demonstrate that the firm strategy profile (to which these boundedly rational consumer strategies are a best response) consists of precisely those firm strategies which are themselves a best response to the consumer and to all of the other firms. That is, we will demonstrate the existence of a class of Nash equilibria in an overall market game in which firms compete in prices, which has not yet been discussed. This new equilibrium is not at odds with the uniqueness proof for the Bertrand equilibrium, because it does not satisfy the conditions of that proof's hypotheses. Here, we will be explicit about making the particular demand curve for each firm's product an endogenous result of consumer strategy choice. The structure of the Bertrand model assumes a single demand curve. This single demand curve assumption is usually thought of as being itself the result of the assumption that consumers will employ the utility maximization *algorithm* (number 1, above) as their strategy choice. As we have pointed out, if we wish to characterize the full set of Nash equilibria in the market game, this assumption is not appropriate², because this algorithm is not a necessary condition for equilibrium.

In order to begin to describe this new equilibrium, we must first consider the justification for why we ought to consider it plausible that consumers would employ a

² Although, as we have also pointed out, that assumption is required if we wish to describe only the set of *subgame perfect* Nash equilibria.

boundedly rational demand strategy which is distinct from the utility maximization algorithm. We therefore return to the idea of the budget tree.

For all but the most trivial of problems, the budget tree is likely to be a very large object. Clearly, if the entire structure of the tree were known by the consumer, and if the consumer had the computing power to fully assess the tree by comparing all the terminal nodes, the algorithm for obtaining an optimal solution is well-understood. But the question is, how realistic is it to believe that these assumptions are all true?

As an analog to this problem, consider a different complicated problem, in which it is readily apparent that the true optimal sequence of choices is not known, literally by anyone: the game of chess. Although this is a more complicated scenario than the budget tree by nature of the fact that chess is a strategic situation involving the interaction of two players, the fact remains that we can appropriately conceptualize of the game as a tree, and the task of players is to find the best set of sequential choices to navigate their way down that tree towards an optimal outcome.

In this setting of chess, it is not the least bit controversial or heretical to claim that even the very best players (both human and machine) are, the vast majority of the time, not playing optimally—or more specifically, are not playing in a manner consistent with subgame perfection. It is certainly not the case that any player, even today's most powerful computer programs, are able to play the game of chess using backward induction, by meticulously calculating the value of every path and subtree, and then selecting the one optimal sequence of moves that will guarantee victory. Moreover, by virtue of the fact that chess is in fact a finite game of full information (there are a finite number of pieces, and, by rule, the game ends in a draw if 50 moves elapse without either

player capturing an opponent's piece), we know in principle, by Zermelo's theorem, that it must be the case that either white has a strategy which will ensure that white will never lose, or else black has such a strategy. However, in practice, we have obviously seen many games, even at the very highest levels of skill, were white wins (and black loses), *and* likewise very many games where black wins (and white loses) instead. This pair of observations is obviously inconsistent with the idea that the players are playing chess in a way that is consistent with what we would call "full rationality."

Now of course, this doesn't imply that chess players are playing the best that they can, given the binding constraints on their rationality, or on their computational abilities. But this is the entire point of the current paper: these rationality constraints play a non-trivial role in the way in which complicated strategic situations actually unfold. Though it has, essentially, been *proven* that a typical highly-skilled chess player is not using backward induction to solve their game tree, we will, for now, merely *hypothesize* that a typical boundedly rational consumer is not using backward induction to solve their budget tree. We will offer a candidate non-backward-induction strategy, and will begin to investigate the consequences of a consumer's employment of that strategy.

In order to do so, we will begin by attempting to draw some additional intuition from our chess example, and we will use one tree that is definitely too large to assess rationally (the game of chess) as a metaphor for what we think might be another such tree (the budget tree).

Clearly, no one actually knows the set of Nash equilibria for a game as complicated as chess. Players do not know the optimal mapping from the set of possible board states to the set of possible moves. And yet, players play chess anyway. What does

this imply for us as economists, if our goal is to be able to describe and predict behavior? Are we lost in the “wilderness of irrationality”? Maybe not. Due to the structure of the game itself, we might be able to make some general observations concerning the way in which players make choices in a setting where the full game structure is too complex to be considered at once.

Broadly speaking, we know that chess players are goal-oriented. Individual moves do not seem to be made with an eye towards whether those moves are located on some optimal path, fully enumerated from opening move to checkmate. Rather, moves seem to be dedicated towards achieving some *intermediate* goal (a hallmark of rational behavior, even despite the fact that real-world chess players are not “fully rational”). Intermediate goals may be things such as control of the center of the board, forking or skewering of opponents pieces to either capture material or at least limit the opponent’s plausible moves, etc.

In pursuing these intermediate goals, players might plausibly consider different candidate moves, visualizing the state of the board that would result from the candidate move and the opponent’s likely response (possibly several moves into the future). When comparing the attractiveness of different potential moves, players probably use a limited version of backward induction, in the sense that they are essentially selecting, or attempting to select, the most favorable future state of the board when they select a best move at any point in the game. However, the values or rankings assigned to the different candidate states of the board are conspicuously *not* derived from the full enumeration and evaluation of all the possible sub-trees following the candidate state. In some way, players seem to assign *directly* an intrinsic value to each intermediate board state that

they consider. The preferences over potential future board states cannot be perfectly derived from the precise evaluation of the underlying structure of the game (or else players would be truly playing optimally). In some sense, preferences over future board states must be primitive, in that they are not *derived* directly from any other set of preferences (such as those over terminal nodes of the game, for instance). The player's actual preference rankings of these intermediate states are the result of both the skill and the experience of the player, and the more "accurate" the player is in his rankings, the more successful he is likely to be in terms of winning chess games.

However, this approach to playing complicated games does not seem to be modeled in any existing concept in game theory. When describing a game, we refer to the set of players, the strategy set of each of those players, the payoff matrix, etc. In essence, a player is defined by his available options and the final payoffs. There is no structure that accounts for the possibility of differences in skill or in rationality of the players, even though it is precisely these differences which account for why "good" players will consistently beat "bad" players, even if a game which is nearly symmetrical, such as chess. Differences among players in the "internal" representation of the actual game are likely a better tool for predicting actual *outcomes* of games than simply looking at the set of equilibria. In other words, in complicated strategic situations, player behavior is probably guided by *relatively* simple models of the actual game, rather than by a hyper-rational evaluation of the actual game itself.

On some level, we know that reasonably smart agents will form simplifying models of complicated situations to help guide their behavior, because that is, in fact, precisely what we do as economists and as scientists. If we wish to give ourselves an

accurate picture of actual consumer behavior when the consumer tries to make the best choice in a setting as complicated as a budget tree (complicated in the sense that a typical tree is likely to have a truly enormous number of nodes and branches), we should consider the possibility that these consumers do precisely what we would do: form a simplifying model to help guide their choices.

Drawing intuition from the case of the chess player, we can now return to our budget tree, and hypothesize how a reasonably “smart” consumer might try to solve the tree with which he is faced. As described above, the tree represents a sequence of decisions to be made concerning the potential purchase of many goods. Two features (at least) of a rational agent ought to show up in this decision process. Firstly, the consumer should be goal-oriented, in the sense that he is *attempting* to make the sequence of purchase decisions which yield the highest possible satisfaction, given her income, preferences, prices, and, perhaps, her scarce computational power, as well. Secondly, in making any individual decision, she ought to, on some level, weigh the costs of a potential purchase against the benefits of that purchase, deciding to buy only when her estimate of the benefits are at least as great as her estimate of the costs.

We can use the concept of a value function as a way of generating a decision rule which will satisfy both of the above. By construction, we shall assume that, at any node on the budget tree, the consumer selects a candidate bundle to consider purchasing. This candidate bundle, a vector in \mathbb{R}^L_+ , essentially specifies the direction that her consumption path will follow through the consumption space, should she decide to purchase the candidate bundle. A decision must be made where the consumer will express a preference between two future states:

1) Where a “yes” decision has been made at that node: the consumer accrues some incremental utility from the consumption of the partial bundle of goods bought at that node, and then finds herself at the next node down the tree along the branch defined by the partial bundle she just purchased. Her available wealth at this next node has been reduced by the price of the purchase made.

2) Where a “no” decision has been made: the consumer does not accrue any additional utility, but also expends no wealth, and then finds herself considering a different possible partial consumption bundle, but at the same node.

Basically, the consumer does not purchase anything until she finds a partial consumption bundle which takes her to a place in the consumption space (augmented by her remaining wealth), which she prefers to her current location in her consumption-wealth space.

A reasonably rational agent will, like the chess player, make her decision in a way consistent with her preference ranking over those two subsequent states, but the manner in which she ranks those states is not *necessarily* consistent with the imputed ranking dictated by the fully rational assessment of her underlying preference function over goods. In the extreme case, a consumer’s value function for each state will be calculated by the full and correct summation of the incremental utility gained at all nodes, along the truly optimal path to the best terminal node, beginning at the root of each subtree. In this case, the consumer who uses this decision rule is guaranteed to select an accumulated final consumption bundle which is equal to a bundle which would have been selected by

backward induction. Again, with the assumption of full rationality, this approach is equivalent to the standard approach.

However, what if our consumer is less than fully rational? What if she makes mistakes in calculating the value function for each sub-tree, or what if she lacks the computing power to perfectly calculate the true value of each sub-tree? Is it possible to still be able to make concrete predictions concerning the consumption behavior of an agent who is imperfectly, though still reasonably, rational?

The consumer's decision reduces to her answer to the following question: was the incremental utility from a "yes" decision worth the income which could have been saved by instead deciding "no"? At the time that the decision must be made, the concrete trade-off is between the potential purchase of the good(s) in question, and the actual money necessary to make the purchase. The consequences of varying degrees of rationality come into play only when the consumer attempts to place a value on the money spent, in terms of the utility which could have been gained if that same amount of wealth had instead been spent on a different good. In other words, rationality plays a role only in determining how accurate the consumer is in weighing the incremental utility gained from a potential purchase, versus the true opportunity cost of making that purchase. To the extent that a consumer is not fully rational, there is a breakdown in the ability of that consumer to actually calculate a purchase's true opportunity cost.

In the absence of the computational ability necessary to perform the standard optimization algorithm on the budget tree, the consumer will have to do what any rational agent would do in a setting where she faces a complicated decision which he may not be

able to fully assess all of his options: she will use a model to simplify the problem, and to recommend her *best guess* as to what the best decision would be.

One candidate solution technique for this type of tree, which might plausibly be employed by imperfectly rational consumers, would be to allow the purchase price of any good to, in some way, act as a stand-in for the actual fully accurate calculation of a good's opportunity cost, in terms of the utility to be gained by consuming the next-best alternative forgone when making any purchase. In other words, if consumers behaved in such a way as to estimate the minimum amount of utility that one could expect to gain, at the margin, for each dollar spent, then the cost-benefit comparison made at each node becomes much simpler. This implies that the actual decision rule used by consumers is, essentially, "Is the good that I am considering buying *worth* the price I have to pay for it?" In this case the word "worth" refers to a more direct comparison between utility and money, rather than the more intricate comparison between the utility of the good, and that of its true opportunity cost (the other things that could have been *bought* with that money), as we have in the standard model. Again, the estimated "value," in utility terms, of money, is intended by the consumer to serve as a stand-in for the fully rational calculation of the value of the relevant sub-trees that the consumer *would* make, if only they had the computational power to do so.

Consider how consumers utilizing this type of simplification might solve the tree. Qualitatively, the decision rule used by these imperfectly rational consumers would operate in a manner quite similar to the fully rational consumer who is able to perfectly calculate the value function for each sub-tree. At each node, the incremental utility gained by a possible purchase is known, and agents simply compare that utility to the

estimate of the opportunity cost of that purchase. In this case, however, the estimate of a potential purchase's opportunity cost is solely a function of the good's price, and the consumer's estimate of the utility value of a dollar. In the standard case, the opportunity cost of a purchase is a function of the consumer's preferences over all goods, as well as the prices of all goods, and the set of all purchase opportunities remaining. To put it another way, the fully rational consumer uses the true (calculated) valuation of each node, while the less rational consumer uses only an estimated valuation of each node.

For a given estimate of the utility value of a dollar (call this value μ , measured in utility/\$), the decision rule will say to make a "yes" decision whenever the incremental utility gained is greater than or equal to the bundle's price, multiplied by μ . In this case, the consumer judges that the purchase is "worth" its price, or that the purchase is a good deal, and so will choose to buy. Otherwise, if the incremental utility of the potential purchase is strictly less than μ times price, the consumer will opt not to buy. In either case, the consumer must choose not to buy if her available wealth is less than the bundle's price. In the case that the consumer chooses not to buy a particular bundle at a particular intermediate node, she then considers purchasing a different bundle. She will continue considering different candidate bundles at the present node until she either finds one she can say yes to (in which case she moves on to the resulting next intermediate node and repeats the process), or until she has exhausted all of her options (at which point further consumption ceases altogether).

It should be noted that, for a given price vector, and for an appropriately chosen value of μ , this decision rule will guarantee that an optimal bundle (from the standard version of the problem) will be chosen in the tree version of the problem. To see why this

is so, recall that one of the consequences of the standard solution to the CBP is that, at an optimal bundle, the marginal utility per dollar spent on each unit of each type of good must be equal for all goods, otherwise, total utility could be increased by spending a little less on the good with the lowest MU/\$, and a little more on the good with the highest MU/\$. So, it is correct to say that, for any optimal consumption bundle in the standard problem, there is exactly one number which describes the marginal utility per dollar of all goods which appear in that bundle with strictly positive quantity. *If* consumers utilizing the simplified decision rule described here to guide their choices on the budget tree were to calibrate their estimate of μ “correctly,” so that the value of μ were the same as the MU/\$ of all goods in the optimal bundle under the standard solution, then they will choose “yes” only for partial bundles with marginal utility per dollar greater than or equal to those goods in the standard bundle, and choose “no” for any and all other goods. Therefore, by the time the decision process terminates (when the consumer has exhausted all or almost all of her wealth), the only goods which would have been selected into the final consumption bundle are the same goods which would have been selected into one of the optimal bundles in the standard problem.

Therefore, this model used by the consumer (using “model” in the sense that this particular decision rule is intended to serve as a simplification for the actual value function used to preference-rank sub-trees) has the advantage of simplifying the complicated budget tree to the point that a decision may be reached at each node based not on any intricate calculations of the value function of a large sub-tree, but based solely on the two pieces of information which are available and directly observable at that node (the good’s price and its potential incremental utility/marginal utility). In other words,

this process is *forward-looking*, in the sense that it prescribes a decision at each and every node, *without* the necessity of any calculations of the value of anything further on down the tree than the present node. Furthermore, this has the additional attractive feature that it is also *useful*, in the sense that, for an appropriately chosen μ , this decision rule delivers a final bundle which is truly optimal under a given price vector.

Despite these similarities between the two formulations of the consumer budget problem, we shall show that, in the case of agents who use the simplified “model” as their decision rule, the nature of the strategic interaction among firms competing in prices is fundamentally altered, with respect to the standard solution. Specifically, the utilization of the simplified model as a demand strategy by the consumer changes the shape of the firm best response functions for the pricing decisions of producers of the goods bought by the consumer. It is these resulting altered competitive incentives among firms which ultimately enable the otherwise less-than-fully rational demand strategy to constitute a best response by consumers in the final strategy profile.

In the sections that follow, we hypothesize here that there may be an upper limit as to how large or complex a tree can be for a reasonably rational consumer to be able to solve it using backward induction. At some point, the number of distinct terminal nodes which must be compared (not to mention the number of distinct histories to each of those nodes) becomes too large for the consumer to practicably evaluate fully³. Conceptually, it is as if there is some horizon, beyond which the consumer simply cannot see. If the tree is small enough so as to not stretch to that horizon, then the consumer can employ backward induction without any trouble, and can thereby guarantee herself an optimal outcome. On

³ Especially during the timeframe in which a single sequential consumption decision is typically made.

the other hand, if the tree is larger than some critical size, then we can think of that tree as stretching over the consumer's horizon. They cannot see the entire tree, and cannot contemplate all of its nodes and branches at once. As a result, if the consumer is to make any decision at all at the nodes that she can see, she must do so using a tactic other than backward induction. In this paper, we will attempt to develop a model in which boundedly rational consumers make consumption choices in this sort of "over-the-horizon" setting. We will assume that the entire budget set is an object which is too large to fully assess at one time, and so the consumer must make smaller, incremental decisions, employing only a limited amount of information at each step along the way. We will begin to describe what we shall call the "satisficing demand strategy," a process by which the consumer makes purchase decisions, not as the result of a truly optimizing process, but as the result of a process where the consumer selects a minimum acceptable level of satisfaction, and then buys whatever items she happens to come across which meet or exceed that satisfaction threshold.

CHAPTER III

THE CONSUMER SATISFICING DEMAND STRATEGY

The Optimizing Consumption Strategy as a Sequential Process

At the heart of the difference between the optimizing and satisficing demand strategies are the differences in the way the concept of opportunity cost is expressed. In a truly optimizing process, the consumer must, explicitly or implicitly, compare every available option, and select an option which is most-preferred among the set of available choices. In this way, the actual action choice prescribed by an optimizing process is guaranteed to be at least as good as every other available action choice.

In a satisficing process, however, the consumer selects a minimum acceptable level of satisfaction, and then simply chooses the first action choice she encounters which provides at least that minimum acceptable level. This implies that, unlike an optimizing process, a satisficing process is inherently path-dependent. If we define set S as the set of all action choices which *could* be selected by the satisficing process, and set O as the set of all action choices which could be selected by an optimizing process, then if S is non-empty, it must contain O . If $S = O$, for some set of parameters describing the overall budget problem, then we will say that the satisficing process is *narrowly equivalent* under those parameters. If $S = O$, for *every* set of parameters which might describe the overall budget problem, then we will say that the satisficing process is *strategically equivalent* under those parameters. If O is a proper subset of S , then we will say the satisficing process is *less discriminating* than the optimizing process under those parameters. Again, the notion of a path is of vital importance here, since any element of S may be selected, if

it is the first such element on which the satisficing consumer happens to consider, even though some elements of S may be strictly preferred to others. By definition, the consumer must be indifferent between all elements of O , so path or order of consideration of the available set of options does not matter under an optimizing process.

Nonetheless, to emphasize the point that the *fundamental* difference between optimizing and satisficing is indeed the concept of opportunity cost, we will take a short detour, and attempt to describe what an optimizing process would look like, if we imposed the requirement that that process be sequential in nature.

With every consumption path that we will consider in this paper, we will assume the requirement of the *irreversibility of consumption*. That is, once an object has been consumed, it may not be *un*-consumed. Although we might imagine a situation in which a consumer might buy a product, and then later change her mind and, say, return that product to the store for a refund, we will say that this situation does not actually reflect what we really mean by “consumption.” In other words, no item is truly consumed until it is no longer possible to return, nor otherwise undo the choice to purchase, that item. This will imply that any path through the consumption space which is to be considered a *consumption path* must be non-decreasing in the quantity of each individual good.

A consumption path will begin at the origin, and then move through the consumption space (again, in non-decreasing fashion) until it terminates at a point representing some combination of goods. This terminal combination of goods will be the final consumption bundle selected by the sequential process. The *final* bundle will therefore be determined by the summation of a sequence of arbitrarily small *partial*

bundles, and the sequence or history of those partial bundles is precisely what we mean when we refer to the consumption path.

In order for this type of consumption path to be consistent with true optimization, there must be a fully rational cost-benefit comparison which takes place for every sequential step along this path. Here, this cost-benefit comparison may be described in terms of the standard notion of opportunity cost. Consider some arbitrary consumption vector \mathbf{X} , which begins at the origin and is fully contained within the budget set. If we interpret this vector as one step along the consumption path, then in order to assess whether or not this step is consistent with fully rational optimization, we must compare the costs and benefits of taking this step. A simple framework for performing this comparison is to consider which consumption opportunities are preserved, and which consumption opportunities are forgone, once this sequential step is taken. In other words, under the irreversibility assumption, if the consumer purchases the partial bundle \mathbf{X} , some elements of the budget set will still be attainable afterwards, and other elements of the budget set will not longer be attainable afterwards. Specifically, if the consumer purchases \mathbf{X} , then the set of still-attainable consumption opportunities, call it A , is defined⁴ as $A \equiv (\mathbf{X} + \mathbb{R}_+^N) \cap \beta$, where N is the total number of distinct goods available, and β is the budget set. The set of consumption opportunities forgone, call it $\sim A$, is the complement of A in β .

Quite simply, any partial consumption bundle \mathbf{X} is consistent with full optimization if and only if the most-preferred element of A is at least as good as every element of $\sim A$.

⁴ \mathbb{R}_+^N is the closure of the positive orthant.

Likewise, any *path* through the consumption space may be thought of as the summation of a sequence of arbitrary partial consumption vectors ($\mathbf{X}^1, \mathbf{X}^2, \mathbf{X}^3, \dots$). Any path is consistent with optimization if and only if, for all i , $\sum_{n=1}^i \mathbf{X}^n$ has the property that the most-preferred element of A^i is at least as good as every element of $\sim A^i$, where A^i and $\sim A^i$ are defined analogously to A and $\sim A$, but for the summation of i individual partial consumption vectors, rather than for a single partial vector \mathbf{X} (fig. 2).

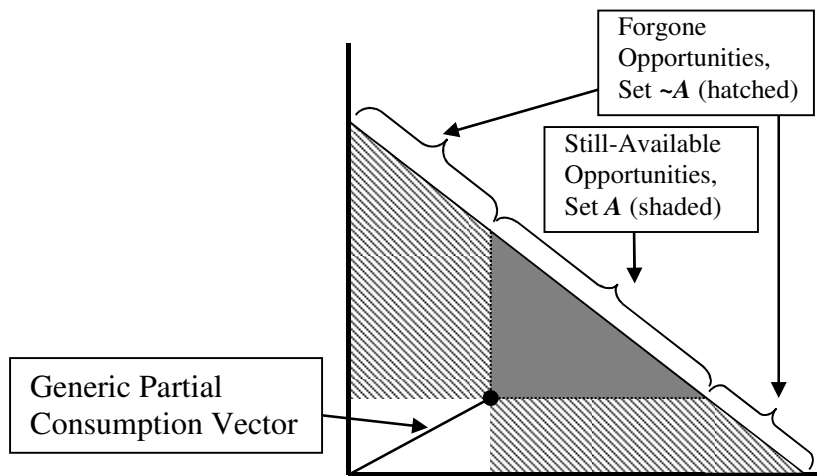


Figure 2. Sequential Optimizing Consumption.

Obviously, this consumption path formulation of the consumer's optimization process does not add very much to our understanding of standard consumer theory. The process needed to identify the most-preferred element of A or A^i (and then also to compare that element to every element of $\sim A^i$) is equivalent to the process needed to identify the most-preferred element of β . In other words, our attempts to describe the optimization process in terms of a sequential consumption path amount to little more than

saying “select the optimal bundle by finding a point of intersection⁵ between the budget hyperplane and the highest attainable indifference curve, and then select literally any non-decreasing path from the origin to that point.” Clearly, this adds nothing to our understanding of consumer optimization, per se. In fact, it is probably implicit in the standard demand model that there will be some path from the origin to an optimal bundle that the consumer actually follows as she makes sequential purchases; it is just that the path itself is trivial, and the location of the set of optimal bundles is all that matters.

But the reason that we have taken the trouble to outline this sequential process is precisely because of the lack of important differences between the sequential description of the optimizing process and the standard, path-independent description of the same. Since satisficing processes are *inherently* path-dependent, and since we generally think of optimizing processes as path-independent, we want to be able to demonstrate that any differences between optimizing demand and satisficing demand are, in fact, due to differences in the role played by the concept of opportunity cost, and *not* due in any way to the fact that one process is path-dependent and one is not.

Now that we have (trivially) described the optimizing demand process as a sequential process, we may more directly compare it to the inherently sequential satisficing demand process.

⁵ A point of tangency if preferences are strictly convex and there is an interior solution to the utility maximization problem.

The Sequential Satisficing Process Outlined

All vectors are N -dimensional, where N is the number of distinct goods available for purchase by the consumer. The process begins at period 1 and continues through period T . “Time” is notional, and the value of T is determined endogenously.

Let $\mathbf{0}$ be a vector whose every element is zero.

Let \mathbf{P} be the price vector.

Let \mathbf{O}^t be a vector representing the quantities of each of the N distinct goods which have already been consumed as of the beginning of period t . Assume that $\mathbf{O}^1 = \mathbf{0}$.

\mathbf{O}^T is the consumer’s final consumption bundle.

Let $\mathbf{X}^t \geq \mathbf{0}$ be a sequential *candidate* consumption vector which is under consideration during round t of the consumption process.

Let $\mathbf{X}^t \geq \mathbf{0}$ be a sequential *realized* consumption vector which is added to the physical consumption at the end of period t , so that consumption follows the consumption accumulation function given by $\mathbf{O}^{t+1} = \mathbf{O}^t + \mathbf{X}^t$.

The non-negativity requirement for \mathbf{X}^t and \mathbf{X}^t is the consequence of the assumption of the irreversibility of consumption. Furthermore, we will require that $\mathbf{X}^t \neq \mathbf{0}$, since we wish that each iteration of the satisficing process actually consider purchasing *something* (it is essentially meaningless to consider an iteration in which no purchases are contemplated). We will allow for the possibility that $\mathbf{X}^t = \mathbf{0}$, meaning that, although some purchases were *considered* in period t , none were actually *made*.

Let $\Pi(\mathbf{O}^t)$ be the path-generating process for the candidate consumption vector, so that $\Pi(\mathbf{O}^t) = \mathbf{X}^t$ for all t .

Let $\mathfrak{R}(\mathbf{X}^t, \mathbf{O}^t)$ be the decision rule for sequential purchases, so that $\mathfrak{R}(\mathbf{X}^t, \mathbf{O}^t) = \mathbf{X}^t$.

Together with the price vector and the consumer's wealth and preferences, $\Pi(\mathbf{O}^t)$ and $\mathfrak{R}(\mathbf{X}^t, \mathbf{O}^t)$ will determine the consumers actual consumption path, including the value of T and the final realized consumption bundle \mathbf{O}^T . See Figure 3, following page.

Thoughts on the \mathfrak{R} Function

There are many possibilities concerning the specific nature of the \mathfrak{R} function. A central hypothesis of this paper is that, for a typical consumer, the entire budget set is likely to be an object which is too large to fully assess in a completely rational manner. From a modeling perspective, we would like to be able to explain consumer behavior in such a way as to reflect the practical impossibility of the comparison of all affordable bundles, while retaining the feature that individuals probably act rationally within the constraints imposed by their limited computational abilities. To put it simply, our present way of thinking about how consumption decisions are made involves a fundamental asymmetry between the two sides of a cost-benefit comparison: Consumers are very good at completely assessing the value of goods and bundles that they have directly in front of their faces, but are unable to fully assess the value of the set of all alternatives. Intuitively, this idea is perhaps best explained in a cardinal utility framework; consumers can correctly compute the amount of satisfaction that any particular purchase might give them, but are not able to know with certainty whether that amount of satisfaction is

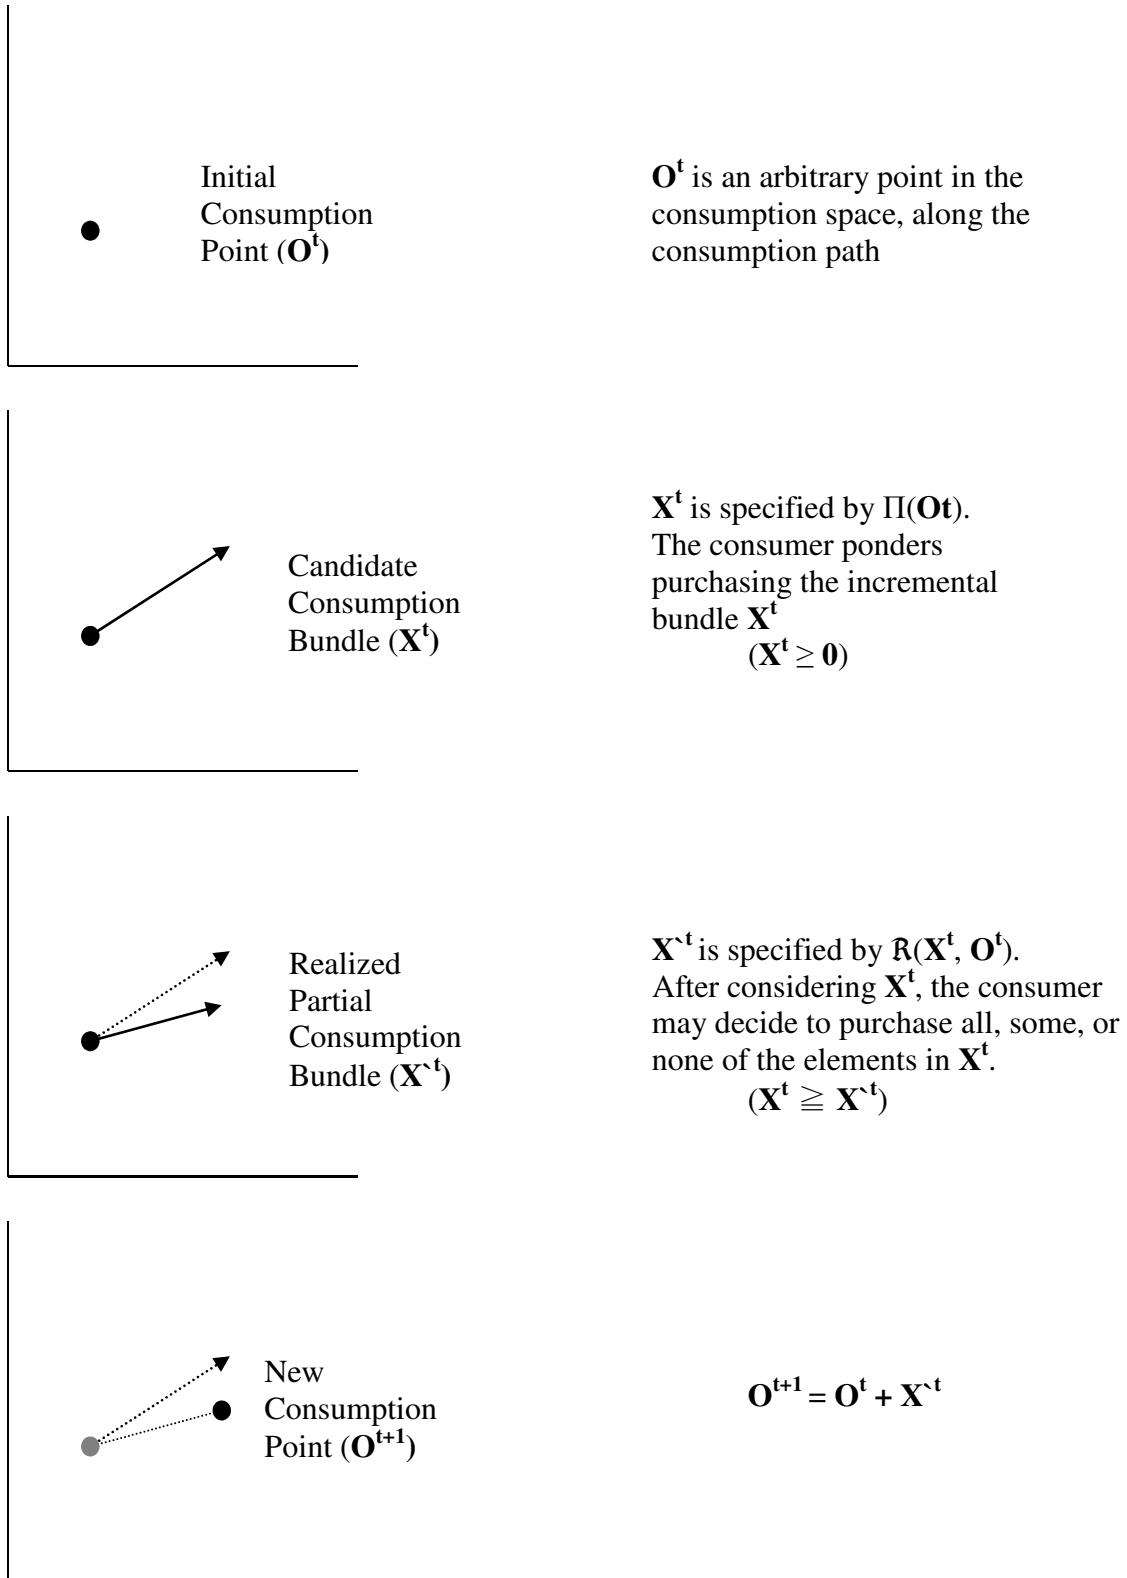


Figure 3. One Generic Iteration of the Sequential Satisficing Process.

greater than or equal to the maximum satisfaction from the set of all feasible alternative bundles. The satisficing consumer must, therefore, select a minimum acceptable level of satisfaction, and then choose to buy bundles which come under her consideration only if they meet or exceed that satisficing threshold.

Note that the satisficing criterion (the minimum acceptable level of satisfaction) is the conceptual analog to the idea of opportunity cost in an optimizing process, as far as the role that it plays in the relevant performance test. An optimizing consumer will (possibly) take a given action if and only if that action is at least as good as the next-best available alternative. A satisficing consumer will choose a particular action that is under consideration if and only if the satisfaction that action yields meets or exceeds the satisficing criterion. If the satisficing criterion itself were *defined* in terms of the highest possible⁶ level of satisfaction, given the entire set of available alternatives, then there ceases to be any difference whatsoever between optimizing and satisficing.

The R Hypothesis

We would like the decision rule employed by the satisficing consumer to reflect the idea that the consumer is unable to fully calculate the true opportunity cost of every potential purchase, but is otherwise rational. When making purchase decisions, she must form some simple *model* which estimates the true (but unknowable) opportunity cost of those purchases, and must use that model or estimate as the yardstick by which she measures potential purchases to separate the “good” bundles from the “bad” (or more

⁶ Of course, this requires that the consumer has sufficient computation capacity to be able to fully assess the entire set of available alternatives.

precisely, the “good enough” bundles from the “not good enough”). Like the chess player who learns to estimate the value of intermediate states of the board (rather than calculating them through backward induction), and uses those estimated valuations to guide her choice of moves throughout the game, we assume that this estimate of opportunity cost is “learned” through experience in some way (as yet unspecified). We further assume that the estimate is a “passably close” approximation to the real opportunity cost, in the sense normally used when defining the term “satisficing.”

Nonetheless, once a consumer has formed this estimate of opportunity cost (in utility terms), we should like to require that any subsequent decision-making based on this estimate be fully consistent with a comparison of benefits and (estimated) costs. In other words, if a consumer uses a (possibly incorrect) estimate of the opportunity cost of a decision, rather than the calculated true cost, the decision process can be thought of as boundedly rational. However, if the consumer ever decides to purchase any goods for which the marginal benefit is strictly less than the estimated cost, it would seem more appropriate to describe this behavior as *irrational*, rather than boundedly rational. To formalize this idea, we will define the following:

The \mathbf{R} Hypothesis:

If a satisficing rule involves forming an estimate of the opportunity cost of a potential purchase, then that rule satisfies the \mathbf{R} hypothesis if both of the following are true:

- 1) The rule never allows any purchases to be made such that any part, element, or subset of that purchase yields a marginal benefit less than the estimate of its opportunity cost.
- 2) If the rule is applied to a candidate bundle for which any part, element, or subset of that purchase yields a marginal benefit greater than or equal to its estimated opportunity cost, then the rule must not fail to actually purchase that part, element, or subset of the candidate bundle, assuming it is affordable.

In other words, we believe that consumers behave constrained rationally, up to their imperfect estimate of the opportunity cost of any choice. If the decision rule is not in some way based upon an estimate of opportunity cost, then the rule does not satisfy the \mathcal{R} hypothesis.

By focusing attention on decision rules which satisfy the \mathcal{R} hypothesis, we can begin to construct a model of consumer behavior which has clear conceptual upper and lower bounds placed on the level of consumer rationality. Consumers are not so rational as to be able to fully assess the entire budget set at once, and so must use an estimate of the true opportunity cost of any purchase. But, once having formed that estimate, they are able to apply it to a decision-making process in such a way as to guarantee that they will never buy any good or service which does not appear to satisfy a cost-benefit test. That is, they are able to behave (constrained) optimally, so that their decisions are no less rational than their estimate of opportunity cost. One implication of this idea is that consumers will be savvy enough to avoid making partial purchases in which one subset of the purchase

has insufficient marginal utility, given its estimated opportunity cost. In other words, if \mathbf{X}^t is decomposable into smaller parts, then the consumer ought to only purchase the *parts* of \mathbf{X}^t which are sufficiently valuable (at the margin) to stand on their own. The consumer should never decide to purchase parts of \mathbf{X}^t which do not individually pass the underlying cost-benefit test. In practice, this implies that we must be careful to construct our \mathcal{R} functions in such a way as to avoid the possibility that excess marginal utility for one part of \mathbf{X}^t compensates for, or cross-subsidizes the purchase of, any other part of \mathbf{X}^t which is insufficiently valuable to have been purchased on its own merits.

Examples of the \mathcal{R} Function

In this section, we will begin to consider what the \mathcal{R} function could or should specifically look like. We will examine a series of potential forms, and then briefly discuss the merits and shortcomings of each. Ultimately, we will not necessarily settle on one universal specification of the \mathcal{R} rule, intended to validly apply in every conceivable circumstance. Instead, we will hope to begin to outline some of the major issues which ought to be considered as we try to implement a specific decision rule which is consistent with the \mathcal{R} hypothesis, for the parameters of any single particular problem we may wish to consider.

One way to loosely categorize the possibilities is to think of the space of all possible \mathcal{R} functions as somehow associated with a continuum of assumptions concerning the level and sophistication of consumer rationality. On one extreme, we might conceive of a decision rule which is completely a-rational; a rule which completely accepts whatever candidate consumption vector it is ever asked to pass judgment upon. In other

words, we might define a completely non-discriminating satisficing rule, which we will call \mathfrak{R}_0 :

$$\mathfrak{R}_0(\mathbf{X}^t, \mathbf{O}^t) \equiv \mathbf{X}^t = \mathbf{X}^t, \forall (\mathbf{X}^t, \mathbf{O}^t)$$

The level of consumer rationality implicit in this rule is low enough to strain the applicability of the term “satisficing,” since it is not apparent that this rule even exhibits a minimum acceptable level of performance or satisfaction (at least, that is, without imposing some external restrictions on the set of possible sequences of \mathbf{X}^t vectors, on the utility function, or on both). It certainly will not satisfy the \mathfrak{R} hypothesis, as it makes no non-trivial attempt to compare costs and benefits. Although \mathfrak{R}_0 is unlikely to ever provide a reasonable representation of actual consumer satisficing behavior, it still may be illustrative to at least consider its existence in principle.

More discriminating (and therefore, more rational) decision rules are of course also possible. We will spend considerable time discussing a particular class of these functions, which are intended to model the behavior of boundedly rational consumers.

This class of satisficing rules of thumb emphasized in this paper is based on the notion of marginal utility per dollar. The consumer selects a minimum level of marginal utility per dollar which any actual purchase must meet or exceed. We will call this benchmark value μ , and it will represent the satisficing criterion that any *potential* purchase must satisfy if it is to become a *realized* purchase. As described above, this class of rules employs a performance test to any candidate bundle. The test is similar to a

cost-benefit test; the consumer buys the candidate bundle if its marginal utility meets or exceeds the satisficing criterion, and does not buy the bundle otherwise.

One way to think of the justification for basing the satisficing rule on the idea of marginal utility per dollar is to compare satisficing with the constrained optimization problem of the fully rational solution. At an optimal bundle, the value of the Lagrangian multiplier λ measures the true shadow value of wealth: the rate at which utility increases as the budget constraint is relaxed slightly. If we make certain standard assumptions concerning the nature of the utility function, then implicitly, it is *as if* the optimizing consumer had imputed a level of “intrinsic” value to each unit of wealth (specifically, λ , measured in utility per dollar), and then performed a simple cost-benefit test on each item in the budget set. For each item, the opportunity cost can be thought of as simply the price of that item, multiplied by λ . It is a necessary condition for inclusion of an item in an optimal bundle that the utility of that item meets or exceeds the opportunity cost. If there is exactly one most-preferred bundle, then this is a sufficient condition as well.

In fact, if preferences are strictly convex and continuous, then there is no element in an optimal consumption bundle (under a given price vector) which does not satisfy the notional cost-benefit test described above, or else it would be possible to improve overall utility while expending the same amount of wealth, by substituting away from these elements in favor of elements of the budget set with higher marginal utility per dollar. Likewise, *every* item in the budget set whose marginal utility per dollar is strictly greater than λ is necessarily in the final consumption bundle. Fundamentally, these statements are true because the constrained optimization problem inherently finds the value of λ for which both of these statements are true.

As such, it is appropriate to think of λ as capturing a great deal of information about the given utility maximization problem. Specifically, λ may be thought of as a *utility index*: a scalar value which summarizes the multi-dimensional budget set, and describes the maximum value of the affordable elements. In a very real sense, λ represents the actual opportunity cost, in utility terms, of purchasing \$1 worth of any good, under the parameters of the UMP. If any of the parameters should change (prices, wealth), then of course the value of λ will immediately change as a result. Of course, this process is still fully rational and globally optimal, because λ itself is *computed* as a function of the fundamental parameters of the utility maximization problem (wealth, prices, and preferences). It is a fundamental feature of the optimization process that the “correct” value of λ is always calculated for any possible set of parameters. Of course, μ is also a utility index, although it is not a perfectly calculated value of the shadow value of wealth. Instead, this index is presumably learned in some way through experience. The difference between μ and λ , as employed in a cost-benefit test, will lead to potentially important consequences in the strategic interactions which take place at the market level, between price-competing firms and the set of consumers.

This specific interpretation of the optimization process is the primary motivation for construction a class of satisficing rules which are based upon the idea of marginal utility per dollar. Essentially, the consumer will form an estimate μ of the utility value of a dollar, and then use that *estimate* (rather than the true, computed value λ) as a stand-in for the actual opportunity cost of each \$1 worth of purchases. The same cost-benefit test as above will be applied to a sequence of candidate partial consumption bundles, and those bundles passing the test will be purchased, and those not passing the test will not be

purchased. The process will continue until the consumer either exhausts her wealth, or until there are no longer any candidate consumption bundles available which can pass the cost-benefit test.

This will allow the consumer to make a *one*-dimensional comparison between any item she wishes to purchase and its (estimated) opportunity cost (equal to monetary cost times μ). This is opposed to the optimizing strategy, which inherently requires that a *multi*-dimensional comparison be made between the item under consideration, and every other element of the budget set. Furthermore, for certain utility functions, if the estimate is *correct*, so that $\mu = \lambda$ for a given parameterization of the budget problem, the satisficing strategy will be narrowly equivalent to the optimizing strategy (More on this in the next section).

Alternatively, if the estimate μ is not necessarily correct, but is instead merely “close enough” to the true value λ , then the utility of the final consumption bundle selected will be “close enough” to the utility of a truly optimal bundle. The phrase “close enough” here represents the minimum acceptable level of performance of the satisficing rule, and the selection of a particular μ , in effect, defines what that minimum acceptable level of performance is.

To continue exploring the possibilities offered by this type of satisficing rule, consider the following concrete example, which we shall refer to as \mathbf{R}_1 .

For each candidate consumption vector \mathbf{X}^t to which the decision rule \mathbf{R}_1 is applied, the consumer will calculate the incremental utility which would be gained *if* that

bundle \mathbf{X}^t were consumed. That is, the consumer calculates the difference between total utility realized with and without the incremental bundle. Define the difference Δ_1 as

$$\Delta_1(\mathbf{X}^t, \mathbf{O}^t) \equiv U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t)$$

If $\Delta_1(\mathbf{X}^t, \mathbf{O}^t)$ is sufficiently large, then \mathcal{R}_1 will select that bundle for purchase. More specifically \mathcal{R}_1 is defined as the following:

If $\Delta_1(\mathbf{X}^t, \mathbf{O}^t) \geq \mu (\mathbf{P} \cdot \mathbf{X}^t)$, then $\mathcal{R}_1(\mathbf{X}^t, \mathbf{O}^t) = \mathbf{X}^t = \mathbf{X}^t$.

Otherwise, $\mathcal{R}_1(\mathbf{X}^t, \mathbf{O}^t) = \mathbf{0}$.

Clearly, \mathcal{R}_1 is a function of μ , so different values of μ will lead to different effective decision rules. Consequently, the decision rule \mathcal{R}_1 ought to be properly indexed in some way by the choice of μ that it employs (i.e., we should use the notation \mathcal{R}_1^μ). For now, we will consider this to be implicit, and will suppress the notation which indexes for μ .

Also, consider the level of rationality represented by this decision rule. It is clearly significantly more sophisticated than \mathcal{R}_0 , since for the first time we now have *some* estimate of opportunity cost, and some attempt being made to compare costs and benefits so as to reach a somewhat efficient allocation of scarce resources (wealth).

However, this rule still exhibits a level of rationality that is likely to be too low, if our goal is to describe the behavior of “reasonably” rational consumers; this rule does not satisfy the \mathcal{R} hypothesis.

To see why, consider the following. We can see that there is a certain troublesome feature of \mathcal{R}_1 : the possibility exists that surplus utility, over and above the threshold value specified by μ , from one good in the bundle \mathbf{X}^t might compensate for insufficient utility from another good, and the consumer might end up expending scarce wealth on goods that do not yield a high enough level of utility, relative to μ .

As an example, consider a consumer who derives a constant 10 utils worth of utility per each unit of good X she consumes, but derives 0 utility from any level of consumption of good Y. Assume the price vector is (1, 1), she has 100 units of wealth, and $\mu = 4$. If the consumer were to employ \mathcal{R}_1 to judge whether or not partial consumption bundle (2, 2) were “good enough” to justify purchasing that bundle, she would find that the marginal utility of good X was high enough to justify purchasing the entire bundle:

$$\Delta_1((2, 2), \mathbf{O}^t) = 20 \text{ for any } \mathbf{O}^t$$

The total cost of (2, 2) is 4.

So, the incremental utility gained by the bundle \mathbf{X}^t , 20, exceeds the estimate of the bundle’s opportunity cost in utility terms, 16 (which is the number of dollars spent in purchasing \mathbf{X}^t , 4, multiplied by the benchmark estimate for minimum utility per dollar, $\mu = 4$).

Since $20 > 16$, \mathbf{X}^t satisfies \mathcal{R}_1 in this case.

Accordingly, \mathcal{R}_1 will select the entire bundle (2, 2) for purchase, despite the fact that both units of good Y were completely useless, but came at a total cost of 2 units of

wealth. We submit that it is not satisfactory to consider a demand strategy which supposedly reasonably rational consumers would ever choose to make such a purchase as described here. The specific violation of the \mathcal{R} hypothesis happens, of course, because this rule allows for good Y to be purchased, despite the fact that, in utility terms, the marginal benefit of Y is zero, which is strictly less than its estimated marginal cost ($P_y * \mu$, or 4 per unit). Again, the reasoning underlying the \mathcal{R} hypothesis is that, even if a consumer cannot fully compute the opportunity cost of this bundle, they should nonetheless possess sufficient rationality to know not to purchase anything for which the marginal benefit is less than the best guess of the marginal cost. In order to address this concern, we need to consider a still more sophisticated satisficing rule, which we will call \mathcal{R}_2 .

Given that we should expect that our consumers should be able to look at a candidate consumption bundle both as a whole, *and* in part, we need to be able to allow for the possibility that the consumer consider each distinct element of candidate consumption vectors, and should be able to discard portions of those bundles which do not satisfy the cost-benefit test, while retaining those portions of the bundle which do.

The following is a description of satisficing decision rule \mathcal{R}_2 . For each candidate consumption vector \mathbf{X}^t to which the decision rule is applied, the consumer will consider the incremental or marginal utility offered by each component of that consumption vector, and decide to purchase only those components for which the incremental (marginal) utility, divided by the cost (price) of that component, meets or exceeds the critical value μ .

The satisficing rule \mathcal{R}_2 used here will be purely subtractive. That is, it will only pare down \mathbf{X}^t in order to obtain \mathbf{X}^* . \mathcal{R}_2 will only potentially remove elements from \mathbf{X}^t and will never add anything to it. As a result, for any period t , we must have either $\mathbf{X}^t \geq \mathbf{X}^*$ or $\mathbf{X}^t = \mathbf{X}^*$.

The rule \mathcal{R}_2 makes a utility comparison between $U(\mathbf{O}^t)$ and $U(\mathbf{O}^t + \mathbf{X}^t)$, but does so component-wise. For any vector \mathbf{Q} , define Q_i as the i^{th} component of \mathbf{Q} , and define \mathbf{Q}_i as the vector whose i^{th} element is Q_i , and whose every other element is zero. When employing \mathcal{R}_2 , the consumer calculates $U(\mathbf{O}^t + \mathbf{X}_i^t) - U(\mathbf{O}^t)$ for each $i \in [1, N]$, and uses the resulting value as the measure of the incremental utility gained by purchasing the quantity X_i^t of good i . Define this value as $\Delta_{2i}(\mathbf{X}^t, \mathbf{O}^t)$, or simply $\Delta_{2i}(\mathbf{X}^t)$ as a notional shortcut, so that

$$\Delta_{2i}(\mathbf{X}^t) \equiv U(\mathbf{O}^t + \mathbf{X}_i^t) - U(\mathbf{O}^t)$$

$\Delta_{2i}(\mathbf{X}^t)$ will be used by the consumer to measure the incremental value of each individual component of \mathbf{X}^t , and then to compare this value against the benchmark estimate of its opportunity cost, $\mu * P_i * X_i^t$ (the benchmark, μ , for minimum acceptable marginal utility per dollar, multiplied by the number of dollars, $P_i * X_i^t$, required to purchase the quantity of good i present in the candidate bundle). \mathcal{R}_2 will then construct \mathbf{X}^* in a component-wise fashion. For any element of \mathbf{X}^t for which

$$\Delta_{2i}(\mathbf{X}^t) \geq \mu * P_i * X_i^t$$

\mathcal{R}_2 will add X_i^t as the i^{th} element the realized partial consumption bundle \mathbf{X}^t . For any element of \mathbf{X}^t for which $\Delta_{2i}(\mathbf{X}^t) < \mu * P_i * X_i^t$, \mathcal{R}_2 will place a 0 as the i^{th} component of \mathbf{X}^t . In this way, \mathcal{R}_2 will avoid the some of the possibilities of wasteful purchases that were possible with \mathcal{R}_1 , and thus provide us with a more reasonable way of modeling the behavior of boundedly rational consumers. The expression $\Delta_{2i}(\mathbf{X}^t) \geq \mu * P_i * X_i^t$ describes the decision rule in terms of a cost-benefit comparison, but it is obviously equivalent to express the same rule in terms of a marginal utility per dollar form:

$$(\Delta_{2i}(\mathbf{X}^t) / X_i^t) / P_i \geq \mu$$

That is, purchases which yield a marginal-utility-per-dollar which is greater than or equal to the threshold level μ will be selected, and those that do not meet or exceed μ will not be selected. Still, there is at least one obvious shortfall of decision rule \mathcal{R}_2 , as well, as pertains to the \mathcal{R} hypothesis. One of these shortfalls has to do with the magnitude of the candidate consumption vectors. It is entirely possible that, even in cases where $\Delta_{2i}(\mathbf{X}^t) \geq \mu * P_i * X_i^t$, for some i , the rule could violate the \mathcal{R} hypothesis by consuming too many units of good i .

Consider the case where the utility function is additively separable, and marginal utility in good i is decreasing according to $\partial U / \partial x_i = (10 - x_i)$, and $P_i = 1$. If $\mu = 5$, then the consumer ought to purchase up to 5 units of good i . Specifically, if X_i^t was less than or equal to 5, then there is not an immediate problem with respect to the \mathcal{R} hypothesis, as \mathcal{R}_2 will correctly accept the entire candidate quantity of i , and $X_i^t = \hat{X}_i^t$. However, if X_i^t was, say, 6 units of good i , then we potentially have a situation where the utility of the first 5

units of i sufficiently compensates for the lower-than-threshold level of utility of the sixth unit, just as excess utility from good X “subsidized” the inefficient purchases of Y in the previous example. Consider a concrete example:

If $O_i^t = 0$, and $X_i^t = 6$, then $\Delta_{2i}(\mathbf{X}^t) = 42$. This exceeds the estimate of the opportunity cost of X_i^t , which is 30 (since $X_i^t = 6$, $P_i = 1$ and $\mu = 5$). Therefore, \mathbf{R}_2 will select $X_i^t = X_i^t = 6$. However, the 6th unit of good i , by itself, does not pass the cost-benefit test. The additional utility of the 6th unit is 4.5, and its estimated opportunity cost is 5. If good i must be purchased in blocks of 6 units, then there is no problem. However, if good i can be purchased by the single unit, then our boundedly rational consumer ought to know better than to purchase this 6th unit, and hence we have demonstrated that \mathbf{R}_2 violates the \mathbf{R} hypothesis.

There are two ways to address this particular concern. The first, is to shift emphasis from the discrete to the continuous version of this idea, and consider this idea in the limit as we let $\|\mathbf{X}^t\| \rightarrow 0$ for all t . Therefore, the issues related to the magnitude of \mathbf{X}^t become irrelevant, as the consumer considers each differential unit of each good on its own. This is the tactic we will eventually take in this paper, as we later describe the geometry needed to derive the consumer’s satisficing demand curve.

Alternatively, we could resolve this \mathbf{R} hypothesis issue, and continue to consider the discrete version of this rule, if we redefined \mathbf{R}_2 to allow for the possibility that X_i^t could be an intermediate value between X_i^t and zero. For example, given some \mathbf{O}^t , if it is the case that X_i^t is divisible into smaller units, then the consumer probably ought to be attributed with sufficient computational ability to be able to sequentially apply the decision rule to each of the smallest feasible units on the interval $[0, X_i^t]$. In this way, the

consumer can select only the number of units for which the benefit exceeds the (estimated) cost, and decline to purchase any units for which the opposite is true. Whereas \mathbf{R}_2 applied the same estimated cost-benefit test as \mathbf{R}_1 to each individual component of \mathbf{X}^t , rather than to the entire vector, this new rule \mathbf{R}_3 is defined in such a way as to apply the same cost-benefit test as \mathbf{R}_2 , not only to each component, but also to the smallest feasible unit of each component. That is, \mathbf{R}_3 would calculate $\Delta_{2i}(\mathbf{X}^t)$, not just for each component of \mathbf{X}^t , but for each smallest unit of each component of \mathbf{X}^t .

If good i is available in continuous quantities, such as -- hypothetically-- gallons of gas or pounds of ground beef, then it makes sense to describe the satisficing criterion in terms of the partial derivative. In this continuous case, $\mathbf{R}_3(\mathbf{X}^t, \mathbf{O}^t) = \mathbf{X}^t$, where the i^{th} component of \mathbf{X}^t is largest number (or supremum) on the interval $(0, X_i^t]$ for which $\partial U/\partial x_i \geq \mu * P_i$, or zero if no such quantity exists.

If a good is not continuously divisible, but must instead be purchased in discrete units, such as jars of tomato sauce or number of whole turkeys, etc., then define the smallest indivisible unit of good i as the number I_i . For example, I_i for pineapples⁷ would then likely be 1 (pineapple), while I_i for cans of beer might be 6 (cans). Also define \mathbf{I}_i as the vector whose i^{th} element is I_i , and whose every other element is 0. If the good truly is only available in discrete quantities, then it will be assumed that X_i^t for that good must be some whole number multiple of I_i , as will X_i^t .

⁷ This is assuming that pineapples are sold by the count, not by weight. If sold by weight, then this good might or might not be more properly thought of as being continuous. Some peripheral concerns exist concerning how best to model this situation, though we do not intend to address them here.

If we wish to define \mathbf{R}_3 in this situation, we must do so in terms of $\Delta_{3i}(\mathbf{X}^t, \mathbf{O}^t)$, where:

$$\Delta_{3i,n}(\mathbf{X}^t) \equiv U(\mathbf{O}^t + n \cdot \mathbf{I}_i) - U(\mathbf{O}^t + [n-1] \cdot \mathbf{I}_i)$$

$\mathbf{R}_3(\mathbf{X}^t, \mathbf{O}^t) = \mathbf{X}^t$, where the i^{th} component X_i^t , X_i^t , is defined as the $n \cdot I_i$, where n is the largest integer on the interval $[1, X_i^t / I_i]$ such that $\Delta_{3in}(\mathbf{X}^t) \geq I_i \cdot P_i \cdot \mu$. X_i^t is defined to be zero if no such n exists.

Going forward, we will assume that any subsequent decision rules we may wish to consider will implicitly have this “incrementalist” approach, and can therefore select intermediate quantities between 0 and X_i^t , as necessary.⁸ Basically, this allows for the consumer to be able to “optimize”⁹ over a single, one-dimensional segment (the interval $[0, X_i^t]$), while still being unable to perform the more sophisticated multi-dimensional true optimization problem over the entire budget set.

When describing subsequent decision rules \mathbf{R}_x , as well as their corresponding $\Delta_x(\mathbf{X}^t)$ functions, we will assume that each is expressible in terms of a similar “ n ” and “ I_i ” notation to convey the idea that increments between 0 and X_i^t are possible. However, for simplicity and clarity, as well as to more easily emphasize the difference between ensuing decision rules, we will assume that the fact that intermediate values may be found for each good is implied. We will therefore focus our ensuing notation on highlighting the interaction between the various components of the consumption vector,

⁸ The same effect would be achieved by restricting the magnitude of \mathbf{X}^t , so that each X_i^t was always the smallest feasible unit of good i . If \mathbf{X}^t always only involved the smallest feasible unit of each good, then \mathbf{R}_2 and \mathbf{R}_3 would necessarily be equivalent. Some extra technical concerns would need to be addressed, though, if we were to decide to restrict the set of \mathbf{X}^t vectors in this way, though.

⁹ Obviously, the consumer is not necessarily *optimizing*, per se. They are doing the best they can over the specific interval $[0, X_i^t]$, given their particular estimate of opportunity cost.

keeping in mind that, in the continuous version of the satisficing process, there ceases to be any relevant difference between \mathcal{R}_2 and \mathcal{R}_3 .

In addition to the quantity concerns described above, the more pressing objection to both \mathcal{R}_2 and \mathcal{R}_3 at the moment is that these rules, while still likely better approximations of consumer behavior than \mathcal{R}_1 , will not allow us to adequately describe behavior in situations where there are interactions among distinct goods in the utility function. If the utility function is additively separable, then \mathcal{R}_3 (as well as the continuous version of \mathcal{R}_2) will actually satisfy the \mathcal{R} hypothesis. But if the utility function involves interactions between distinct goods, this will not be the case. Specifically, if we wish to consider situations where the consumer uses a satisficing process to select bundles of goods which include perfect complements or perfect substitutes, for instance, merely assessing incremental utility in a component-wise manner will not allow us to correctly measure the proper marginal utility of items in a candidate consumption bundle. If the utility function includes 2 goods J and K which are perfect complements, then \mathcal{R}_3 will never allow the consumer to purchase any positive amount of either good (let alone both). When the consumer considers the marginal utility of J, she will always find that the marginal utility of J is equal to zero, and will accordingly never decide to purchase any additional units of J. Similarly for good K. This is true even if it is the case that the consumer might derive significant surplus over and above the estimated opportunity cost through the consumption of both goods together. If we wish to employ a rule which is capable of dealing with this objection, we will need to consider a still more sophisticated decision rule, \mathcal{R}_4 .

For any vector \mathbf{Q} , define Q_i as the i^{th} component of \mathbf{Q} , and define \mathbf{Q}_{-i} as the original vector \mathbf{Q} , but with the i^{th} component Q_i replaced by a zero (so that $\mathbf{Q}_{-i} = \mathbf{Q} - \mathbf{Q}_i$). When using \mathcal{R}_4 , the consumer calculates $U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t + \mathbf{X}_{-i}^t)$ for each $i \in [1, N]$, and uses the resulting value as the measure of the incremental utility gained by purchasing the quantity X_i^t of good i , conditional on also jointly purchasing the rest of the items in vector \mathbf{X}^t . Define this value as $\Delta_{4i}(\mathbf{X}^t, \mathbf{O}^t)$, or simply $\Delta_{4i}(\mathbf{X}^t)$ as a notional shortcut, so that

$$\Delta_{4i}(\mathbf{X}^t) \equiv U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t + \mathbf{X}_{-i}^t)$$

(Note that $\Delta_{4i}(\mathbf{X}^t)$ and $\Delta_{3i}(\mathbf{X}^t)$ will be equivalent if the utility function is additively separable.)

$\Delta_{4i}(\mathbf{X}^t)/X_i^t$ can be thought of as the *conditional marginal utility* (CMU) of good i , given that the consumer already has consumed bundle \mathbf{O}^t , and is going to purchase all of the other components which make up vector \mathbf{X}^t . If we contemplate the continuous version of this idea by considering the limit of $\Delta_{4i}(\mathbf{X}^t)/X_i^t$ as $\|\mathbf{X}^t\| \rightarrow 0$, we see that the conditional marginal utility is equivalent to the left-hand partial derivative of the utility function with respect to good i , evaluated at $(\mathbf{O}^t + \mathbf{X}^t)$.¹⁰

The partial derivative of the utility function, with respect to good i , evaluated at $(\mathbf{O}^t + \mathbf{X}^t)$ is defined as:

$$\partial U / \partial X_i = \lim_{h \rightarrow 0} [U(\mathbf{O}^t + \mathbf{X}^t + h \cdot \mathbf{X}_i^t / X_i^t) - U(\mathbf{O}^t + \mathbf{X}^t)] / h$$

¹⁰ There might be some technical bugs yet to work out with this statement, and the lines which follow it. It is probably necessary to be more precise about whether we are taking $\lim_{\|\mathbf{X}^t\| \rightarrow 0}$ or $\lim_{X_{it} \rightarrow 0}$. These are not necessarily the same thing.

The left-hand partial derivative is the left-hand limit in the above equation, which also implies that h is negative. Therefore, we have:

$$(\partial U / \partial X_i)^- = \lim_{h \rightarrow 0^-} [U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t + \mathbf{X}^t - |h| \cdot \mathbf{X}_i^t / X_i^t)] / |h|$$

which is equivalent to:

$$\lim_{X_i^t \rightarrow 0} \Delta_{4i}(\mathbf{X}^t) / X_i^t$$

The motivation for using the left-hand partial derivative as part of our performance test for partial bundles is as follows. First of all, in cases where the utility function is differentiable at \mathbf{O}^t , the left-hand partial derivative will of course be equal to the full partial derivative, by definition. That is, if the partial derivative is defined, then it will always have the same value as the left-hand partial derivative. In cases where the overall derivative is undefined, such as in cases where marginal utility has a jump discontinuity, it may still be possible for the left-hand partial derivative to be perfectly well-defined. We may therefore capture relevant information given by the utility function that we would be unable to utilize if we restricted our attention to the partial derivative. We can appeal to the intuitive structure of the sequential satisficing process in order to consider why the left-hand partial derivative contains the “correct” information for the consumer’s performance test. We recognize that the consumer needs to compare total levels of utility both with and without the incremental unit of good i (in order to determine whether or not that unit of good i is “worth it”). However, if there are

interactions amongst distinct goods in the utility function, then the utility comparison ought to be able to account for these interactions. Basically, the satisficing process will accomplish this in the same way that the optimizing process does: by evaluating bundles of goods rather than *only* considering the one-dimensional marginal effect of altering a single quantity at a time. In fact, the conditional marginal utility concept, in the context of the sequential satisficing process, will actually do both. We will evaluate the marginal effect of altering a single quantity, conditional on larger bundle being purchased at the same time. In this way, the consumer will be able to allow for utility to be derived through the simultaneous purchase of complementary goods, but will also be able to separate out the individual contribution of each component of the marginal bundle. That is, conditional marginal utility will allow the consumer to avoid buying too many of any single complement (i.e. avoid the cross-subsidization of “too many” units of one good by the surplus utility of the bundle as a whole). Consider the following example:

Assume the utility function is $U(x,y) = \min(x,y)$.

The partial derivative of the utility function with respect to good x is 1 if $x < y$, is 0 if $x > y$, and is undefined if $x = y$.

However, the *left-hand* partial derivative for good x is defined everywhere. It is 1 if $x \leq y$, and it is 0 if $x > y$.

Also, we argue that the left hand partial derivative allows a more valid comparison of marginal utility. If the consumer already has bundle $\mathbf{O}^t = (5,5)$, and is considering purchasing the candidate consumption bundle $\mathbf{X}^t = (1,1)$, what is the appropriate way to measure the marginal utility of the additional unit of x and y ? As we

would normally apply the idea of marginal utility, we would find the marginal utility of x by finding $[U(\mathbf{O}^t + \mathbf{X}_i^t) - U(\mathbf{O}^t)] / X_i^t$, or $[U(6,5) - U(5,5)] / 1$. This is, of course, zero for this utility function. However, an alternative calculation would be the *conditional* marginal utility of x . By “conditional,” we mean that marginal utility of x , conditional on \mathbf{X}^t . To perform this calculation, we find $[U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t + \mathbf{X}_{-i}^t)] / X_i^t$, or $[U(6,6) - U(5,6)] / 1$. In the latter calculation, we find that the conditional marginal utility of x is 1. In this way, CMU can give us a more valid measure of the marginal utility of complementary goods in a bundle¹¹.

Furthermore, CMU gives us a mechanism for ensuring that the \mathbf{R} hypothesis is satisfied in situations where complementary goods are purchased. Again, as above, assume that $U(x,y) = \min(x,y)$. Further assume that the price vector is $(\frac{1}{2}, \frac{1}{2})$ and that $\mu = \frac{1}{2}$. For a candidate consumption vector $\mathbf{X}^t = (2,1)$, the total cost of the bundle is 1.5, the incremental utility of the bundle is 1. Therefore \mathbf{X}^t would satisfy \mathbf{R}_1 , since $1 \geq (\frac{1}{2}, \frac{1}{2}) \cdot (2,1) * \mu = 0.75$, and $\mathbf{X}^t = \mathbf{X}^t$. This is true despite the fact that the 2nd unit of x in \mathbf{X}_i^t obviously yields zero additional utility, but still costs a positive amount of wealth. Again, this demonstrates that \mathbf{R}_1 violates the \mathbf{R} hypothesis for this utility function.

However, because of the complementarity between the two goods, neither component of \mathbf{X}^t will pass the performance test conducted by \mathbf{R}_3 . Since positive marginal utility requires additional units of *both* x and y , $\Delta_{2i}(\mathbf{X}^t)$ will be zero for both $i = x$ and $i = y$. Therefore, \mathbf{R}_3 violates the \mathbf{R} hypothesis for the opposite reason that \mathbf{R}_1 did; it ignores the opportunity to purchase a bundle which offers strictly greater additional utility than its estimated opportunity cost. Relative to what we might think of as

¹¹ Though, this also introduces a sort of “double-counting” problem, which is addressed below.

“reasonable” consumer behavior, \mathcal{R}_1 buys too much, and \mathcal{R}_3 buys too little, in this example.

\mathcal{R}_4 can begin to address both types of \mathcal{R} hypothesis violations, because it employs CMU in its performance test. First of all, \mathcal{R}_4 avoids the pitfall of \mathcal{R}_1 , because it will not allow for the purchase of the 2nd, useless unit of x in \mathbf{X}^t . Specifically, the left-hand partial derivative of the utility function, with respect to x , evaluated at $(\mathbf{O}^t + \mathbf{X}^t)$ will be zero on the interval $(1, 2]$, and will be 1 on the interval $[0, 1]$, given that $X_y^t = 1$. When \mathcal{R}_4 is applied (under the assumption of incrementalism, described in the $\mathcal{R}_2/\mathcal{R}_3$ section), it will reject any amount of x which is strictly greater than X_y^t , and so it is able to avoid the pitfall of \mathcal{R}_1 . Furthermore, once the “excess” of good x has been trimmed, \mathcal{R}_4 will be able to consider the remaining bundle $(1, 1)$. It will accept the remaining bundle in its entirety, because $\Delta_{4i}(\mathbf{X}^t)$ is high enough in $(1, 1)$ for both $i = x$ and $i = y$, and thus avoid the pitfall of \mathcal{R}_3 .

Notice, despite the fact that we are employing a left-hand derivative (or its discrete analog), this concept of marginal utility is not necessarily backward-looking. Since each iteration of the satisficing process begins at some \mathbf{O}^t , and then considers the effect of adding a partial vector \mathbf{X}^t , conditional marginal utility does measure the change in utility that results from the addition of further consumption. Specifically, the comparison involves the difference in utility level caused by an increase in good i consumption from O_i^t to $(O_i^t + X_i^t)$. Loosely, the process is forward-looking from the perspective of \mathbf{O}^t and backward looking from the perspective of $(\mathbf{O}^t + \mathbf{X}^t)$.

\mathcal{R}_4 still introduces its own worries with respect to the \mathcal{R} hypothesis, however. Essentially, by looking at conditional marginal utility, \mathcal{R}_4 “double counts” the marginal

utility offered by a partial bundle which contains complementary goods. Using the same utility function as above, $U(x,y) = \min(x,y)$, we can easily provide an example where \mathcal{R}_4 violates the \mathcal{R} hypothesis:

$$\mathbf{O}^t = \mathbf{0}$$

$$\mathbf{P} = (1,1)$$

$$\mathbf{X}^t = (1,1)$$

$$\mu = 1$$

In this example, $\Delta_{4i}(\mathbf{X}^t) = 1$ for both $i = x$ and $i = y$. Therefore, both elements of \mathbf{X}^t will pass \mathcal{R}_4 's performance test, and $\mathbf{X}^t = \mathbf{X}^{\wedge t} = (1,1)$. However, in this example, \mathbf{X}^t as a whole would not have even passed \mathcal{R}_1 . The total marginal utility gained by $\mathbf{X}^t = \mathbf{X}^{\wedge t}$ is 1, while the total marginal cost of this partial bundle is 2. If $\mu = 1$, then the bundle as a whole provides strictly less marginal utility than its estimated opportunity cost. Therefore, its purchase represents a violation of the \mathcal{R} hypothesis. This is due to the fact that the concept of conditional marginal utility assigns the full amount of added satisfaction derived from any partial bundle to *both* goods x and y .

Therefore, in order to formulate a rule which is consistent with the \mathcal{R} hypothesis, we need to define a still more sophisticated satisficing rule, \mathcal{R}_5 . This rule will be identical to \mathcal{R}_4 , except that the consumer will be aware of any relevant utility interactions between goods, and will also apply the performance test to all combinations of goods for which there exists such an interaction. When the consumer considers purchasing combinations of items which are related in the utility function, she will first test each quantity of each item individually, as in \mathcal{R}_4 . But she will then further test the combined quantities of those

goods, as if they were jointly considered a single good. In the previous example, \mathbf{R}_5 would, just like \mathbf{R}_4 , first pare down the original \mathbf{X}^t vector from (2, 1) to (1, 1), since the CMU of good x is zero if $x > y$. \mathbf{R}_5 's individual component test approves both elements of (1, 1). However, \mathbf{R}_5 would then further test (1, 1) by evaluating incremental utility along the direction of the vector (1, 1). This is found by calculating $U(1, 1) - U(0, 0) = 1$ and comparing it to the total cost $(P_x, P_y) \cdot (1, 1) = 2$. Since the estimated opportunity cost (defined as $\mu \cdot \text{total cost}$) exceeds the marginal utility, \mathbf{R}_5 will reject the entire \mathbf{X}^t bundle, thus avoiding any violation of the \mathbf{R} hypothesis in this case.

Also, there is a similar concern with \mathbf{R}_4 with respect to the situation where the utility function includes goods which are perfect substitutes. Assume there are two goods I and J which are perfect substitutes. The utility function is such that the consumer derives 1 util of utility, so long as she consumes at least one unit total of either good, but receives no additional utility from further consumption beyond one unit.

That is:

$$U(x, y) = \begin{cases} 0 & \text{if } (x + y) < 1 \\ 1 & \text{if } (x + y) \geq 1 \end{cases}$$

If $\mathbf{O}^t = (0, 0)$ and $\mathbf{X}^t = (1, 1)$, $\mu = 0.5$, and $\mathbf{P} = (1, 1)$, then if we apply \mathbf{R}_4 we will get $\mathbf{X}^t = \mathbf{0}$. The CMU of both x and y is zero, since $U(\mathbf{O}^t + \mathbf{X}_{-i}^t) = U(\mathbf{O}^t + \mathbf{X}^t) = 1$ for both $i = x$ and for $i = y$. However, by purchasing either (1, 0) or (0, 1), the consumer would increase her total utility by an amount greater than μ times the cost of her bundle. Therefore, failure to make one and only one of these two purchases constitutes a violation of the \mathbf{R} hypothesis. As with the issues related to the magnitude of the \mathbf{X}^t vectors,

mentioned in relation to \mathbf{R}_2 , the concerns about the under-consumption of perfect substitutes will not be relevant if we restrict our attention to only the continuous case of the satisficing process. However, if we wish to address these concerns in our formulation of the discrete version of this process, we need to also define \mathbf{R}_5 appropriately.

Once again, the following correction is not necessary in the continuous version of this model, but if we wish for the discrete version to be well-defined and consistent with the \mathbf{R} hypothesis, we must also define \mathbf{R}_5 and $\Delta_{5i}(\mathbf{X}^t)$ such that:

$$\Delta_{5i}(\mathbf{X}^t) \equiv U(\mathbf{O}^t + \mathbf{X}^t) - U(\mathbf{O}^t + \mathbf{X}^{A_{-i}t})$$

where $\mathbf{X}^{A_{-i}t}$ is defined such that its j^{th} component is

$$\left\{ \begin{array}{l} \mathbf{X}_j^{A_{-i}t} = \mathbf{X}_j^t, \text{ if } j > i \\ \mathbf{X}_j^{A_{-i}t} = 0, \text{ if } j = i \\ \mathbf{X}_j^{A_{-i}t} = \mathbf{X}_j^t, \text{ if } j < i \end{array} \right.$$

$\mathbf{R}_5(\mathbf{X}^t)$ is applied, in ascending order from $i = 1$ to $i = N$ in the construction of $\mathbf{X}^{A_{-i}t}$.

An Intuitive Sketch of the Satisficing Process Under \mathbf{R}_5

Imagine that a consumer makes periodic trips to the supermarket throughout the month. At the beginning of each trip, there is an accumulated bundle of goods that she has already purchased prior to making the current trip (This bundle may be the zero vector, if the current trip is her first trip of the month). Call this accumulated bundle \mathbf{O}^t . Based, in part, on the composition of \mathbf{O}^t , the consumer will select a basket of goods to consider purchasing during her current trip. She will walk through the aisles, and select

various quantities of various goods, and place them in her shopping cart. Call the function which specifies which quantities of which goods are placed into the cart $\Pi(\mathbf{O}^t)$, and call the bundle of goods which ends up in the shopping cart \mathbf{X}^t . The bundle in the shopping cart consists of the goods that she will consider purchasing during her current trip. Once she has selected her candidate basket of goods, the consumer will ask herself if she is *sure* that she wants to buy each of these items. She must make a final decision about whether she actually wishes to purchase all of the items in her cart. In order to make this decision, she will conduct a performance test for each item in her cart, and as necessary, for combinations of items. One by one, she picks up items from her basket, and holds them in her hand. She is then able to perfectly assess the conditional marginal utility of each of these items. For each item, she knows how much additional utility she would derive by purchasing *everything else* in her basket, and she knows by how much her utility would increase over that amount if she were to also the item in her hand. She knows the conditional marginal utility of each item, and she knows the price of each item. If the conditional marginal utility is at least as great as its price multiplied by μ (recall that $\text{Price} \cdot \text{Quantity} \cdot \mu$ is her estimate of the opportunity cost of any purchase), then she will decide to actually purchase the item in her hand. If the conditional marginal utility is less than price times μ , then she will decide not to purchase the item in her hand, and will remove it from her cart and place it back on the shelf. Assume that she conducts this performance test for the smallest feasible unit of each type of good in her basket, so that she may alter the quantity of any good in her basket, by deciding to put some of any particular good back on the shelf while also keeping some in her cart. Once she has completed this process for each individual good, she repeats the process for combinations

of goods that have potential interactions between them in the utility function. That is, if her basket is initially full of hot dogs and hot dog buns, she will first decide whether she has too many hot dogs (individually) or too many buns (individually), relative to the amount of the other complement present in her basket. She will remove any of either which do not pass the performance test. Next, she will conduct a *joint* performance test of the combination of both complements. If the number of hot dogs and hot dogs buns provides a jointly sufficient amount of marginal utility, relative to the total price of both goods, then she will purchase the remaining contents of her basket. If the combination of hot dogs and buns does not pass the joint performance test, then she will remove dog-bun quantities until the joint performance test is satisfied¹².

Finally, note that it is not necessary for the above description to be a *literal* description of the behavior of a typical consumer. In fact, it is probably more valid to imagine that the “basket” is a virtual basket, and the comparison of conditional utility takes place in the consumer’s mind, rather than the consumer physically placing all potential items into a basket, and physically taking them out one by one to evaluate their conditional marginal utility.

In this way, we can allow our satisficing rule \mathfrak{R}_5 to make a valid judgment in cases where there are interactions in the utility function among different goods (complementarities, substitutabilities).

¹² Possible further adjustments need to be made if the complementarity between the goods is non-linear. For now, we will not address these types of concerns. Our basic approach is to require that, whatever the satisficing rule is that the consumer employs, it must be consistent with the \mathfrak{R} hypothesis. The details and mechanics of constructing such a rule, relative to a certain utility function, may be left unaddressed until that particular problem is modeled. In other words, the precise construction of an \mathfrak{R} hypothesis-consistent rule for *every* conceivable utility function need not be conducted *now*. We can cross some bridges as we come to them.

Additional Comments on the \mathbf{R} Function

Once again, our purpose in the preceding discussion was not necessarily to derive a single, universal decision rule which is applicable in every conceivable situation. There probably do exist certain reasonable utility functions for which it can be easily shown that \mathbf{R}_5 potentially violates the \mathbf{R} hypothesis. Our intent was to highlight, through examples, some of the issues which need to be considered when attempting to derive a specific decision rule for use in a specific example. Going forward, we will assume that, whatever utility function we are employing in any particular example, all we need to do in order to analyze the problem is to find a specific \mathbf{R} function which satisfies the \mathbf{R} hypothesis for that particular utility function. Through the remainder of this paper, we will use \mathbf{R}_5 , unless otherwise noted. We will keep in mind, though, that other examples not discussed here may require a different rule in order for the resulting satisficing strategy to be consistent with the \mathbf{R} hypothesis. Going forward in this paper, we will drop the subscript on \mathbf{R} , and will adopt the convention that, if we are using the just the unmodified symbol “ \mathbf{R} ” to denote our decision rule, then we are implicitly claiming that said decision rule satisfies the \mathbf{R} hypothesis for the problem in question.

Finally, we may point out that, in principle, the idea of the satisficing rule of thumb actually nest the fully optimizing demand strategy, as a special case. If we define a particular rule \mathbf{R}^{FR} , such that the rule itself is a function of all of the parameters of the consumer Utility Maximization Problem (prices, wealth, preferences), we can clearly describe a sequential satisficing strategy that will necessarily be strategically equivalent to the true optimizing demand strategy for any path-generating process in which all \mathbf{X}^t vectors are strictly > 0 . Specifically, if \mathbf{R}^{FR} identifies the set of most-preferred affordable

consumption bundles, and then simply approves all components of any \mathbf{X}^t vector which leave at least one element of the most-preferred set still attainable (in the “ $A/\sim A$ ” sense, described above), and disallows any components which would lead to no elements of the most-preferred set still attainable, then we would necessarily have a satisficing rule¹³ and strategy which are indistinguishable from the true optimizing strategy. Of course, this situation would not really capture the essence of the term “satisficing.” We will assume, as we have all along, that consumers lack the computational capacity to employ such a sophisticated satisficing rule as \mathbf{R}^{FR} .

The Path-Generating Process Π

In order to describe the sequential consumption process, we need some construct in place to describe the order in which partial consumption bundles are evaluated by the decision rule \mathbf{R} . That is, conditional on the consumer already having chosen to consume some bundle \mathbf{O}^t , we need to specify what the next sequential candidate bundle \mathbf{X}^t will be.

In other words, we need to define the path-generating process Π .

We can do so by conceiving of the consumption space as a vector field. That is, to every point in the consumption space, we associate a corresponding vector. The associated vector will be the candidate consumption vector \mathbf{X}^t for that particular point in the consumption space \mathbf{O}^t . As we defined previously, this means that $\Pi(\mathbf{O}^t) = \mathbf{X}^t$. Under the assumption of the irreversibility of consumption, each of the associated vectors must be $\geq \mathbf{0}$ (implying also that they all be $\neq \mathbf{0}$, as stated in the initial description of the sequential satisficing process). There will be an analogous field defined below, $\Pi^{\setminus}(\mathbf{O}^t) =$

¹³ Including an arbitrary tie-breaking rule, if necessary.

\mathbf{X}^t , which describes the *actual* consumption path which results from the application of \mathcal{R} to \mathbf{O}^t and $\Pi(\mathbf{O}^t)$. Again, as described previously, the vectors in the Π , the \mathbf{X}^t s, may be = $\mathbf{0}$, but must be non-negative in every component, so that $\mathbf{X}^t \geq \mathbf{0}$ for all t .

We may think of the Π field as being a preference function, of sorts. Given that the sequential consumption process has already arrived at point \mathbf{O}^t , the consumer would like for consumption to continue in the direction specified by $\Pi(\mathbf{O}^t)$. Like a standard utility function over bundles of goods, the Π field itself may be, in principle, arbitrarily complex, ill-behaved, or mischievous. But in practice, we will generally restrict our analysis to certain well-behaved forms. Our hope is to use a few tractable examples to begin to draw general conclusions and find patterns in how this Π field construct impacts the consumption process.

Once we have defined the Π field, we can then conduct our entire demand analysis using only the information contained in three separate functions of the same domain (namely, the consumption space). That is, the consumption space contains three distinct layers of information, and each layer plays a fundamental role in the evaluation of the sequential satisficing process. The three layers are:

- 1) The cost function: for any point \mathbf{O} , the monetary cost of purchasing the consumption bundle represented by \mathbf{O} is very simply given by the inner product of the consumption vector with the price vector \mathbf{P} . (Cost = $\mathbf{O} \cdot \mathbf{P}$)
- 2) Preferences: The utility function (and all of its relevant derivatives)
- 3) The Π field: The vector field, $\Pi(\mathbf{O}^t)$, which gives the associated candidate consumption vector, \mathbf{X}^t , for each point in the consumption space. This function

tells us the order in which sequential candidate bundles will be evaluated by the chosen decision rule \mathbf{R} . The Π field will, in conjunction with \mathbf{R} , determine the evolution of the actual consumption path, from the origin to the final consumption bundle \mathbf{O}^T .

The first two layers of information are standard and familiar, while the third is novel. Together, these three layers of information make the satisficing process possible. The third specifies what bundles will be assessed by the satisficing rule, and in what order, while the first two allow for a comparison of the (estimated) costs and (true) benefits of those bundles. Along with a given price vector \mathbf{P} and initial condition that consumption begins at the origin, this vector field Π defines candidate consumption as a simple dynamical system in wealth¹⁴. Later, we will incorporate the decision rule \mathbf{R} , and will define a related vector field $\Pi^*(\mathbf{O}^t)$, which will define the *actual* consumption path as a simple dynamical system.


Consider the Π vector field in isolation. Starting at any point \mathbf{O} in the consumption space, Π defines a flow trajectory which uses point \mathbf{O} as its initial condition (fig. 4). This process may be either discrete (in which case the magnitude of the associated vectors is important) or continuous (in which case the magnitude is not important, for our purposes). We shall focus primarily on the continuous case in most of what follows, and will accordingly ignore the magnitude of all associated \mathbf{X}^t vectors. This system will exhibit no momentum in the consumption path, and the algebraic description of the consumption path can usually be found by solving a simple differential equation.

¹⁴ It is possible that we are not using the precisely correct terminology here. What we mean is that the “time” dimension may be interpreted as the consumer’s remaining wealth.


For what follows, we will generally assume that the most important initial condition is the origin; the consumer starts with nothing, and must purchase everything she needs in each period. Since our goal is ultimately to derive demand, this assumption is appropriate. Even if we wished to use this construct to describe a situation in which the consumer has some initial endowment (or previously accumulated stock) of consumption goods at period zero of our model, we should still conceive of the consumption path in the current period starting at the origin, and then impose the appropriate translation or shift of the utility function in order to properly reflect the marginal utilities of subsequent purchases beyond the initial stock¹⁵.

Since we will be primarily concerned with the flow given by Π which begins at the origin, we shall refer to this flow as the *principal path* or *original path* of the vector field. The principal path will be an important construct in determining the actual demand functions. A simple interpretation of the principal path is that it describes the physical history of incremental purchases which would result if the consumer's behavior was not bound by any constraints (especially the worth constraint). That is, if there were no constraints on the consumer's purchase behavior, then physical consumption would evolve simply by following along the principal path through the consumption space, ad infinitum. In practice, the vector field provides the physical path that consumption will follow, as well as a complete set of rules by which the actual consumption path will be

¹⁵ As an example, assume that $U(x) = 1/x$, and that the consumer has an initial endowment of 1 unit of good x . If we wish to describe *demand*, or what the consumer will purchase in the current period, we should not begin the consumption path at a quantity of 1, since that first unit is not actually purchased in the current period. Instead, we should still require that the consumption path begin at the origin, but shift the utility function to $U(x) = 1/(x+1)$.



For any
vector
field...



There is a
particular
trajectory
defined by
that each
point in the
field.

Figure 4. Vector Fields and Consumption Trajectories.

updated/determined should the principal path be constrained or altered by either the budget hyperplane or by the satisficing rule \mathfrak{R} .

Consider two simple fields, which are chosen primarily for their simplicity and tractability:

The “identity field,” (fig. 5) in which to each n-tuple \mathbf{Q} in \mathbb{R}^N_+ , we associate the same vector \mathbf{Q} .




Figure 5. Identity Field (\mathbb{R}^2).

The “inverse identity field” (fig. 6) in \mathbb{R}^2_+ , in which to each pair (x,y) in \mathbb{R}^2_+ , we associate the vector (y,x) .




Figure 6. Inverse Identity Field (\mathbb{R}^2).

Since the domain of each of these fields is consumption space, or the (closure of the) positive orthant, both of these fields generically satisfy the irreversible consumption requirement that each partial consumption vector \mathbf{X}^t be $\geq \mathbf{0}$. Despite the simplicity of these two fields, they each exhibit the troublesome feature that the vector associated with the origin is the zero vector, $\mathbf{0}$. Note that, not only does this imply that the principal path for each of these fields is the degenerate case of a single point (the origin itself), but these fields do not satisfy the assumptions of the satisficing process itself. It is required, by construction, that each associated candidate consumption vector \mathbf{X}^t be non-negative in every component, and also *not equal to zero*. We must overcome this problem by assumption, if we are to employ either of these simple fields as a path-generating process. For everything that follows, we will establish a convention that any field Π for which it is

the case that $\Pi(\mathbf{0}) = \mathbf{0}$, we will implicitly assume that we replace the zero vector with some other vector which is $\geq \mathbf{0}$. For simplicity, we will use the δ vector in place of the zero vector $\mathbf{0}$ at the origin (if δ is some arbitrarily small but positive number, then the δ vector is defined as the vector whose every element is δ). So, for example, when we refer to the “identity field” henceforth, we will actually be referring to the field defined by $\Pi(\mathbf{Q}) = \mathbf{Q}$, if $\mathbf{Q} \geq \mathbf{0}$, and $\Pi(\mathbf{Q}) = \delta$, if $\mathbf{Q} = \mathbf{0}$. Similarly for the “inverse identity field” in \mathbb{R}^2 . Once we have made this “ δ modification,” both of these fields (in 2-space) will have the identical principal path: the line $y = x$.

Beginning to Outline the Geometry of the Satisficing Demand Process

Recall that we have previously defined \mathcal{R}_0 as the completely non-discriminating satisficing rule:

$$\mathcal{R}_0(\mathbf{X}^t, \mathbf{O}^t) \equiv \mathbf{X}^t = \mathbf{X}^t, \forall (\mathbf{X}^t, \mathbf{O}^t)$$

If we assume for the moment that the satisficing rule employed is the completely non-discriminating rule \mathcal{R}_0 , so that the rule fully approves any and all candidate consumption vectors which are brought under its consideration, then the actual consumption path will simply follow along the principal path of Π until the consumer exhausts all of her available wealth. In other words, if \mathcal{R} is non-discriminating, then the final consumption bundle selected by the satisficing process will simply be the point of intersection of the primary path and the budget hyperplane (fig. 7).

Qualitatively, this is precisely the geometry that will allow us to derive the demand curve for a satisficing consumer: the Π field specifies a principal path, and actual consumption begins at the origin and follows that path, until it is no longer able to do so due to the presence of some binding constraint. In this particular example, the only such constraint is the budget constraint (again, due to the non-discriminating nature of \mathcal{R}_0). Below, we will introduce other constraints which are the result of more restrictive satisficing rules beyond \mathcal{R}_0 . For now, however, we will make some observations on the role played by the Π field in selecting the final consumption bundle \mathbf{O}^T . Most importantly, it is obvious that the nature of the Π field itself will play a role in the selection of the final consumption bundle.

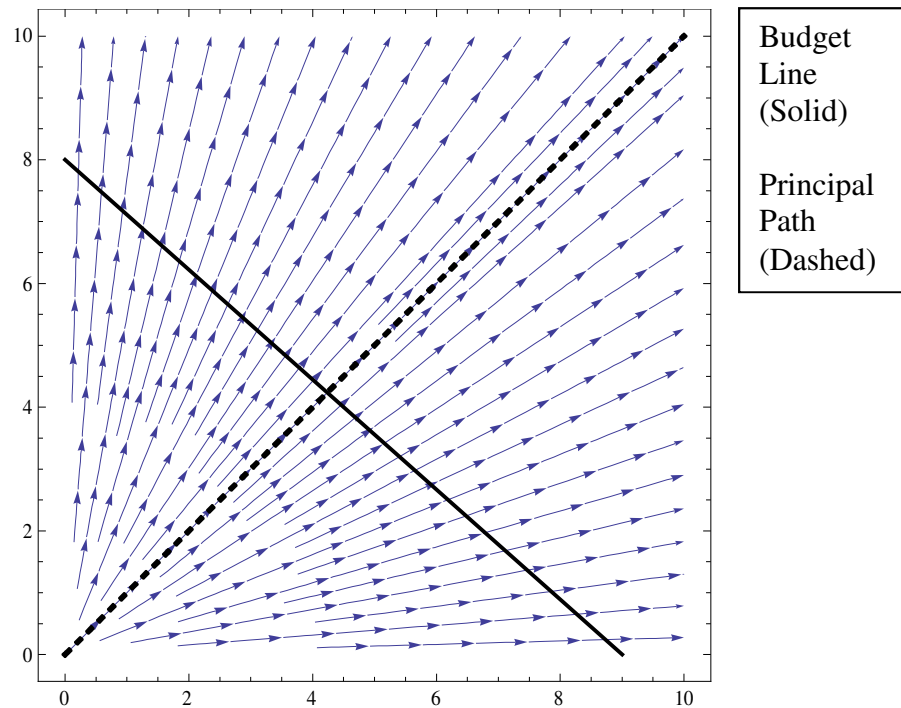


Figure 7. Budget Constraint and the Consumption Path.

Clearly, under \mathcal{R}_0 , changing altering the principal path of the vector field will alter the realized final consumption bundle. In fact, in this particular case, all that is necessary to achieve a different final bundle is to replace the vector which is associated to the origin. In 2–space specifically, we can replace the δ vector (δ, δ) with a more general definition of the associated vector $(\alpha*\delta, \delta)$, where α is any number on $[0, \infty)$, as our defined value of $\Pi(\mathbf{0})$. The δ vector, as previously defined, obviously corresponds to the case where $\alpha = 1$, but we can select *any* bundle¹⁶ on the budget hyperplane by altering the value of α (fig. 8).

Obviously, it is possible for the Π field itself to discriminate amongst consumption bundles, even if the decision rule does not do so in any substantive way. In fact, it is necessary that this be the case; even under \mathcal{R}_0 , *some* final bundle is going to be chosen by the satisficing process. If \mathcal{R}_0 does play a role in deciding which final bundle is selected, then the Π field must do so by default. Like the decision rule, the Π process can in principle be “tuned” to reflect varying levels of sophistication on the part of the consumer. That is, we can imagine very sophisticated Π fields, which utilize a great deal of available information, and lead to a final consumption bundle which is either truly optimal, or at least very close to optimal a great deal of the time. We can also imagine very simple Π fields, which reflect a low level of rationality on the part of the consumers, and which do not do a good job, by themselves, of ensuring a close-to-optimal final bundle. We can imagine a continuum between the two extremes.

¹⁶ Clearly, to include both endpoints of the budget line, it is also necessary to allow for altering the component in which α appears, as in $(\delta, \alpha*\delta)$ in place of $(\alpha*\delta, \delta)$.

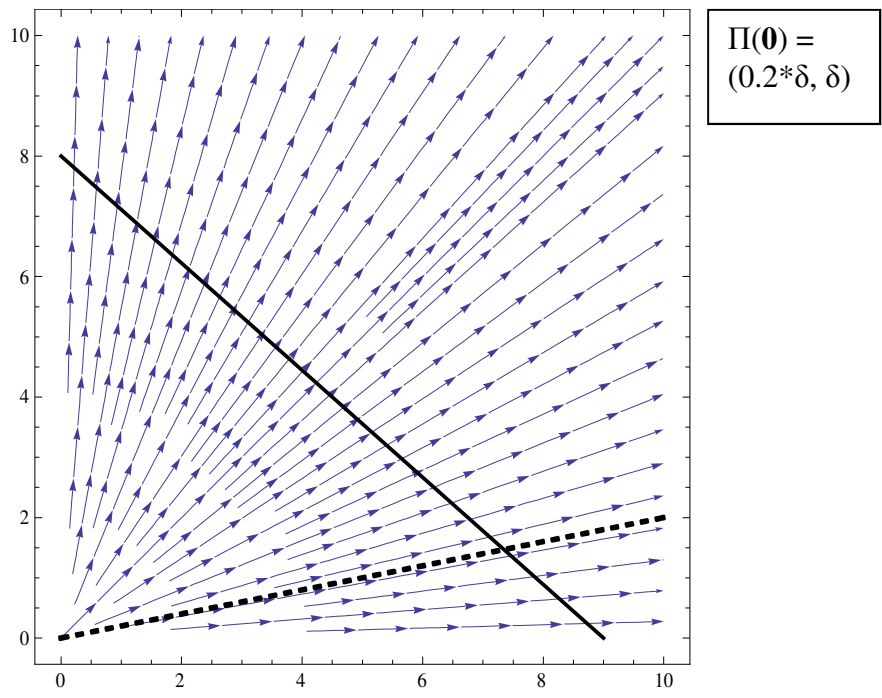



Figure 8. Varying the Vector Associated to the Origin.

Also like the decision rule, it is possible to construct a Π field in such a way as to guarantee *strategic equivalence* between the satisficing and optimizing strategies (We can refer to this “fully rational” field as Π^{FR}). As a result, there is a second way in which the satisficing process actually nests the optimizing process as a special case. Basically, the process described above, in the section entitled “The Optimizing Consumption Strategy as a Sequential Process,” is an example of a path generating process which is necessarily strategically equivalent to the optimizing strategy. Recall that this process basically involves first identifying the set of optimal affordable bundles, and then selecting any one of many possible (non-decreasing) paths from the origin to any member of that set. Of course, in order for this strategic equivalence to hold, it is a necessary condition that the Π field itself be a function of the price vector and the utility function. That is, if the Π field is to guarantee a truly optimal consumption bundle under \mathbf{R}_0 , for all possible price vectors, than the field itself must change as the parameters of the consumer budget problem (prices, wealth, preferences) change. Moreover, the field must always change in a precisely-chosen way which guarantees that the principal path will intersect with the set of optimal bundles at the budget hyperplane. In other words, $\Pi^{\text{FR}}(\mathbf{O}^t)$ exists, but it is actually more accurately written as $\Pi^{\text{FR}}(\mathbf{O}^t, \mathbf{P}, U(\mathbf{O}^t), w)$.

Recall that a central hypothesis of this paper is that these informational requirements, which must be satisfied in order to allow for such a sophisticated field as Π^{FR} to be employed by the consumer, are too unreasonably unwieldy to be thought of as describing the behavior of a realistic consumer. Accordingly, despite the fact that the basic setup of the satisficing can be thought of as nesting the optimizing strategy, true optimization does not truly reflect the spirit of what we mean by the term “satisficing.”

We will assume that the consumer does not employ either \mathbf{R}^{FR} or Π^{FR} as a component of a satisficing strategy.

On the other extreme of the “rationality spectrum” of the set of all possible Π fields is what we will define as a *naïve field*. A field will be said to be naïve if and only if it is not a function of the price vector or of the utility function itself. That is, if a field changes the mapping from the consumption space to the set of \mathbf{X}^t vectors, as the set of optimal bundles changes (which is, in turn, the result of a change in \mathbf{P} , in wealth, or in the utility function), then that field will not be said to be naïve. Consequently, under \mathbf{R}_0 , a naïve field *might* lead to a satisficing demand process which is narrowly equivalent to the optimizing process, for some specific parameters of the consumer budget problem. But a naïve field cannot, in general, be strategically equivalent to the optimizing process, under \mathbf{R}_0 .

In general, what we intend to describe when we use the term “satisficing strategy” will be a combination of an imperfectly rational decision rule (a rule which is not \mathbf{R}^{FR} , but is instead based upon an estimate of the opportunity cost of any purchase), in combination with a naïve vector field Π .

We can use this type of strategy, as defined, to derive consumer demand. To do so, we will use the Π field, and then we overlay an analytic locus which describes a constraint on the consumption behavior. In the preceding example, the budget line represents the only constraint on consumption behavior. Using the same domain (the consumption space), we have superimposed objects of two types: the vector field and the budget hyperplane. The interaction between these two objects will, in part, determine the consumer’s final demand bundle.

We will use a similar structure to complete the description of the geometry of the satisficing process. We will continue with the practice of starting with the Π field, and then overlaying constraints upon the field and its principal path. Here, we will consider another type of constraint, which we will call the *worth constraint*. The worth constraint is a physical manifestation of the satisficing rule \mathfrak{R} , and represents the locus of all points for which the conditional marginal utility of any partial purchase is just equal to its price times the benchmark “utility index” value μ . Though there is but a single satisficing rule employed by this process, there will be at least as many distinct worth constraints as there are distinct goods available for consumption (and strictly more worth constraints than goods, if there are any complementarities among those goods). In other words, each worth constraint represents the application of the satisficing rule to one specific margin of consumption. For points in the consumption space at or outside (relative to the origin) the locus of points describing any particular worth constraint, actual consumption may not continue in the direction specified by that particular consumption margin. This is demonstrated below.

By overlaying a system of worth constraints on the consumption space, we can use the Π field as well as the constraints to derive an *actual consumption path (ACP)*. In general, the ACP will be a piecewise combination of the Π field’s principal path and an algebraic description of one or more of the worth constraints. It can equivalently be thought of as the principal path of the *modified* vector field, $\Pi^t(\mathbf{O}^t) = \mathbf{X}^t$.

The Worth Constraints

The shape and location of the worth constraints will depend upon the satisficing rule \mathcal{R} (including the value of μ), the utility function, and the price vector. The price vector and μ will have a quantitative effect, primarily determining the position of the constraints (changes to the price vector or to μ will shift the location of one or more constraints). The utility function has both a quantitative effect, as well as a qualitative effect on the nature of the worth constraints; it determines the shape of the constraints, as well as the position. Each worth constraint will coincide with one of the level curves of a particular partial derivative of the utility function. Each worth constraint describes the frontier at which the decision rule just begins to censor specific components of any potential candidate consumption vectors.

Three Illustrative Cases

Consider three cases for the utility function. While clearly not an exhaustive list of possibilities, the following three cases can be used to illustrate how the worth constraints will depend upon the utility function.

Case 1: Additive Separability

If the utility function is additively separable, the marginal utility of each good depends only on the quantity of that good consumed. Therefore, the worth constraint for each good will always be a hyperplane normal to the axis measuring that good. In two dimensions, this is easy to illustrate; given P_i and μ , there will be exactly one critical

value for consumption of good i at which the rule \mathcal{R} is just satisfied (with no slack). The worth constraint for good i (again, given parameters given P_i and μ), will be the line through that critical value and perpendicular to the i -axis.

Case 2: Perfect Substitutes

If the utility function is of the form $U(x, y) = u(A \cdot x + B \cdot y)$, then the worth constraints will be linear, and downward sloping. In fact, they will coincide exactly with the indifference curves.

Case 3: Perfect Complements

These constraints will have two parts: a segment connecting the origin with some point (α, α) , where α is some number ≥ 0 , and then a segment beginning at (α, α) and extending indefinitely in the positive direction perpendicular to one of the axes.

Because of the discontinuity in the marginal utility function for perfect complements, the worth constraints in this case can be thought of as being composed of two parts: the “inside” constraint, and the “outside” constraint. The “inside” constraint reflects the fact that each good will have zero marginal benefit if there is an insufficient amount of the other good already being consumed. The “outside” constraint reflects the fact that, in the range of quantities for which altering the amount of one good does affect the value of the utility function, then conditional marginal utility of doing so still must be at least as great as μ times price.

Worth Constraints in Action

The underlying idea here is that the Π field simply indicates the order in which the decision rule \mathcal{R} is applied to partial consumption bundles. The physical manifestation of the decision rule, as visible in the consumption space, is the system of worth constraints. Worth constraints represent the possible ways in which the consumer might value individual goods, or combinations of goods, in any specific partial consumption bundle. At a minimum, there will be one worth constraint for each good available to the consumer. But there will also be an additional worth constraint for each relevant interaction between goods (i.e., if goods X and Y are complements, then there will be a worth constraint for X, another for Y, and a third for X-and-Y jointly).

The worth constraint for “good x” will always coincide with one of the level curves of the partial derivative of the utility function with respect to x. Specifically, it will be the level curve for which marginal utility is equal to μ times the goods price. As mentioned above, these level curves, and therefore the constraints themselves, will have qualitatively different shapes, depending upon the nature of the utility function.

The worth constraints shape the actual consumption path, using the primary path of the Π field as a “starting point.”

For example, consider the identity field in 2-space. If we assume that the utility function is additively separable, and is given by $U(x, y) = \ln x + \ln y$, the price vector for goods (x, y) is $(1, 1)$, and the value of μ employed by \mathcal{R} is $1/5$. Clearly, by construction, this satisficing process will exhibit a binding worth constraint at a quantity of 5 for each good. At any quantity below 5 (for either good), the marginal utility divided by the price will exceed the μ -threshold, and so any partial consumption bundle in that range will pass

the \mathfrak{R} test with respect to that good. Geometrically, we can represent this by plotting the locus of points where this worth constraint is *just* binding, so that $\partial U/\partial x = \mu * P_x$ and $\partial U/\partial y = \mu * P_y$. Shown below is the identity field with just the “y” worth constraint superimposed (fig. 9).




Figure 9. The Worth Constraint.

If we ignore the x worth constraint for the moment, we can envision what the actual consumption process will look like with this primary path and this single worth constraint. Consumption begins at the origin, and follows the primary path ($y = x$). We can conceive of every sequential step of arbitrary smallness as being one iteration of the satisficing process. In the range below the constraint, the satisficing rule approves everything in each candidate consumption bundle, so that candidate consumption and

actual consumption are equivalent over that range. However at the constraint (and above it, if for the sake of argument, it were possible for consumption to ever consumption to enter that range), any possible candidate consumption bundle would *not* pass the \mathcal{R} test with respect to good y , since marginal utility divided by price would be strictly less than μ . Therefore, the satisficing rule would require that any quantity of good y which is present in any candidate bundle X would be removed prior to purchasing the actual partial bundle X^* . Accordingly, actual consumption will only proceed in a direction orthogonal to the y -axis, once we have reached the portion of the consumption space where the y worth constraint is binding. More specifically, in a region where the worth constraint binds for only one good y , any actual consumption vector X^* will be equal to the orthogonal projection of the candidate consumption vector X^t , onto the hyperplane formed by the remaining axes other than the y -axis. It is as if the consumption path “runs into a wall” formed by the worth constraint, and cannot continue in the direction of the wall, but can continue in a modified direction, along the wall.

An alternative way of describing the same law of motion is to assume that the worth constraint actually operates directly on the vector field $\Pi(\mathbf{O}^t)$ itself. That is, for all points in the consumption space which lie beneath the worth constraint (ie, between the constraint and the origin), the vector field remains unchanged. But for all points which lie on or above the worth constraint for good y , the y -component of the associated vector in the Π field is censored, and replaced with a zero. The original, unmodified vector field is $\Pi(\mathbf{O}^t)$. However, the new, modified vector field (which takes into account the consequences of applying the decision rule \mathcal{R} , and of the potentially binding worth

constraints) is not $\Pi(\mathbf{O}^t)$, but is instead $\Pi^{\wedge}(\mathbf{O}^t)$. The difference¹⁷ being that $\Pi^{\wedge}(\mathbf{O}^t)$ specifies the sequence of *realized* partial consumption vectors, \mathbf{X}^t , rather than the sequence of *candidate* partial consumption vectors \mathbf{X}^t . Recall that one difference between these two types of partial consumption vectors is that \mathbf{X}^t may be equal to the zero vector, while \mathbf{X}^t may not.

In this way, it is not necessary to define any new law of motion that rigorously specifies how to find the actual path when the primary path intersects with a constraint. This is shown on the following page (fig. 10), with the worth constraint shown as a dashed black line.

Introducing the x worth constraint, and performing the same vector field modification, we can show the complete vector field $\Pi^{\wedge}(\mathbf{O}^t)$, which describes actual law of motion for consumption in this example (fig. 11).

Note the important feature that, for these constraints (and if we ignore the budget constraint for now), the final consumption bundle will necessarily be (5, 5). Furthermore, this is the case even if we replace the vector associated to the origin with literally *any* other vector which is $> \mathbf{0}$, or strictly positive in every component. Interestingly, this consumption bundle will necessarily result (again, ignoring the budget constraint for the moment) for nearly *any* vector field which satisfies the irreversibility of consumption requirement (each associated candidate partial consumption vector is $\geq \mathbf{0}$). Indeed, the only fields which will possibly *not* lead to consumption bundle (5, 5) are those which have regions where some associated vectors are parallel to at least one axis. In this case,

¹⁷ More accurately, $\Pi^{\wedge}(\mathbf{O}^t)$ is actually $\Pi^{\wedge}(\mathbf{O}^t, \mathbf{R}, U(\mathbf{O}^t), \mathbf{P}, w)$




Figure 10. The Worth Constraint and the Π Field.

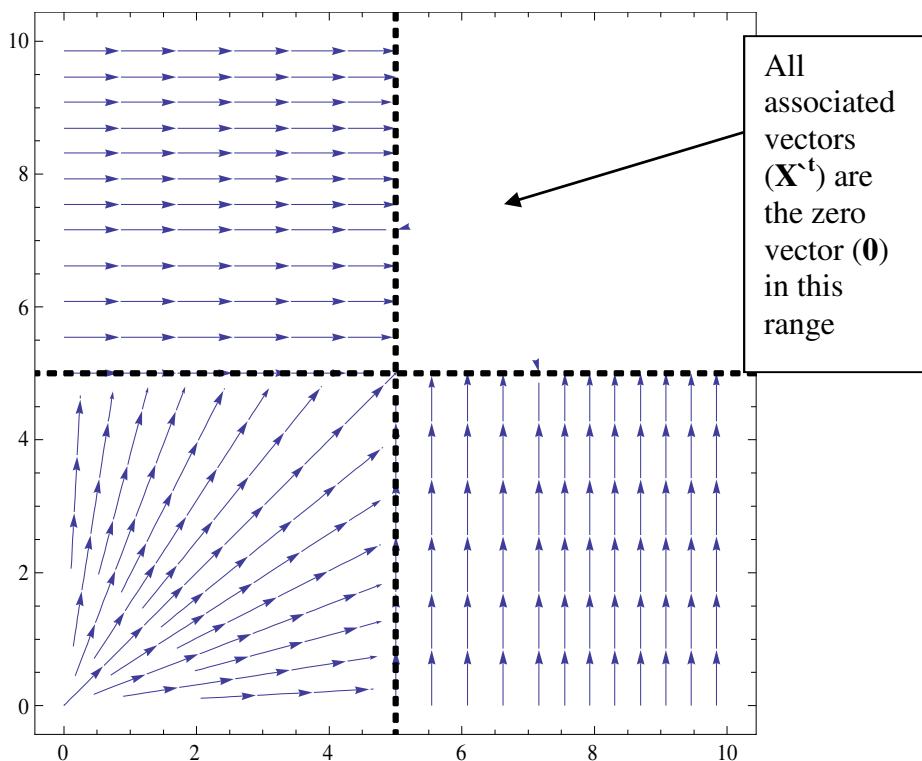


Figure 11. Complete System of Worth Constraints.

there are regions in the consumption space where the consumer simply would not *consider* purchasing any additional units of at least one particular good, under any circumstances. If consumption were described by such a field, it would be as if it simply would not occur to the consumer to even entertain the idea of purchasing additional units of that particular good. While we claim that this situation is unlikely to be a realistic description of actual behavior in most cases, we also point out that, without controversy, we can be certain that the consumption bundle (5,5) is necessarily chosen if the Π field satisfies the stronger assumption that each vector be strictly $> \mathbf{0}$, and the budget constraint is not binding.

Also, note that each individual worth constraint only censors the Π field with respect to a particular component. As a result, each worth constraint will only prevent further movement in one particular direction. If we were to shift the X worth constraint out from $x = 5$ to $x = 7$ in our previous example, the principal path would still intersect the y worth constraint at (5,5). However, consumption would still continue; it would move along the y worth constraint until consumption was restricted in all directions, at (7,5). The actual consumption path will therefore still terminate at the intersection of both worth constraints, even if one begins to bind before the other (fig. 12).




Figure 12. Final consumption bundle is (7,5) if the worth constraints are adjusted as shown. The Actual Consumption Path is in heavy black, the worth constraints are dashed, and the final consumption bundle is the point indicated.

Although this claim should be obvious, below we include a few illustrative examples (fig. 13) of distinct fields, with the same system of worth constraints, which all lead to the same consumption bundle (5 , 5).

The constant among all four panels in figure 13 is of course that the decision rule \mathbf{R} , through the system of worth constraints it imposes on the consumption space, is what actually shapes the consumption path and determines the final consumption bundle \mathbf{O}^T . This is true despite the fact that the satisficing strategy is employing a completely a-rational naïve vector field in each case. In other words, whatever rationality is present in these four particular satisficing strategies is manifested in the form of the decision rule,

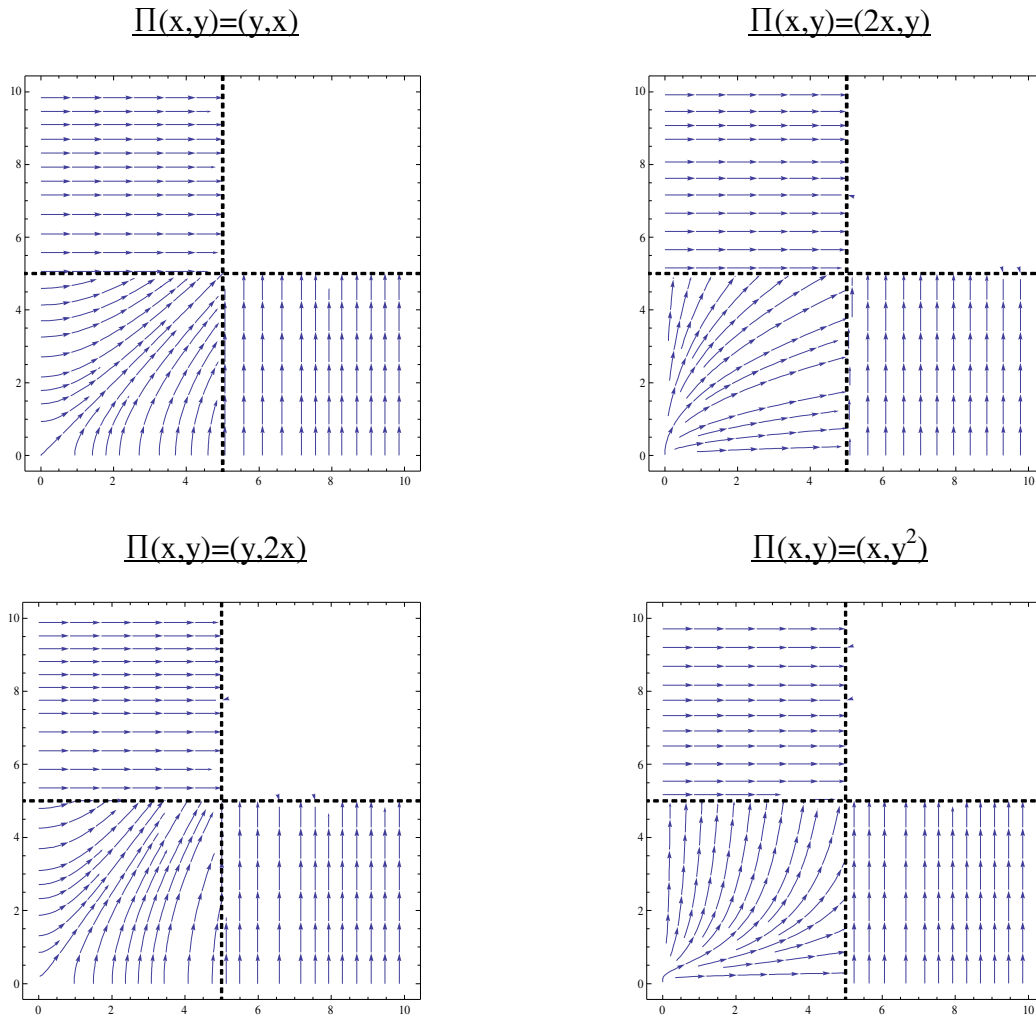


Figure 13. Worth Constraints Shaping Differing Π Fields.

and not the path-generating process. As a result, the particular choice of Π field is largely irrelevant to determining the final bundle in this case (again, so long as the budget constraint is not binding where the any of the worth constraints are slack). Because of this fact, we should feel more comfortable allowing for an essentially arbitrary choice of the system of worth constraints it imposes on the consumption space, is what actually shapes the consumption path and determines the final consumption bundle \mathbf{O}^T . This is true

despite the fact that the satisficing strategy is employing a completely a-rational (non-decreasing) Π field as we analyze specific examples of satisficing strategies, *at least* in cases where the utility function is additively separable. We will eventually show that, even in cases where the choice of vector field does influence the final consumption bundle selected, this does not, in general, prevent the possibility of any arbitrary non-decreasing vector field from forming a part of a Nash equilibrium in the overall market.

Notice, of course, that the fields here have no impact on the consumption bundle. All of the discrimination is done by the system of worth constraints, or more fundamentally, by the decision rule \mathbf{R} . In this case, any field which has all associated vectors strictly $> \mathbf{0}$ will necessarily lead to bundle (5, 5). This is precisely the reason why we might not object to the assumption that the satisficing consumer employs a naïve path-generating process; the (bounded) rationality of the satisficing strategy is still present, it just results from the decision rule rather than from the sequence of candidate consumption vectors.

Of course, if we do take the budget constraint into account, it is no longer *necessarily* the case that each of these distinct fields will lead to the same consumption bundle. But they will still all lead to the same consumption bundle, if that common bundle, (5,5) in this case, lies in the affordable region of the consumption space. If that bundle is not affordable, then consumption will terminate at some point along the budget hyperplane, and the location of that point will potentially be different for different vector fields. However, we will show in the next section that this situation (where it is possible for different fields to lead to different consumption bundles, under a utility function

which is additively separable) is not consistent with profit maximization, and so therefore cannot occur in Nash equilibrium.

With qualitatively different utility functions, we will get qualitatively different worth constraints, even using the same decision rule \mathcal{R} . For example, if instead of additively separable utility, we had a situation where two or more goods were perfect substitutes, then we will see a different shape of the worth constraint for each good.

For example, assume that the utility function for goods x and y is given by

$$U(x, y) = 100 \ln(x + y)$$

And the price vector is $(P_x, P_y) = (2, 3)$

And the satisficing rule \mathcal{R} uses $\mu = 5$

Then the locus which describes the “ x ” worth constraint is found by solving:

$$\partial U(x, y) / \partial x = P_x \cdot \mu$$

$$100 / (x + y) = 2 \cdot 5$$

$$\mathbf{y = 10 - x}$$

And the locus describing the “ y ” worth constraint is found similarly:

$$\partial U(x, y) / \partial y = P_y \cdot \mu$$

$$100 / (x + y) = 3 \cdot 5$$

$$\mathbf{y = 20/3 - x}$$

Each of these loci coincides with a particular indifference curve,

$U(x, y) = 100 \cdot \ln(10)$, and $U(x, y) = 100 \cdot \ln(20/3)$, respectively (fig. 14, fig. 15).




Figure 14. The system of worth constraints divides the consumption space into 3 regions (outward from origin): Fully unconstrained consumption, Y-constrained-only consumption, and fully constrained consumption.




Figure 15. The same graph as above, with the Actual Consumption Path (ACP) and the final consumption bundle, $(20/3, 20/6)$, superimposed.

This example has been constructed primarily to illustrate the mechanics of how the satisficing process works under this particular utility function. But it also reveals an interesting feature of this strategy, which will be discussed in more detail in the next section; there are certain price vectors to which the satisficing strategy is not capable of being a best response, for any level of μ . Obviously, given this utility function, this particular satisficing strategy (\mathbf{R} and Π) cannot possibly be a best response to the given price vector, nor to any price vector where $P_x \neq P_y$. Since the two goods are perfect substitutes with distinct prices, full rationality requires that the consumer purchase only the cheaper good (and none of the more expensive good), regardless of the total amount of wealth available. Clearly, this construction will, in general, have a region in which consumption of both goods is not bound by a worth constraint. As a result, the satisficing process will generally always involve purchasing at least some amount of every perfect substitute good for which the corresponding component of the \mathbf{X}^t vectors is positive, even if the prices of those substitutes are not equal. We will address this concern in the next section, where we will demonstrate that, if the consumer views two products as perfect substitutes, and if the consumer employs a satisficing demand strategy, then it will not be profit-maximizing for firms to charge prices which differ from each other. Thus, we will eventually suggest that, if consumers, as a group, were to employ a satisficing strategy, then they will be unlikely to face such a troublesome price vector in the first place¹⁸.

Finally, the third illustrative (though certainly not exhaustive) case involves a utility function in which goods are perfect complements. As discussed in the section dealing with the \mathbf{R} rule, the discontinuity in the marginal utility function complicates the

¹⁸ It is already well-known to be the case that profit-maximization implies that no such “unequal” price vector would prevail in equilibrium if all consumers were employing an optimizing demand strategy.

application of the decision rule somewhat, in that we must explicitly consider the left-hand partial derivative of the utility function when measuring conditional, rather than absolute, marginal utility. If the quantities of goods J and K enter the utility function as an argument of a Min function, then there is will one constraint (previously alluded to as the “inside constraint,” as it refers to the “inside” function of the compound function $U(x,y) = (f(\text{Min}(x,y))\dots)$ that is binding only if $J \geq K$. If $J > K$, then a marginal decrease in the quantity of J will not change utility (Equivalently, a marginal increase in consumption, equal to \mathbf{X}^t , will have exactly the same change in utility as a marginal increase in consumption of \mathbf{X}_{-J}^t). If a change in the quantity of J does not change the evaluation of the inside function, $\text{Min}(x,y)$, then J’s conditional marginal utility is zero. Hence the conditional marginal utility of good J is zero if $J > K$. The inside constraint is therefore described by the line $J = K$. Any \mathbf{X}^t vectors for points below this line (so that $J > K$), will have the J element replaced with a zero in the Π field (fig. 16).

If changing the quantity of J does change the value of the inside function, then the inside worth constraint is not binding. Next, the satisficing consumer must decide whether the conditional marginal utility of each good (which, in general, will be positive if the inside constraint is not binding) is sufficiently high, given the price of good J, and given μ . This is accomplished by assessing changes in the “outside” function of the compound utility function, and the locus of points for which good J provides just enough conditional marginal utility to satisfy \mathbf{R} will accordingly be referred to as the “outside worth constraint.” If it is the case that $J \leq K$, then the (conditional) marginal utility of consumption depends entirely upon the quantity of good J. Accordingly, there will be some critical value for the quantity of J, given P_J , the value of μ , and the nature of the




Figure 16. The Identity Field, with the inside worth constraint shown for good X, under $U(x,y) = f(\text{Min}(x,y))$. $\text{Min}(\)$ is the “inside function” while $f(\)$ is the “outside function.”

outside function, which just satisfies the decision rule \mathcal{R} . As a result (in 2-space) the outside worth constraint for good J will be a line that is perpendicular to the J axis, and is located at the critical value of J (fig. 17).

For any utility function with $\text{Min}(x,y)$ as an argument, the concept of CMU will result in double-counting of marginal utility, if applied only to each element of the \mathbf{X}^t vector. As described in the previous section detailing \mathcal{R}_4 , another constraint is needed to test not only the individual components x and y , but also joint components involving specific combinations of x and y . For the perfect complement case, where each quantity enters the utility function only as an argument of the $\text{Min}(\)$ function and nowhere else,

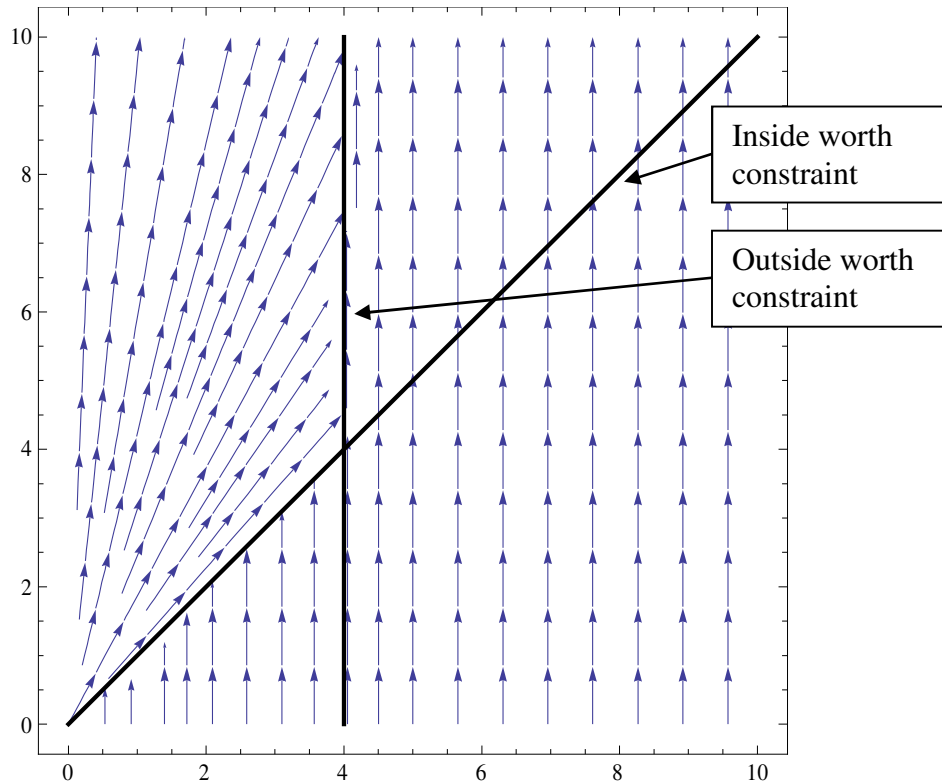


Figure 17. The Identity Field, with the inside and outside worth constraints shown for good X. The outside constraint binds at $x = 4$.

this can be accomplished by imposing an additional outside constraint, in which the CMU of good x is divided by the sum of the two prices ($P_x + P_y$) rather than just its own price.

Since the CMU for each complement will be equal to the marginal utility of the bundle (X_i^t, X_i^t) , then we can achieve a valid cost-benefit comparison by using the total price of the bundle, rather than just the price of one good, in our estimate of opportunity cost. Since prices are always strictly positive, this second outside constraint will always be necessarily more restrictive than the original inside constraint. Therefore, we may disregard the fact that there are, strictly speaking, two distinct outside constraints. Instead we shall refer only to a single outside constraint, whose position is given by the quantity satisfying $\Delta_{5i}(\mathbf{X}^t)/(X_i^{t*}(P_x + P_y)) = \mu$ (fig. 18, fig. 19).




Figure 18. The complete Π^* field. The principal path is the segment beginning at the origin and ending at (4,4).




Figure 19. Same graph as above, with the inside worth constraints and peripheral streamlines removed, showing only the outside constraints and the principal path of the Π^* field.

Under a given satisficing rule \mathcal{R} , the principal path of the Π field specifies the actual sequential path that consumption will follow. We may think of this as a simple dynamical system, where the t dimension has the natural interpretation of the amount of wealth which the consumer has expended at each point along the path. The Π field, along with the decision rule \mathcal{R} , will define the Π field, and its principal path. In order to complete this description of the satisficing demand strategy, we simply need to specify how far along the principal path of Π the actual consumption path will proceed. This is as simple as specifying the amount of wealth available to the consumer. This can be accomplished, quite simply, by finding the point of intersection of the principal path of the Π field and the budget hyperplane. Graphically, if we superimpose the budget line in any of our previous three examples, the final consumption bundle \mathbf{O}^T is simply the point where the principal path of Π crosses the budget line.

This geometry fully describes the satisficing demand process. It incorporates both types of binding constraints on consumer behavior, the budget constraint and the worth constraints, as well as the path generating process Π . The final consumption bundle selected by the satisficing demand strategy will be the point of intersection between the principal path of the Π field and the budget line. Shown below are three examples (fig. 20, fig. 21, fig. 22). The budget line is in bold, and the system of worth constraints are shown dashed.




Figure 20. The budget constraint may be either more restrictive than the worth constraints...




Figure 21. ...exactly as restrictive as the worth constraints...




Figure 22. ...or less restrictive than the worth constraints.

But in any case, the final consumption bundle is determined by the interaction of the Π field, the worth constraints, and the budget constraint¹⁹. In the next section, we will allow prices to vary, thereby systematically changing the location of the system of worth constraints and the budget hyperplane. By doing so, we will be able to derive individual demand curves, firm profit functions, and firm best-response functions, which will finally allow us to demonstrate the existence of a new class of Nash equilibria in the overall market game.

¹⁹ Or just by the ACP, if it does not intersect the budget hyperplane.

CHAPTER IV

THE COMPLETE MARKET, WITH SATISFICING CONSUMERS

After having outlined the satisficing strategy itself, we now want to begin to analyze the performance and the implications that strategy, within the strategic context of a complete market. The markets in this section will consist of a representative consumer, and two (or more) firms, each of whom produces one good. The relationships among the products of distinct firms will depend upon the nature of the consumer's utility function. We will begin by allowing the price of an individual good to vary, in order to demonstrate how to derive the satisficing demand curve for an individual good. From there, we will impose a firm cost structure, and then use the demand curve in order to derive the firm profit function. Next, we will construct the best-response functions, which give the profit-maximizing price for each firm, as a function of the price(s) of the competing firm(s). This will enable us to demonstrate the existence of a new class of Nash equilibria in markets in which firms compete in prices. The new equilibrium will be a function of the consumer worth parameter μ (as well as of the Π function). This implies that, for any set of preferences, consumer wealth, and firm production technology, there will be an infinitude of distinct Nash equilibria in the overall market game. That is, assuming that consumers satisfice, and firms compete in prices, then equilibrium prices will not necessarily be fully determined by the underlying structural parameters of the problem (technology, wealth, preferences), but instead will be dependent upon the

satisficing rule employed, and more specifically by the consumer's estimate of the utility value of a dollar, μ . Importantly, in a wide set of circumstances, this estimate μ will be self-fulfilling. Whatever the estimate of μ that a consumer chooses as a part of her satisficing strategy, that value will end up being correct under the final equilibrium price vector, once firms compete in prices. Furthermore, despite the ex ante less-than-full rationality exhibited by a satisficing strategy, we will demonstrate several distinct cases where the use of such a strategy by the consumer will end up inducing price competition between firms which has the effect of rationalizing the satisficing strategy itself. That is, despite not being a true algorithm for finding an optimal consumption bundle under *any* possible price vector, the satisficing strategy, as previously defined, has the important feature that it constitutes a best response to the best response to itself, and will therefore form a part of a Nash equilibrium strategy profile in the complete market game.

Our first step is to begin to describe the satisficing demand curve. We will assume that every firm produces only a single good, and we will derive a separate demand curve for the product of each firm, even if the representative consumer views the output of one firm as a perfect substitute for the output of another firm. Recall from the previous section that of the three fundamental structures of the satisficing demand strategy (Π field, satisficing decision rule \mathcal{R} /system of worth constraints, and budget constraint), we have assumed that one of these, the Π field, is independent of the price vector, while the other two are each a function of prices. In later work, we expect that it will be of interest to explore the consequences of allowing the Π field itself to vary, either as a function of prices, or as a separate object of competition amongst firms (for example, exploring the hypothesis that a primary beneficial effect of advertising is to alter the Π field in favor of

the firm which is doing the advertising). That is, we will later wish to explore the hypothesis that firms might prefer to compete, not in prices, but in the space of influencing, possibly through advertisement, the order in which consumers consider purchasing goods under a satisficing framework. For now, we will continue with our assumption that the Π field is naïve, and does not change as the price vector changes. From the consumer's point of view, this implies once again that it is the decision rule (and its resulting worth constraints) which discriminates amongst potential consumption bundles, and not the Π field itself.

The budget constraint obviously shifts as prices change (holding wealth constant), as does the system of worth constraints (holding μ and \mathbf{R} constant). By varying the price of a single good, we observing the resulting changes in the effect of the constraints on the satisficing process, we can derive a demand curve for every firm and every good in a straightforward manner.

Specifically, the demand curve for any individual good can be shown to be derived from a combination of (at least) two distinct loci. In the consumption space, varying the price of any good will have the effect of shifting the position of both the budget constraint and the worth constraints. In general, these two types of constraints will not shift at the same rate or in the same direction, so it is the net effect of both shifts which will allow us to see how quantity demanded varies with price. For any good, the quantity demanded for that good will be defined as the minimum of two critical values: the quantity at which the worth constraint binds for that good, and the quantity of that good at which the budget constraint binds overall. Both of these critical values are determined, in part, by the manner in which the actually sequence of partial purchases is

accumulated (i.e. the ACP). Looking at each of these quantities separately, we can describe how each critical quantity varies as the price of the good in question changes. Plotting either critical value as a function of price gives us a partial description of the demand curve.

First, consider a simple example of the worth constraint. The worth constraint represents the consumer's rule of thumb for making a yes/no decision over any potential purchase of an arbitrarily small unit of a particular good, or of a particular subset of a partial bundle of goods. The critical value at which the worth constraint just begins to bind may be found by graphing the quantity of the good on the horizontal axis, and marginal utility divided by that good's price, graphed on the vertical axis. In addition, if we plot a horizontal line at the consumer's estimate of the utility measure of the good's marginal opportunity cost (defined as $\mu * P_x$), then the intersection of the marginal utility curve and the $\mu * P_x$ line determines the highest quantity of the good in question that the consumer will be willing to purchase. As mentioned above, this value ignores concerns about whether the consumer has sufficient wealth to be able to purchase the good in question. Increasing the value of μ employed by the consumer will shift the $\mu * P_x$ line up, and will therefore decrease the quantity for which the worth just constraint binds. Decreasing the value of μ will have the opposite effect. More importantly, increases or decreases in the good's price will have an identical impact on the location of the intersection point, and hence the critical quantity of consumption.

It will be shown below that the marginal utility of the good in question may also depend on quantities of other goods consumed. If the quantity consumed of any one good is related to the quantity consumed of another good (as specified by the Π field and the

ACP), then the marginal utility of good “X” at quantity “x” will be the partial derivative of the total utility function, evaluated at the consumption bundle on the ACP that contains x units of good X. In two dimensions, this implies that the quantity demanded “y” of good “Y” is implicitly defined as a function of the quantity demanded “x” of good “X” by the Π field and the ACP. As a result, the relevant measure²⁰ of marginal utility for good X is $\frac{\partial U[x, y(x)]}{\partial x}$.

Accordingly, we can find the first critical value for the quantity of good X (fig. 23), the critical value which corresponds to the worth constraint, by solving the following for x:


$$\frac{\partial U[x, y(x)]}{\partial x} = \mu * P_x$$


Figure 23. Deriving the Worth Constraint Locus in Price-Quantity Space.


²⁰ Or more accurately, the *left-hand* partial derivative of the utility function, with respect to good x.

The curve represents the marginal utility of X, while the horizontal line is defined by $\mu * P_x$. Changes in P_x shift the horizontal line up or down. The quantity of good X at the point of intersection is therefore a function of the price, and will be a key component in the derivation of the satisficing consumer's demand curve (see below). If $U(\cdot)$ is not additively separable, then the marginal utility of one good may depend in part upon the quantity of a different good. Therefore, changes in the Π field or the ACP will have the result of changing the shape or position of the curve above (for example, as $y[x]$ changes), and therefore will also change the critical value of good X for a given μ and a given P_x .

In addition, there is a second critical value of consumption, which is determined by the budget constraint. The quantity of good X for the point where the ACP intersects the budget line defines the second critical value. Both of these critical values will be a function of good X's price. For any price, the corresponding quantity demanded must be the minimum of these two critical values. As a matter of notation, we will refer to the first critical value for good X as " x^1 " and the second critical value for good X as " x^2 " (and similarly, of course, " y^1 " and " y^2 " for the critical values of good Y) (fig.24).

Deriving the Demand Curve


Given the two types of constraints on consumption behavior (the budget constraint and the worth constraints), it is possible to use the geometry of these constraints, along with the Π field, to derive the individual demand curve for any good available to the consumer. Since an individual consumer will continue to purchase differential units of each good up to the point where at least one constraint just begins to bind, we can use



A decrease in the price of one good will shift the budget line out, leading to an increase in the budget constraint critical value for that good, from x^2 to x^{2*} .


A price increase will have the opposite effect.

(ACP is shown here as the curved line) In this case, neither good is bound by its worth constraint at the prices and quantities shown.



The ACP is a function of the price vector, so it is possible for the consumption of one good to be capped at the quantity for which its worth constraint is binding (as shown here for good Y). This does not alter our analysis for finding the budget-constraint critical value for good X, though it will, in general, lead to a higher quantity of good X demanded at any price, relative to the above case where good Y is not worth-bound.

The new x^2 after the price change is determined by the point of intersection of the ACP and the budget line.



In the case where the ACP reflects the fact that the good for which we are deriving the demand curve (here, good X) is worth-bound at a quantity within the budget set, we have a choice about how to define x^2 . We could continue to use the quantity where the ACP intersects the budget line (x^2), or we could use the point where the ACP *would have* intersected the budget line (x^{2**}) if we could ignore the worth constraint for good X.

We will use x^{2**} , but this choice does not alter the ultimate shape of the demand curve.

Figure 24. From Consumption Space To Demand Curve Loci.

each constraint to calculate two critical values for consumption, based upon the other parameters of the consumer's budget problem: wealth, prices, and preferences. Specifically, the consumer will employ a process by which she sequentially accumulates differential units of each type of good available, and will stop purchasing additional units of good X once she either runs out of money, or she has already consumed a large enough amount of good X so that additional purchases of X do not produce a sufficient amount of marginal utility relative to the estimated opportunity cost of that good (in other words, when it does not become "worth it" to buy any more good X at the going price). Therefore, the first critical value, x^1 , is the amount of good X which corresponds to the point of intersection between the marginal utility curve and the $\mu \cdot P_x$ line. This is the point at which the consumer will stop buying any more units of good x, *because of the fact that she no longer feels that the benefit exceeds the cost*. The second critical point, x^2 , which is also a function of the price of good X, is the point at which the ACP intersects the budget line. This is the point at which the consumer will stop buying any more units of good X, *because she has exhausted her income*. Since either constraint becoming binding will prevent the consumer from purchasing any more good X, we can use these constraints to find the quantity demanded for good X at each and every price of good X. Specifically, at any price of good X (given the price of good Y, etc.), the quantity demanded at that price must be the minimum of the two critical values, and as mentioned, each of these values will depend directly upon the price of good X.

The demand curve is therefore constructed in a piecewise fashion: for "low" prices, the demand curve is the locus of price-quantity pairs for which the budget constraint just binds; for "higher" prices, the demand curve is the locus of price-quantity

pairs for which the worth constraint just binds. The point of intersection of these two loci will, in general, form a kink or corner in the individual demand curve, and this “kink point” will be a focus of later discussion. The shape and location of each of these loci, and therefore the shape and location of the demand curve itself, will also depend on the particular sequential path followed by consumption, the ACP, which is in turn derived from the Π field and the decision rule \mathbf{R} , in the manner described in section II.

Individual Loci and Demand Curve

So far, we have described the location of the two critical values in term of the marginal utility space and the consumption space. In order to derive the demand curve, we must translate these critical values into price-quantity space. This is easily done simply by for solving for each critical quantity as a function of price for the good in question.

Consider an example where the consumer’s utility function is of the form:

$$U(x,y) = A \ln [Bx + Cy + D]$$

and we will derive the demand curve for good X using the following parameterization:

$$u[x,y] = 3000 \ln [(x + y)/10 + 1]$$

$$\partial u / \partial x = 300 / [(x+y)/10 + 1]$$

$\Pi(x, y)$ is the identity field; the line $y = x$ is its principal path

$$w = 100$$

$$\mu = 10$$

$$P_y = 5; (P_y \text{ is the price of the competing firm})$$

The first critical value of x (called “ x^1 ”) is found by solving the following system of equations²¹ for x :

ACP: $y(x) = \text{Min}[x, y^1]$

Worth Constraint: $\partial u[x, y(x)]/\partial x = \mu * P_x$

For the above parameterization, this gives the value of x^1 as a function of P_x , as shown below. This is the locus of all price-quantity pairs of good X for which the worth constraint just binds (fig. 25):




Figure 25. X^1 as a function of price (P_x).

Likewise, the second critical value, x^2 , is found by solving for x in the following system of equations:

ACP: $y(x) = \text{Min}[x, y^1]$

Budget Constraint: $w = x * P_x + y * P_y$

²¹ Note that x^1 and y^1 are defined in terms of each other.

Strictly speaking, we ought to consider the possibility that good X may become worth-bound if the ACP reaches x^1 at a point within the budget set. We would therefore otherwise wish to adjust our algebraic description of the ACP to account for this possibility. However, when we realize that the calculation of x^1 already accounts for this contingency, we recognize that, in practice, we are able to ignore this with respect to our definition of x^2 . In other words, just as we ignored any concerns over whether the consumer actually had sufficient income to purchase the quantities of good X in question when we calculated x^1 , here we are able to ignore any concerns about whether the quantities of good X actually satisfy \mathbf{R} when we calculate x^2 . It is important to realize that we are not able to ignore such considerations for any good *other* than the good for which we are presently deriving the demand curve. If any other goods are worth-bound within the budget set, that fact absolutely must be accounted for in our description of the ACP as we derive both x^1 and x^2 .

Recall that this is essentially the same issue as whether to use x^{2^*} or $x^{2^{**}}$ as the second critical value, as mentioned of the third panel Figure 24. Again, the choice of whether we should use x^{2^*} or $x^{2^{**}}$ is largely a matter of taste, as it will have absolutely no effect on the ultimate quantity demanded for good X, since we have defined quantity demanded for good X at a given P_x as $\text{Min}[x^1, x^2]$. In neither case (x^{2^*} or $x^{2^{**}}$) will x^2 ever be less than x^1 if the X worth constraint is binding for a particular price and consumption bundle (fig. 26).




Figure 26. X^2 as a function of price (P_x).

Superimposing both loci shows the outline of the demand curve (axes inverted) (fig. 27):

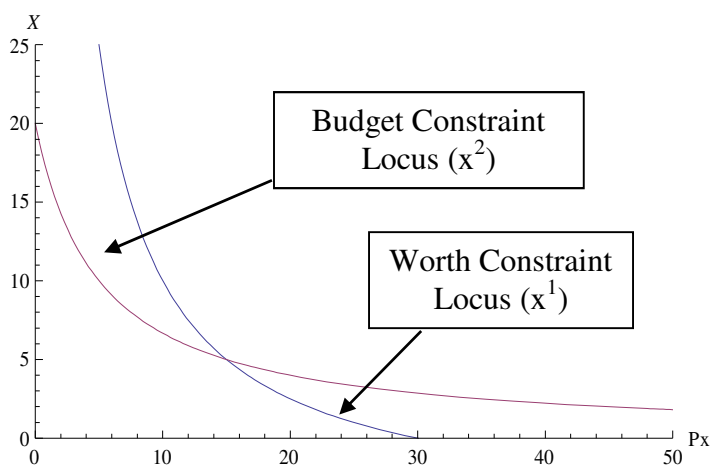


Figure 27. Both Loci Superimposed.

Plotting the minimum of $\{x^1, x^2\}$ as a function of P_x gives the complete demand curve for product x , given the parameters of the problem (fig. 28).




Figure 28. Satisficing Demand Curve for good X (Axes Transposed).

Transposing the axes presents the demand curve in standard format (fig 29):




Figure 29. The standard presentation of the satisficing demand curve. The “Worth Constraint Locus” forms the upper segment of the demand curve, while the “Budget Constraint Locus” forms the lower segment.

As a result of being derived from the minimum of two loci, the individual demand curve will, in general, have a kinked shape. This is qualitatively identical to the demand

curve facing an oligopolist, as described in 1939 IO papers by Sweezy and by Hall and Hitch. Both Sweezy and Hall and Hitch found modest empirical support for this shape of the demand curve, at least to the extent that several business owners indicated in survey responses that they *believed* they actually faced a demand curve such as this, and conducted their pricing decisions accordingly (most notably, several business owners responded to the questionnaire that, despite the recommendations made by the concept of marginal cost pricing, they set their own prices based, in part, on some concept of average total cost. This leads to higher prices realized in the market than what would have otherwise been predicted by standard models of price competition. When confronted with economic theory that predicted higher profits resulting from lower prices, several respondents declined to lower price due to a belief that doing so would not sufficiently increase their quantity sales. This is belief is entirely consistent with the predictions made by a satisficing demand curve.). Stigler later found no empirical evidence for a kinked demand curve in several oligopolistic industries, though it should be pointed out that all of the markets he looked at involved wholesalers and intermediate suppliers. As such, the buyers in these markets were all large firms in reasonably concentrated industries, who were not price-takers. As a result, all of the firms in Stigler's study would be a poor fit for the hypothesis of a satisficing consumer (as described in [field paper]). We will show below that the satisficing strategy is likely more appropriately employed by a price-taker than by a consumer with market power.

Returning to our model, we can begin to describe competition between firms, by using our newly constructed demand curve. We can find the profit-maximizing price for this firm once we assume a cost structure. We can use this solution (which will obviously

be a function of the competing firm's price) to construct a best response function in the price game for each firm. Interestingly, we can show very easily that, in 2 dimensions, and for linear ACPs of the form $y = a*x$, $a > 0$, it is never profit-maximizing for a firm to set its price along the lower portion of the demand curve.

Price-quantity pairs along this lower portion of the demand curve satisfy the following two equations, so we can easily solve for both x and y as a function of prices and the parameter a .

$$\begin{aligned} w &= p_x x + p_y y \\ y &= a*x \quad (a > 0) \end{aligned}$$

$$\begin{aligned} x &= \frac{w}{p_x + a*p_y} \\ y &= \frac{w}{\frac{p_x}{a} + p_y} \end{aligned}$$

Differentiating total revenue with respect to price for each firm shows the result:

$$\begin{aligned} TR_x &= \left(\frac{w}{p_x + a*p_y} \right) p_x \\ \frac{\partial TR_x}{\partial p_x} &= \frac{(p_x + a*p_y)w - p_x w}{(p_x + a*p_y)^2} \\ \frac{\partial TR_x}{\partial p_x} &= \frac{w}{(p_x + a*p_y)} - \left(\frac{w}{(p_x + a*p_y)} \right) \left(\frac{p_x}{p_x + a*p_y} \right) \\ \frac{\partial TR_x}{\partial p_x} &= \left(1 - \left(\frac{p_x}{p_x + a*p_y} \right) \right) x \geq 0 \end{aligned}$$

The inequality in the final line is a strict inequality if p_x , p_y and x are all strictly positive.

Likewise, for good Y:

$$TR_y = \left(\frac{w}{\frac{p_x}{a} + p_y} \right) p_y$$

$$\frac{\partial TR_y}{\partial p_y} = \frac{\left(\frac{p_x}{a} + p_y \right) w - p_y w}{\left(\frac{p_x}{a} + p_y \right)^2}$$

$$\frac{\partial TR_y}{\partial p_y} = \frac{w}{\left(\frac{p_x}{a} + p_y \right)} - \left(\frac{w}{\left(\frac{p_x}{a} + p_y \right)} \right) \left(\frac{p_x}{a} + p_y \right)$$

$$\frac{\partial TR_y}{\partial p_y} = \left(1 - \left(\frac{p_x}{a} + p_y \right) \right) x \geq 0$$

The inequality in the final line is a strict inequality if p_x , p_y and x are all strictly positive.

In other words, along the portion of the demand curve where the budget constraint binds and the worth constraint does not, revenue is strictly increasing in price. If costs are not decreasing in quantity, then profits must also be strictly increasing in price, as well. This implies²² that neither firm will ever, in equilibrium, price its product so as to locate along the bottom portion of the demand curve; price must necessarily be set at the kink in the demand curve, or higher. This is the algebraic consequence of the “race to the budget line” story from [field paper], in which muted substitution effects on the part of the consumer induce firms to attempt to capture as much of the consumer’s wealth as possible, as quickly as possible, by attempting to price as high as possible but not higher

²² Assuming a representative consumer. There are a richer set of implications to be derived from consumer heterogeneity in one or more of income, preferences, or particular choice or estimate of the value of μ , but these will not be explored in the current paper.

than the consumer's worth constraint, in contrast to the traditional "race to the bottom" in standard Bertrand competition where each firm tries to undercut the competitor by the smallest possible amount.

Therefore, it can never be a profit-maximizing strategy for a firm to price *below* the price corresponding to the kink in the demand curve, though it *may* sometimes be optimal to price above the kink. In either case, the worth constraint *must* hold with equality, and so the consumer's prior μ will necessarily be confirmed in the final accumulated consumption bundle, if we have price competition and profit-maximization. This is significant, since it guarantees that the equilibrium price-response by each firm to the consumer satisficing strategy will be such that the worth constraint binds. This implies that the marginal utility per dollar in the consumer's final bundle *must* be exactly μ . This result shows that, at a minimum, the consumer satisficing strategy, where purchase decisions are made on the basis of the rule-of-thumb, rather than on the complete "fully rational" optimizing algorithm from standard consumer theory, will be supported by the equilibrium concept of rational expectations equilibrium: consumers based their decisions on the belief that each dollar spent on consumption would yield them *at least* μ units worth of marginal utility (and possibly more). Once they have behaved in this way and assess the results of their actions, they find that, indeed, every dollar did get them at least μ , and the marginal dollar of consumption yielded *exactly* μ . In other words, consumers will necessarily find that, whatever their estimate for the utility value of a dollar, they were correct. At this point, the equilibrium concept supporting the outcome is different for each player in the market game: each firm's action vis-à-vis the other player and the consumer is part of a Nash equilibrium (slight abuse of

terminology here, since technically we can not have a Nash equilibrium if *all* three parties are not playing a best response), while the consumer's actions are justified by a rational expectations equilibrium. More on this below.

We now return to the profit functions for each firm. For simplicity, we will use constant marginal cost throughout this paper. Again, following Sweezy and Hall/Hitch, it has already been shown that, for kinked demand curves, the profit-maximizing price will be the price at the "kink" in the demand curve, for a range of values for marginal cost. This follows from the fact that the shape of the demand curve implies that marginal revenue will be discontinuous at the kink. There will be a range of values of marginal cost such that marginal revenue is lower than marginal cost for all prices below the "kink price," and marginal revenue will be higher than marginal cost for all prices above the kink price. Therefore, the profit-maximizing price occurs exactly at the kink, if marginal cost lies within the given range. The above proof implies that marginal revenue is negative for quantities above the kink quantity (since total revenue is increasing in price below the kink), so for linear ACPs in two dimensions, the lower bound of the range of marginal cost that implies profit-maximization at the kink is not greater than zero. Later we will extend this result to more general situations.

Finally, note that the location of the kink is itself a function of both firms' prices, and so the location of the kink in each firm's demand curve will change with the competitor's price.

Holding wealth and P_y constant, we can graph firm X profit as a function of P_x (fig. 30). This is done numerically (see Appendix).

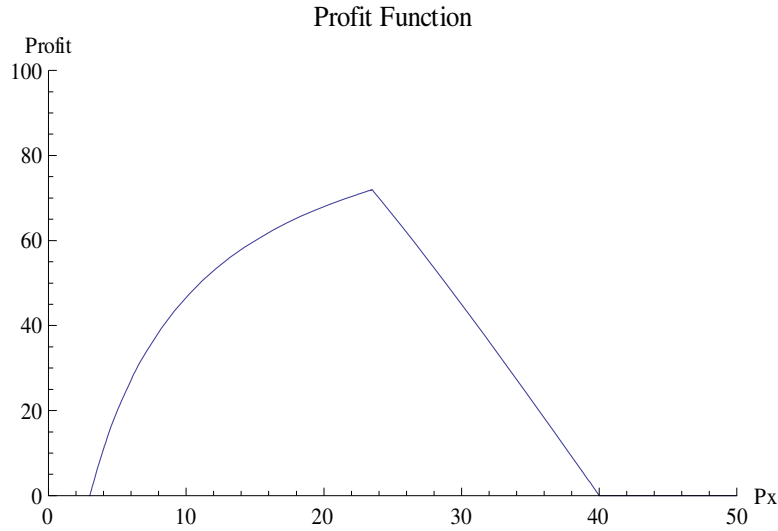


Figure 30. Firm X profit function, for $U(x, y) = 3000 \ln [(x + y)/10 + 1]$, $W = 100$, $\mu = 10$, $P_y = 5$, and $\Pi(x, y) = (x, y)$, with constant marginal cost $c = 3$. Note that the maximum profit in this example is attained by pricing good X at the “kink price” of 15.

Next, by allowing the price of the competitor’s product, P_y , to vary, and graphing the firm X profit maximizing price as a function of P_y , we can derive firm X’s best-response function (fig. 31). This is also done numerically (see Appendix).




Figure 31. Firm X Best Response (as a function of P_y) is shown. Superimposed is the line $P_x = P_y$. By symmetry, this demonstrates the existence of a Nash equilibrium in the firm pricing game at a price profile of (20, 20).

For this example, when each firm charges a price of 20, then each firm is profit-maximizing relative to the behavior of the other firm, and the behavior of consumers. Since neither firm has an incentive to change price (and doing so will lead to strictly lower profit), then, if we can verify that the consumer is purchasing a consumption bundle which is utility-maximizing given her preferences, income, and price vector, then we have demonstrated the existence of a Nash equilibrium in the overall market game in which the consumer plays a satisficing, rather than optimizing, demand strategy (fig. 32).




Figure 32. The demand curve shown for good X, when P_y is equal to 20. By symmetry, the demand curve for good Y when P_x is equal to 20 is identical. Notice that the “kink” is located at the equilibrium price of 20 for each firm. Under the price vector (20, 20), this satisficing consumer purchases the bundle (2.5, 2.5).

First of all, it is straightforward to verify that this satisficing consumer is exhausting all of her available wealth (100) when she purchases the bundle (2.5, 2.5) at prices (20, 20). She spends 50 units of wealth on each good X and Y, for a total of 100. Furthermore, since goods X and Y each enter the utility function as perfect substitutes, it must be the case that the marginal utility of each good is identical, at each and every

possible consumption bundle. In other words, the marginal rate of substitution is exactly equal to 1. Since prices are also equal in this price profile, we also have that the price ratio of the two goods equals 1. We therefore necessarily have satisfied the utility-maximizing conditions that:

- 1) The consumer's MRS = Price Ratio
- 2) She spends all available wealth.

As a result, there is no opportunity for the consumer to achieve a strictly higher level of utility under this price vector. In other words, this particular satisficing demand strategy is clearly a best-response to the price vector (20, 20). Had the consumer used the standard optimizing algorithm as her demand strategy under this price vector, instead of using her satisficing strategy, she would have ended up with the same *action choice*, which in this setting means she would have ended up with the exact same final consumption bundle (2.5, 2.5). The price vector (20, 20) represents profit-maximizing behavior for each firm, given the strategy choices of the consumer and the competing firm, and the consumption bundle yields to the consumer the highest attainable utility under that price vector.

In other words, we have shown that the complete strategy profile constitutes a Nash equilibrium.

Furthermore, since we have a utility function for which both goods X and Y are perfect substitutes, and we have two firms X and Y which compete in prices, we already know,

that the standard Nash equilibrium (which we would obtain on the assumption that the consumer uses an optimizing, rather than satisficing, demand strategy is the Bertrand result: each firm sets price equal to its marginal cost, which in this example was constant at $c = 3$. Clearly, we have qualitatively different pricing behavior being supported in each of these two distinct Nash equilibria, so the differences between them are not trivial. This distinct, “satisficing equilibrium” in the overall market game has not been the subject of any previous research of which the author is aware.

In the previous example, we relied on the fact that the two goods being perfect substitutes implied that the marginal rate of substitution was equal to the price ratio, at the final satisficing consumption bundle. Next we shall follow a similar analysis of a problem with a utility function for which the goods X and Y are not perfect substitutes, and we shall demonstrate that the utility maximization conditions will still hold under the equilibrium price vector, as a fundamental consequence of the satisficing strategy itself.

Consider a similar problem, where we change the utility function so that the goods X and Y are no longer perfect substitutes. Let us leave all of the other parameters of the previous problem unchanged, except for the fact that the utility function is now of the form

$$U(x, y) = A \ln(B \cdot x + 1) + C \ln(D \cdot y + 1)$$

Since we no longer have perfect substitute goods, we are no longer able to invoke the standard Bertrand result that, in equilibrium, $P = MC$. Instead, we must do a bit of math to calculate the equilibrium price vector for differentiated Bertrand competition (which is

what we would have if the consumer employed the standard optimizing demand strategy).

For the parameterization $A = C = 2000$; $B = D = 1/5$; $W = 100$; $c = 3$, it can be shown (see Appendix I) that the equilibrium price vector under differentiated Bertrand competition is (9.39, 9.39).

First of all, we solve the system of equations

$$W = P_x * x + P_y * y$$

$$(\partial U(x, y)/\partial x) / P_x = (\partial U(x, y)/\partial y) / P_y$$

for both x and y , in order to derive the individual demand functions. This gives

$$x = \frac{5(20 - P_x + P_y)}{2P_x} \quad \text{and} \quad y = \frac{5(20 - P_y + P_x)}{2P_y}$$

Firm total revenue is given by multiplying price times quantity demanded, so

$$TR_x = \frac{5(20 - P_x + P_y)}{2} \quad \text{and} \quad TR_y = \frac{5(20 - P_y + P_x)}{2}$$

And the profit function for each firm is

$$\text{Profit}_x = (P_x - c) \frac{5(20 - P_x + P_y)}{2P_x} \quad \text{and} \quad \text{Profit}_y = (P_y - c) \frac{5(20 - P_y + P_x)}{2P_y}$$

Plotting the ArgMax of each profit function, as a function of the other firm's price gives

(fig. 33):




Figure 33. The Bertrand equilibrium price vector, (9.39, 9.39), can be solved for numerically. See Appendix.

When faced with the equilibrium price vector, this optimizing consumer will purchase the bundle (5.32, 5.32).

Next, let us contrast this equilibrium result with the resulting equilibrium strategy profile if the consumer employs a satisficing demand strategy, rather than an optimizing one. As before, let us assume that the satisficing strategy used by the consumer involves $\Pi(x, y) = (x, y)$, and $\mu = 10$. Different assumptions concerning Π and μ will lead to different equilibrium price vectors and consumption bundles. We will discuss this in greater detail later on. At this point, our primary goal is to demonstrate the *existence* of distinct equilibria from the differentiated Bertrand competition result above.

Using the same process as above, we can derive the satisficing demand curve for good X by calculating the corresponding loci for each the two constraints on consumption. Shown here is the demand curve for good X, when $P_y = 9.39$, and the other parameters are as above (fig. 34).




Figure 34. Satisficing Demand Curve for X when $P_y = 9.39$.

This demand function, under the same assumption of constant marginal cost production technology with $c = 3$, gives the following firm X profit function (the profit function below (fig. 35), and the demand curve immediately above are both shown here for the case that the competing firm is charging the differentiated Bertrand equilibrium price of 9.39):




Figure 35. Satisficing Firm X profit function for $P_y = 9.39$.

Once again, allowing P_y to vary, and plotting the profit-maximizing price for firm X as a function of firm Y's price, gives the best response function for firm X (fig 36). The $P_x = P_y$ line is once again shown, to illustrate the location of the equilibrium price of 30.0.




Figure 36. The equilibrium price vector, $(30, 30)$, can be demonstrated numerically. See Appendix. The best response function is flat above $P_y = 40$, because, as we can see from the demand curve, 40 is the “choke price” at or above which firm Y will sell exactly 0 units. Therefore, prices at or above 40 for firm Y do not impact firm X's objective function, and firm X would effectively act as a monopolist.

As before, having shown that charging a price of 30 is profit maximizing for each firm, given that the other firm is also charging 30, and the consumer is utilizing the particular satisficing demand strategy $\{\Pi = (x, y), \mu = 10\}$, we further need to demonstrate that the consumer's satisficing strategy constitutes a best response to this price vector. The first step in demonstrating this is to consider the symmetric individual good demand curves, for the case where the price of the *other* good is 30. This/these is/are shown here (fig. 37):




Figure 37. Once again, observe that the kink price coincides with the equilibrium price of 30.

From this demand curve, we can verify that the corresponding quantity demanded is precisely $5/3$ units of each good. At a price of 30, the consumer spends 50 units of income on each good individually, or 100 total, thereby exhausting her entire income. From the utility function, we know that the marginal utility of each good at this quantity is $2000 / (5 + 5/3) = 300$. Marginal utility (300) divided by price (30) gives precisely 10, which is our value of μ in this satisficing strategy. Here again, since marginal utilities are both equal, and prices are both equal, we have also satisfied the utility maximization condition that $MRS = \text{price ratio}$, since these are both once again equal to exactly 1.

Hence, there is nothing (including both changing any of the parameters of the satisficing strategy, as well as changing from a satisficing to an optimizing strategy) which the consumer can do to unilaterally increase her utility under this price vector. Once again, this demonstrates the existence of a Nash equilibrium in the market game which is distinct from the standard differentiated Bertrand equilibrium.

So far, we have only looked at symmetric examples, and therefore we have been able to demonstrate that the satisficing strategy is a best response for the consumer because the marginal utilities are equal for both goods, as are the prices. In other words, we have been able to demonstrate the existence of a novel Nash equilibrium in the market game because the MRS and the price ratio have both been equal to exactly 1. We next quickly demonstrate that this is not necessary for a satisficing Nash equilibrium to exist in a market.

If we revisited the immediately previous example, but break the symmetry between the two goods by altering the utility function, we can see that there will still exist a distinct Nash equilibrium involving a satisficing strategy on the part of the consumer.

Specifically, if we alter the utility function so that

$$U(x, y) = 2000 \ln(x/5 + 1) + 1000 \ln(y/5 + 1)$$

then we obviously have a situation where the symmetry between the two goods has been broken. It can be shown (see Appendix I) that, if the consumer plays an optimizing strategy, then the differentiated Bertrand equilibrium in the market game includes the price vector $(P_x, P_y) = (12.73, 7.01)$, with the consumer purchasing the bundle $(x, y) = (5.41, 4.45)$. The final marginal utility per dollar spent for each good is 15.1.

If we calculate how the equilibrium changes if the consumer were to employ the satisficing strategy $\{\Pi = (x, y), \mu = 10\}$ once again, we will discover that the equilibrium price vector becomes $(P_x, P_y) = (26.65, 13.35)$, and the consumer purchases the bundle $(x, y) = (2.505, 2.491)$. Notice, despite the fact that the consumer is worse off than in the

optimizing equilibrium (less consumption of each good, less total utility), this strategy profile is nonetheless a Nash equilibrium, since the consumer cannot unilaterally improve her utility by altering her strategy in any way, including by switching from a satisficing to an optimizing demand strategy.

Once again, this is straightforward to verify y:

$$(P_x, P_y) \cdot (x, y) = (26.65, 13.35) \cdot (2.505, 2.491) = 100 = W$$

And

$$U_x(2.505, 2.491) = 266.5$$

$$U_y(2.505, 2.491) = 133.5$$

$$\text{So the MRS} = 1.99625$$

Which is precisely equal to the price ratio, as

$$26.65/13.35 = 1.99625$$

Since the consumer exhausts all of her wealth, and purchases a bundle for which the marginal rate of substitution is equal to the price ratio, she is maximizing her utility, *given the price vector*. Therefore, we have a Nash equilibrium in the market game, and furthermore, it is one which is entirely distinct from the differentiated Bertrand outcome.

Further note that the final marginal utility per dollar for both goods is exactly equal to the value of μ employed by the satisficing strategy, $\mu = 10$.

$$U_x(2.505, 2.491)/P_x = 266.5/26.65 = 10$$

$$U_y(2.505, 2.491)/P_y = 133.5/13.35 = 10$$

It is this last fact, which has been touched upon previously, which offers the best intuitive insight into why a satisficing demand strategy can form a part of a Nash equilibrium strategy profile in the market game.

For convex preferences, there are two conditions which are, jointly, necessary and sufficient for consumer utility maximization:

1. The MRS must equal the price ratio
2. The consumer must spend all her available wealth

Together, these conditions imply that the consumption bundle which satisfies 1 and 2 will lie on the consumer's highest attainable indifference curve. By the construction of the satisficing strategy, condition 1 will *necessarily* be satisfied whenever the price vector is such that the price for every individual good is located on the upper segment of the consumer's demand curve. Recall that this upper segment was defined as the locus of price-quantity pairs for which the worth constraint for that good was just binding. This implies that for any good for which it is the case that the price is located on the upper segment of the demand curve, then marginal utility per dollar (that is, marginal utility of each good at the final bundle, divided by the price of that good) for that good must be exactly equal to μ . If *every* good is priced at a level along the upper locus of the demand curve, then *every* good has marginal utility per dollar exactly equal to μ . This of course implies in turn that every good has marginal utility per dollar equal to that of every other good. In other words, the MRS will equal the price ratio, and condition 1 is satisfied.

Furthermore, *if* the market's representative consumer is employing a satisficing strategy, then price competition among firms will guarantee that all individual prices will in fact be located somewhere along the upper segment of the corresponding demand curve. This is so, with satisficing demand, because upper-locus pricing is a necessary condition for profit-maximization. If it were ever the case that a firm was pricing strictly below the kink price, then it could unambiguously increase profit by raising its price (a proof of this statement is provided below). Price competition and profit-maximization imply that condition 1 will be met if the consumer employs a satisficing demand strategy.

Furthermore, condition 2 is satisfied whenever the final consumption bundle is located strictly on the budget hyperplane. Any price located along the lower portion of the demand curve, by construction, involves a consumption bundle located precisely on the consumption frontier.

The Satisficing Market Equilibrium

To summarize, for a given satisficing demand strategy, the kink in the individual satisficing demand curve is a very special place indeed.

Firstly, as described in Hall and Hitch, the fact that the demand curve itself is not differentiable at the kink implies that there is a discontinuity in marginal revenue at that point. This means that there will be a range of values for marginal cost, all of which imply that the profit-maximizing price and quantity for an individual firm will be at the kink. As long as marginal cost is positive, it will be profit-maximizing to price at least as high as the kink (since marginal revenue beyond the kink quantity and below the kink price will be either zero or negative—see proof below). If marginal cost is less than or

equal to the left-hand marginal revenue, then the profit-maximizing price will be exactly the kink price. If marginal cost is higher than left-hand marginal revenue at the kink, then the profit-maximizing price will occur somewhere on the portion of the demand curve strictly above the kink, located on the worth-constraint locus.

Since the kink point lies on the upper locus of the demand curve, it guarantees that, if the price of the good in question is set at or above this price, marginal utility per dollar spent on that good in the final bundle is exactly equal to μ . If this is true for all goods, then condition 1 for utility maximization is met. Furthermore, profit-maximization implies that it will, in fact, be true that all goods are priced along this portion of the demand curve.

Since the kink also lies on the lower locus of the demand curve, it guarantees that, if the price of the good in question is set at or below this price, then the budget constraint will hold with equality. Furthermore, if the budget constraint holds with equality for one good, then it must necessarily hold with equality for *all* goods (It is impossible to have insufficient income available to purchase an additional differential unit of one good, while at the same time having sufficient income to purchase an additional differential unit of a different good). Thus, if the price of any good is set at the kink price, then condition 2 for utility maximization is met.

Conditions 1 and 2 together imply that the satisficing strategy constitutes a best response on the part of the consumer to the price vector which helped define²³ the satisficing demand curves. In other words, and loosely speaking, the higher segment of the demand curve is the locus of points which satisfies condition 1 of the requirements for utility maximization, while the lower portion of the demand curve is the locus of

²³ By helping to determine the location of the worth constraints in the consumption space.

points which satisfies condition 2. The kink itself is the only point which satisfies both conditions, as it lies in/on both loci. As a result, all firms pricing at their own kink price is a necessary condition for Nash equilibrium in the market game, assuming that the consumer is playing a satisficing strategy.

Claim:

Under satisficing demand, it is a necessary condition for profit-maximization that every firm prices its product along the upper segment of its satisficing demand curve.

Proof:

Assume there exists a firm X, which produces good X, which has priced its product strictly below the kink price for its satisficing demand curve. Call the kink price for good X " P_x^k " and the actual initial price of good X " P_x^0 ".

If $P_x^0 < P_x^k$, then, by construction, the worth constraint for that firm's product is not binding at the corresponding quantity demanded for firm X's product. (This is true because, at the low price, the budget constraint becomes binding upon the consumption of that particular good before the worth constraint does). Therefore, firm X has an opportunity to raise its price at least slightly to $P_x^1 < P_x^k$, and still have it be the case that the worth constraint is not binding for its product.

Furthermore, as long as there is at least one good for which the worth constraint is not binding, this implies that the resulting consumption bundle must be located on the budget hyperplane (at the very least, the consumer would have exhausted all of her income buying additional units of the non-worth-bound good, even if every *other* good were worth bound). This implies in turn that the consumption bundles corresponding to both P_x^0 and P_x^1 satisfy the equation for the budget hyperplane:

$$W = P_x \cdot Q_x + \sum_{i \neq X} P_i \cdot Q_i$$

Solving for Q_x gives

$$Q_x = (W - \sum_{i \neq X} P_i \cdot Q_i) / P_x$$

And firms X's total revenue is

$$TR_x = P_x \cdot Q_x = (W - \sum_{i \neq X} P_i \cdot Q_i)$$

Or simply, firm X revenue equals total wealth minus the total amount of money spent on all other goods combined.

For a given price vector \mathbf{P} , the quantities demanded Q_i of all i distinct goods are given by the ACP.

When firm X increases its price from P_x^0 to P_x^1 , the budget hyperplane will shift inward, and all of the Q_i 's will change in a manner prescribed by the ACP. Denote the quantities demanded under P_x^0 as Q_i^0 's, and the quantities demanded under P_x^1 as Q_i^1 's.

Therefore, as the price of good X increases from P_x^0 to P_x^1 , firm X total revenue changes from

$$TR_x^0 = (W - \sum_{i \neq X} P_i \cdot Q_i^0)$$

to

$$TR_x^1 = (W - \sum_{i \neq X} P_i \cdot Q_i^1)$$

Shifting the budget hyperplane inward towards the origin implies that the resulting consumption bundle, after the price change, will be a point along the ACP which is closer to the origin than the original consumption bundle before the price change. Since the ACP is non-decreasing in the quantity of every individual good, we know that $Q_i^1 \leq Q_i^0, \forall i$.

(The only possible way that any of the Q_i 's can increase in response to an increase in the price of good X is if good X is already worth-bound. By hypothesis of this proof, the worth constraint is not binding for good X at either price, since $P_x^0 > P_x^1 > P_x^k$)

Since, once again, $TR_x = (W - \sum_{i \neq X} P_i \cdot Q_i)$, the condition that

$Q_i^1 \leq Q_i^0$ for all i implies the total revenue earned by firm X after the price increase is at least as large as the total revenue earned by the firm before the price change. At a minimum, total revenue remains unchanged for firm X as a result of the price change, and this only happens if all goods other than X are worth-bound at the consumption bundles corresponding to P_x^0 and P_x^1 , so that all other quantities besides that of good X remain unchanged ($Q_i^1 = Q_i^0$, for all $i \neq X$). If at least one other quantity besides

that for good X does change as a result of the price increase, then total revenue earned by firm X will be strictly higher after the price increase than before.

Regardless of whether firm X revenue increases or remains the same, profit will necessarily increase after the price change from P_x^{k0} to P_x^1 . This is true, because as P_x increases, Q_x necessarily decreases if X is not worth-bound. If marginal cost is strictly positive, then total cost will decrease while total revenue either stays unchanged or increases, as a result of the price change. Hence, firm X profit must increase when firm X raises its price, assuming that it is pricing strictly below the kink price.

Since any firm charging a price strictly below the kink price implies that profits may be increased through a price increase, this demonstrates that all firms pricing their product at or above the kink price is a necessary condition for profit-maximization. QED.

As a corollary to the point that, if the budget constraint binds for one good, then it must also bind for every other good, it should be obvious that similar reasoning applies to the possibility that the budget constraint does not bind for one good. That is, if any firm is pricing its product strictly above the kink price, this implies that the firm is pricing in such a way that the worth constraint binds for its good, but the budget constraint does not bind. This, of course, means that the consumer has extra available wealth that she *could* use to purchase additional units of the good in question, if she so desires. However, she

chooses not to purchase these additional units, because the monetary cost of the would-be additional purchases does not justify the estimated opportunity cost, in utility terms, of the money which would need to be spent. Of course, if she has additional available wealth to potentially purchase more units of this particular good, then she necessarily has additional wealth available to purchase more units of *any other* good, as well. That is, it is simply not possible for the budget constraint to bind on some goods and not on others. This also illustrates the important point that, changes in price by any one firm will potentially lead to a change in demand for all other firms (and more specifically, to a change in the location of the kink price for all other firms).

The Kinked Demand Curves and Nash Equilibrium

If the (single) consumer in the market is playing a satisficing consumption strategy, then under the following conditions, a kinked demand curve for every individual good will be a necessary condition for Nash equilibrium:

- 1) Firm marginal costs are positive everywhere
- 2) Marginal utility is bounded (so as to rule out the possibility that a consumer would ever be willing to exchange 100% of her available wealth for an infinitesimal amount of consumption)
- 3) Preferences are strictly convex

First of all, consider the set of all goods available to the consumer which do not enter the utility function in an additively separable manner. This set may or may not be empty. Regardless of the presence or absence of a distinct kink, or point of non-differentiability, it is necessary that the demand curve for every good in this set be such

that the price of each good lies on both the worth constraint locus and the budget constraint locus, if the consumer is to be playing a best response (that is, if the consumer is to satisfy conditions 1 and 2 for utility maximization). Furthermore, given that each firm has a positive marginal cost, profit maximization is only possible in this circumstance of the two loci of each of the corresponding demand curves are such that the worth constraint locus has a flatter slope than the budget constraint locus.

To see why, recall (from the proof beginning on page 150) that along the budget constraint locus, for any budget hyperplane, and for any non-decreasing consumption path, it must be the case that firm revenue is non-decreasing in price. Since the budget constraint locus is everywhere downward sloping, this implies that marginal revenue is non-positive at all points along the budget constraint. Since utility maximization requires that any candidate Nash equilibrium involve all prices lying on both loci of the corresponding demand curve, we can directly compare the marginal revenue implied by each locus.

In general, marginal revenue at any differentiable point on the demand curve is given by

$$MR = Q^*(dP/dQ) + P$$

So the marginal revenue implied by the worth constraint locus is

$$MR^{wc} = Q^*[dP/dQ]^{wc} + P$$

And marginal revenue implied by the budget constraint locus is

$$MR^{bc} = Q^*[dP/dQ]^{bc} + P$$

Profit maximization requires that each firm set their price such that they sell the largest possible quantity for which marginal revenue is greater than or equal to marginal cost.

Since marginal costs are positive by assumption, and marginal revenue along the budget constraint locus is everywhere non-positive, Nash equilibrium is only possible if marginal revenue is also positive along the worth constraint locus at a point of intersection between the two loci.

Thus, if $MR^{wc} > MR^{bc}$, since at any intersection point, both loci will by definition have P and Q in common, it follows that the slope of the worth constraint must be larger (flatter, since both slopes are negative) than the slope of the budget constraint. In other words, it is necessary for equilibrium that, in absolute value, $[dP/dQ]^{wc} > [dP/dQ]^{bc}$. Hence, in equilibrium, all demand curves (within the set of non-additively separable goods) must have a kink point, or point of non-differentiability, at the equilibrium price.

Finally, consider the set of all goods which do enter the utility function in an additively separable manner. Likewise, this set may or may not be empty. If it is non-empty, then all goods within this set will necessarily exhibit diminishing marginal utility (implied by the assumed convexity of preferences), and so all worth constraint loci for these goods must be downward sloping, regardless of the choice of Π field (since consumption of any other good will not impact the marginal utility of an additively separable goods; only the quantity consumed of the good in question has a bearing on the marginal utility of that good).

Since marginal utility is assumed bounded, each of the goods in this set will have a choke price, at or above which zero units of that good are purchased, for any choice of satisficing strategy. Since the worth constraint locus is asymptotic to the price axis for any choice of Π , there will necessarily always be a range in the neighborhood of the price axis for which the budget constraint locus lies above the worth constraint locus. Since

both loci are decreasing everywhere, if the two loci do intersect, there must be at least one point of intersection for which the budget constraint locus is as steep as or steeper than the worth constraint locus. As above, if there is to be a Nash equilibrium in the market, then the worth constraint locus must be strictly steeper than the budget constraint locus, since marginal revenue is non-positive at all points along the budget constraint locus (regardless of the particular Actual Consumption Path, provided that that path is non-decreasing, as required by the construction of this model).

Without imposing additional restrictions on the utility function, there is nothing to rule out the possibility that there is more than one point of intersection between the two loci. However, if any hypothetical additional points are to constitute possible locations for a Nash equilibrium price for the corresponding good, it must be the case that that intersection point be qualitatively identical to the point described above (i.e., that the demand curve in the neighborhood of such an intersection point be composed of a flatter worth constraint locus portion in the neighborhood immediately above the intersection, and a steeper budget constraint locus in the neighborhood immediately below the intersection). This is true, again, because of the fact that profit is unambiguously increasing in price along the budget constraint locus, for any choice of Actual Consumption Path.

To show why, first notice that, if both loci are downward sloping, then in order for them to intersect, it must be the case that the "upper" locus must be flatter than the "lower" locus in the neighborhood of the intersection point, regardless of which of the two loci is upper and which is lower. Therefore, if there were a point of intersection of the two loci, such that the budget locus was above the worth locus at the intersection, that

point could never be consistent with profit maximization (and therefore never be consistent with Nash equilibrium), since the firm could always strictly increase profit by increasing price. Furthermore, the possibility that, in equilibrium, both loci have the same slope at the intersection is ruled out by the assumption of strictly positive marginal costs (as described above).

Once again, if there is to be a Nash equilibrium in a market with a representative consumer which plays a satisficing strategy, it must be the case that all demand curves have a kink located at the equilibrium price.

For certain preferences or satisficing strategy choices, we cannot presently rule out the possibility that the resulting demand curves behave in ways which are qualitatively different from the demand curves presented in this paper. Furthermore, it is possible in principle, that in certain circumstances, there may not exist a Nash equilibrium involving a satisficing consumption strategy, since we have not yet been able to derive a rigorous proof of existence of such an equilibrium under general conditions (Though we have been demonstrated the existence of such an equilibrium within the context of a non-negligible set of arguably realistic examples). The current lack of formal existence proof notwithstanding, we still have demonstrated that it is a necessary condition that the system of satisficing demand curves behaves qualitatively as we have described them in this paper, at least locally. In other words, if a satisficing Nash equilibrium exists for any specific market, it must involve a system of demand curves which all have a point of non-differentiability at the equilibrium price for each product, with the worth constraint locus forming the portion of the demand curve in a

neighborhood immediately above this kink, and the budget constraint locus forming the portion of the demand curve in a neighborhood immediately below.

Nonetheless, we argue that the model developed in this paper does apply, without ambiguity as to whether a satisficing Nash equilibrium exists, not to all possible cases, but to a reasonably broad set of realistic circumstances. As such, we feel that this model does offer an interesting and potentially important contribution towards describing consumer behavior, and one which cannot be duplicated with existing models based upon subgame perfect consumption behavior.

Aggregation

So far, we have restricted our attention to the case where the demand side of the market is composed of a single, representative consumer. We are accustomed to having a single consumer's behavior serve as a model for an entire population of consumers, because, in an optimizing framework, the necessary conditions for the existence of a positive representative consumer have been derived by Deaton and Mellbauer (1980) and others. However, in a framework where consumers may conceivably choose a satisficing strategy, it is not clear that individual demand curves will aggregate into a market demand curve which is identical to that of an appropriately chosen notional representative consumer. A diverse population of consumers may, potentially, give rise to a market demand curve which is not validly modeled by the demand curve of a single satisficing consumer.

We do not need to justify the representative agent assumption in order to demonstrate that the satisficing equilibrium *exists*. We do, however, need to justify the

assumption in order to feel confident that the equilibrium provides a relevant description of the behavior of actual markets. To defend the applicability of the satisficing Nash equilibrium developed in the previous sections, we now wish to devote some consideration to the issue of aggregation. We should begin by recognizing that, for a population of consumers which is heterogeneous with respect to wealth and/or preferences, we needn't necessarily expect that the summation of the individual demand curves for any particular good look qualitatively identical to a typical individual satisficing demand curve. If individuals are diverse, then for any particular good, it is possible that the "kink price" for that good is different for each individual. As a result, horizontally adding the individual demand curves may lead to a market demand curve which is different from the demand curve of a potential representative consumer in at least one of two ways. Either there will be many individual kinks (one at each price which is a kink price for at least one of the underlying individual curves), so that the aggregated market demand curve is a piecework collection of convex sections connected by kinks, or the kinks tend to smooth out as many different consumers are added (see, for example, the discussion of the "regularizing effects of aggregation" in Mas-Colell, et al., page 122). In either case, it is probable that the fundamentally most important feature of the satisficing equilibrium carries through; the price vector supported by Nash equilibrium is (potentially) higher if consumers satisfice than it would be if consumers all optimized. Nonetheless, in either case, it would also be true that the market demand curves would be qualitatively different from the demand curves of any satisficing consumer, and we would not be justified in using a representative satisficing consumer to describe market behavior.

To address these concerns, and to begin to make the case in favor of a positive representative satisficing consumer, we will argue that the issues raised above (concerning both multiple possible kinks and regularizing smoothing) are possible only in situations where consumers are heterogeneous in their fundamental characteristics (wealth, preferences), *and* are playing satisficing strategies which are chosen arbitrarily. In other words, if we allow ourselves to analyze the individual demand choices of consumers, within the context of the overall strategic setting of the market game, we will have reason to conclude that reasonably rational consumers' demand strategies will be selected so as to induce a remarkable amount of regularity in the market demand curve.

In other words, the situation where the market demand curve is qualitatively distinct from the individual demand curves is not consistent with utility maximization. Basically, we will employ the concept of Nash equilibrium as our aggregation technique.

However, even if the underlying fundamental characteristics of the individual consumers are quite different among the population, we can predict a high degree of order in the stylized model if we consider the implications that each consumer ought to, eventually, be playing a best response to the same price vector. That is, even for distinct preferences and levels of wealth, if all consumers are employing a best response strategy to the prevailing price vector, then the strategy choice of each consumer will result in a system of demand functions for which the primary kink of each demand curve coincides with the prevailing price of the corresponding good.

The intuition underlying this result is identical to that in the previous section, describing why a Nash equilibrium in the representative consumer case must occur where all firms price their product at precisely the kink price. Another way of thinking about the

same idea is that, assuming that all consumers play a satisficing strategy, then in equilibrium, all consumers must be selecting a demand strategy chosen specifically so that the kink of the demand curve for each product coincides exactly with the prevailing price of that good.

Again, the reason for this result has to do with the idea of utility-maximization and that of consumer best response. If there is a situation where an individual consumer has a demand curve for which the kink price does not coincide with the actual price, then the product is priced on either the strict upper portion or the strict lower portion of her demand curve. In either case, one of the two requirements for utility maximization is not met²⁴.

Consider a population of 2 consumers, with distinct wealth and preferences, but with an identical satisficing strategy (Π, μ) .

Consumer 1:

$$\text{Wealth} = 100$$

$$U(x, y) = 3000 \ln[(x/50) + 1] + 6000 \text{Log}[(y/25) + 1]$$

Consumer 2:

$$\text{Wealth} = 200$$

$$U(x, y) = 6000 \ln[(x/25) + 1] + 3000 \text{Log}[(y/50) + 1]$$

Both consumers employ the satisficing strategy using the identity field as their choice of Π , and use $\mu = 10$.

²⁴ I believe this result holds generically, but there are exceptions. For example, it is possible to have a satisficing strategy where the underlying satisficing decision rule is \mathfrak{R}_0 , and the principal path of the Π field is Π^{FR} .

Shown below are the individual demand curves for good X (fig. 38) (note the axes are transposed relative to the standard presentation of the demand curve):




Figure 38. Individual Demand Curves for Heterogeneous Consumers.

Adding (by price) these individual demand curves yields the following (transposed) market demand curve for good X (fig. 39):





Figure 39. Non-equilibrium Market Demand Curve with Heterogeneous Consumers.

Notice that adding the two distinct individual demand curves leads to a shape which is qualitatively different from that of the typical individual satisficing demand curve. Specifically, since each individual curve has a point of non-differentiability occurring at a different price, then the market curve will have a point of non-differentiability at each of those two prices. Also, since we are constructing a market demand curve by adding two simple demand curves together, the market demand curve will have an additional point of non-differentiability located at the choke price of each of the individual demand curves, when such choke prices exist (except, of course, for the highest choke price, which also constitutes the choke price for the market demand curve). This new type of kink, which we will refer to as “secondary kinks,” is a consequence of the aggregation of a finite number of distinct individual demand curves, and is common to both satisficing and optimizing market demand curves. These secondary kinks are not of much particular interest, and we will not go into very much detail describing them or their implications, other than to note that the “direction” of secondary kinks is opposite of that of primary kinks. That is, a segment connecting two points on the demand curve on either side of, and in a neighborhood around a secondary kink will always lie above the demand curve, while a similar such segment for a primary kink will always lie below. As a result, the union of two convex sections of the demand curve which are joined by a secondary kink will still be a convex curve (though it will obviously have a point of non-differentiability) (fig. 40).



We argue that the reason we see heterogeneous satisficing consumers giving rise to market demand curves with multiple primary kinks is that the satisficing strategy used by both consumers in the above illustration is essentially arbitrary. Furthermore, while it makes sense to suspect that allowing consumers to be heterogeneous in strategy choice as well as wealth and preferences (recall, in the above example, both consumers played the same satisficing strategy) will only increase the irregularity of the market demand curve, we will demonstrate that allowing strategy choices to be made rationally will actually improve the behavior of the market demand curve.

It is important to note that strategy choices are not made in a vacuum, they are made within the context of the strategic interaction of the overall market game. So, whatever the particular combination of individual consumer strategies that comprise the market demand curves, firms will select a price for their product which is profit maximizing, given the behavior of the population of consumers, and of competing firms.

Importantly, once firms have each set their price, then the same price vector faces all consumers. In other words, no matter how diverse consumers may be in the dimensions of wealth or preferences, they will always be identical, at least, in the price vector that they face. This common price vector is the feature of the market game that allows heterogeneous consumers to select distinct strategies which ultimately end up leading to a well-behaved market demand curve.

When we conceive of consumers in the standard optimizing model, we are in effect, assuming that consumers are playing a strategy which is a best response to every possible price vector. In the satisficing framework, in order to demonstrate orderly aggregation to a market demand curve which resembles the demand curve of a representative consumer, we will need only to assume that consumers play a strategy that is a best response to a single price vector (i.e., whichever price vector the consumers *actually* face). In fact, so long as all consumers are playing a satisficing strategy which is a best response to the prevailing price vector, then all individual demand curves for all distinct goods will exhibit a single primary kink at the actual price of the corresponding good. Despite the fact that consumers are heterogeneous in wealth and preferences, the fact that they are also allowed to be heterogeneous in strategy choice (including choice of μ , which, like the utility function, is not directly observable) allows the population of consumers to “coordinate” on a particular price vector.

If all consumers have demands curves such that the kink price for every good coincides with the kink price of the same good for every other consumer, then adding the individual demand curves by price will result in a market demand curve which has exactly one primary kink. There may be as many as $N-1$ secondary kinks (where N is the

total number of individual consumers), but we will assume that these types of kinks will not substantively affect our analysis, for two reasons. First of all, if two convex sections of a market demand curve are connected by a secondary kink, then the union of those sections will still be convex. Therefore, if there is only one primary kink in the market demand curve, then the existence of a number of secondary kinks will not change the qualitative feature that the market demand curve and individual demand curves share in common: they are formed by two convex segments joined at a primary kink. Though this is not yet a perfect proof that a positive representative satisficing consumer exists (since we cannot construct an example of an individual consumer's demand curve which has secondary kinks), it is at least a step in the right direction.

In order to demonstrate that, in equilibrium, even a population of heterogeneous consumers will have demand curves which each have a primary kink at the prevailing price (which, as noted previously, also implies that each firm's price is profit-maximizing for an entire range of values of marginal cost), we will rely on similar reasoning to the proof of why a Nash equilibrium in the single-consumer case must involve all products being priced at the kink price (or, equivalently, in all kinks being located at the actual price).

Regardless of the underlying unobservable utility functions²⁵, or of the level of wealth available to a consumer, if a consumer's satisficing strategy is to be considered a best response to the behavior of the other actors in the market game, it necessarily must be narrowly equivalent to the optimizing strategy under the prevailing price vector (though again, it need not be narrowly equivalent under any *other* price vector, and it

²⁵ Although we do require standard assumptions on preferences such as rationality, convexity, local non-satiation, etc.

need not be *strategically equivalent*). The argument why all individual consumers must have demand curves such that the actual price is the kink price is similar to the reason why the single consumer Nash equilibrium must involve all firms pricing at the kink price. If any consumer is buying any good at a point on their demand curve below the kink, then she is not located on the worth constraint locus, and therefore it will not be the case that the marginal utility per dollar is equal for all goods. She can therefore increase her total utility by reallocating her income among the goods she is able to buy. I.e., she has a profitable deviation, and the initial situation could not be part of a Nash equilibrium. Likewise, if any consumer is buying any good at a point on her demand curve above the kink, then she is not spending all of her income, and assuming LNS preferences, can increase her total utility by buying more.

Only if each consumer is buying each good in such a way that every demand curve has a kink at the actual price can it be the case that every consumer is maximizing their utility relative to the price vector they face, which is of course a necessary condition for Nash equilibrium.

For any price vector, in order to find the parameters of a satisficing strategy that can be a best response to that price vector, the consumer can conceivably behave in one of at least two ways. In principle, the consumer can “learn” the best response to a price vector through something like a tâtonnement process. That is, if under a given price vector, the satisficing strategy ever leads to a consumption bundle that does not exhaust all her income, she can revise her estimate of μ downward, so that she has a lower threshold for purchases which satisfy her satisficing decision rule. If she ever has a situation where the marginal utility per dollar is not the same for all goods, then she can

increase her utility by reallocating consumption. Here, she can either increase μ , or alter her choice of Π field (or both) in order to increase her utility. This process needs elaboration, but if this type of tâtonnement process does indeed take place, then it is still going to be the case that the outcome of Bertrand competition with fully optimizing consumers will not necessarily result.

In addition to the possibility of some type of tâtonnement process, it is also possible that, in a multi-period framework, the consumer could learn the parameters of a best response satisficing strategy all in one shot, by optimizing in the first period, and then playing a particular satisficing strategy (which is derived from the first period consumption bundle) in every period thereafter. In this case, the satisficing strategy serves the role as a sort of “memory” of what the consumer did when she was optimizing. In other words, a consumer first truly does optimize with respect to her consumption choices, but having optimized in one period, the consumer next “goes from memory” in deciding which consumption choices to make in some number of subsequent periods.

In this setting, the particular satisficing strategy which is a best response to the first period price vector is easy to calculate, and we will define such a strategy as a “Calibrated Satisficing Strategy (CSS).”

We will introduce a concept which will play a pivotal role in the description of market Nash equilibria, the Calibrated Satisficing Strategy (CSS):

Calibrated Satisficing Strategy (CSS): A consumer’s satisficing demand strategy is a *calibrated satisficing strategy* for some price vector \mathbf{P} , if the strategy is narrowly equivalent to the optimizing strategy under \mathbf{P} , and the worth constraint binds for all goods at the consumption bundle selected by the strategy under \mathbf{P} .

The concept of the CSS can be used to demonstrate the way in which Nash aggregation might affect market behavior.

Notice that the condition of narrow equivalence is a necessary condition for Nash equilibrium. By definition, the consumer cannot be employing a best-response strategy to any price vector if the demand strategy is *not* narrowly equivalent to the optimizing strategy under that price vector.

It can be shown that there exists at least one CSS for any parameterization of the consumer budget problem (that is, for any specific choice of utility function with convex preferences, level of wealth, and price vector).

Claim: A CSS exists for any consumer budget problem with convex preferences

Proof:

Assume a particular convex utility function U^* , level of consumer wealth W , and prevailing price vector P .

Let \mathbf{X} be an optimal bundle²⁶ selected by the standard optimizing demand strategy. Let λ represent the shadow value of wealth of the constrained optimization problem.

Define a satisficing strategy \mathbf{S} to be the satisficing strategy (as defined in previous sections) for which $\mu = \lambda$, and for which the principal consumption path is the ray which begins at the origin and includes the vector \mathbf{X} .

²⁶ There may, in principle, be more than one such bundle. But if so, the total utility $U(\mathbf{X})$ is the same for all such bundles, and all such bundles will share a common shadow value of wealth λ .

Consumption will begin at the origin, and then move along the principal consumption path eventually ceasing precisely at point \mathbf{X} . The budget constraint will bind exactly at \mathbf{X} , since \mathbf{X} being optimal requires that it lie on the budget hyperplane, if preferences are convex. Furthermore, since the shadow value of wealth at \mathbf{X} is precisely λ , the worth constraints will all bind precisely at \mathbf{X} and not before, if preferences are convex.

In other words, strategy \mathbf{S} will select the optimal bundle \mathbf{X} under the given parameters of the consumer budget problem.

Since strategy \mathbf{S} will select bundle \mathbf{X} , and since all worth constraints will just begin to bind at this consumption bundle, then strategy \mathbf{S} is a CSS for the budget problem in question. QED.

Returning to the example of the demand curves of heterogeneous consumers from above, we can now demonstrate that, by calculating a CSS for each consumer, and allowing each consumer to play that particular best response strategy, rather than the arbitrary strategy used by each consumer in the initial example, then the market demand curve derived by adding both CSSs will have a single primary kink located at the prevailing price vector, and will consist of two convex sections which intersect at that kink. In other words, the market demand curve composed of CSSs instead of arbitrary satisficing strategies is qualitatively similar to a single-consumer satisficing strategy.

By taking an arbitrary price vector, in this case $(7,6)$, and finding each consumer's optimal consumption bundle under that vector, given her unique income level

and preferences, we can construct a CSS for each consumer. Plotting the individual CSS's gives the following (figs. 41-44):

Consumer 1:




Figure 41. Consumer 1 Demand Curve for X (under the CSS) with the price of X plotted as the dashed line.




Figure 42. Consumer 1 Demand Curve for Y (under the CSS) with the price of Y plotted as the dashed line.

Consumer 2:




Figure 43. Consumer 2 Demand Curve for X (under the CSS) with the price of X plotted as the dashed line.




Figure 44. Consumer 2 Demand Curve for Y (under the CSS) with the price of Y plotted as the dashed line.

Adding the individual demand curves gives the well-behaved market demand curves we were looking for (fig.45, fig. 46):

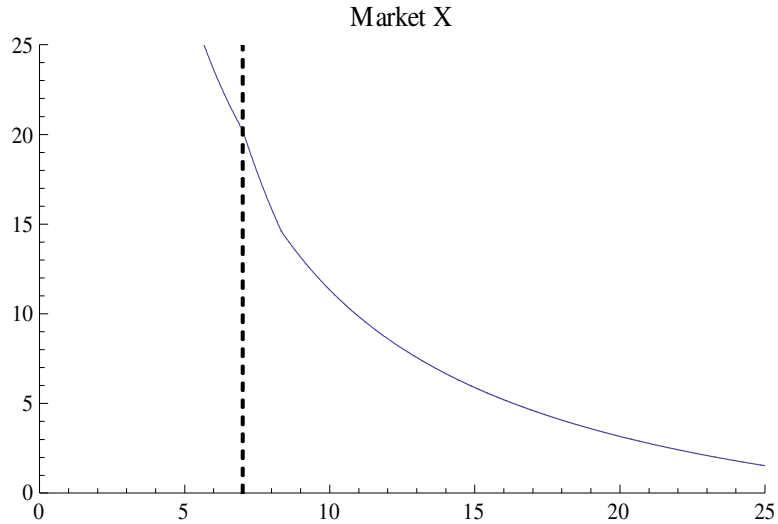


Figure 45. The aggregated market demand curves for good X, resulting from all consumers playing a best-response to the price vector \mathbf{P} (in other words, the market demand curve resulting from all consumers playing a CSS for \mathbf{P}) results in exactly one primary kink, with one convex section above and one convex section (including one secondary kink) below.




Figure 46. The aggregated market demand curves for good Y.

In summary, we can make a reasonable argument that the satisficing demand strategy aggregates to the market level in a well-behaved way (that is, up to a collection of secondary kinks which result from the “graininess” of market demand due to the discrete nature of any population of consumers), assuming that the individual satisficing consumers select their satisficing strategies in a manner which depends upon the prices vector they face. Although the concept of the calibrated satisficing strategy (CSS) is useful here, we have not yet spent any time explaining how we might expect a satisficing consumer to come to know the relevant parameters of a CSS for any particular budget problem. We shall explore this issue in the next section (Section IV: Identification). But for now, we will be content to show that, *if* all consumers play a CSS, then for each good, the individual demand curve kink prices coincide with each other and with the prevailing price of that good, and we have a well-behaved market demand curve which is qualitatively similar to the individual demand curve of a single representative satisficing consumer.

Differential Rationality

In the model as set up so far, it is natural to wonder whether any results obtained depend critically on the idea that agents are, in some sense, “stupid” when functioning in their role as consumers (and lack the computational capacity to optimize, and so must satisfice), and yet are “smart” in their role as producers and managers of firms (and therefore have the ability to work out the precise shape of the satisficing demand curve and its implications for the location of the profit-maximizing price for their own product). This is certainly a topic which deserves discussion. But, along several fronts, we believe

it is possible to demonstrate that this particular issue of differential rationality is not a major concern as far as limiting the applicability of this model's results.

Firstly, there is a scale argument to be made as to why firms might be more likely to have an incentive to be more precise about their choice of pricing strategy than consumers would be about their consumption strategy. Paralleling the argument first laid out by Stigler's (1961) seminal paper on search theory, it should be clear that the payoff to more effective consumption or pricing strategies is increasing in the amount bought or sold. Since firms tend to operate on scales much larger than those at which consumers consume, the level of potential benefit from more sophisticated strategies is likely to be higher for firms than for individual consumers. Assuming that the costs of implementing a fully rational strategy are roughly of the same magnitude in either case (or at least, that the relative difference in implementation costs is smaller than the relative difference in scale), then it would make sense for firms to devote more resources to identifying and employing a sophisticated pricing strategy (that is, a strategy which develops and employs a great deal of information about the nature of the demand system which describes consumer behavior) than what an individual consumer might choose to employ with respect to her consumption strategy (that is, she may devote relatively fewer resources to being diligently certain that she has considered all available alternative consumption choices, or that she has carefully considered the true opportunity of any partial purchase). This difference in scale would account for a distinct level of apparent rationality between firm behavior and consumer behavior, even if we recognize the fact that the consumers are the very same agents who are ultimately running the firms, and

even if we imagine that they have the same *capacity* for calculation and comparison of alternatives in either role.

In addition, there is an agent heterogeneity argument. It is possible that there are differences in the innate abilities and talents of agents, and that some consumers are more adept at making optimal choices than others. A simple version of this argument would involve two types of agents, “smart” and “stupid.” If we assume that the number of “smart” agents is small relative to the number of “stupid” ones, then we can easily envision a scenario where (relatively few) smart agents have roles making the strategic choices for firms, and the larger number of stupid agents have less sophisticated tasks in the operation of firms. In this case, firms could again reasonably be expected to be able to employ a more sophisticated strategy relative to consumers. If the number of smart agents is small enough relative to the overall population of consumers, then it may still be a *reasonable* approximation of the aggregate group of consumers to model demand as primarily deriving from the behavior of the (much larger) segment of less-sophisticated agents.

But, perhaps most importantly, it should be argued that firm behavior in this model is not necessarily fundamentally founded on hyper-accurate knowledge of the demand curve, as we might expect. And likewise, satisficing consumer behavior is not necessarily as brainless as we might expect.

To the first point, we can recall the result from Day (1967), mentioned earlier, that firms may themselves employ a type of satisficing process that converges to profit-maximization, even in the case where they do not know their own profit functions. So, extending that the result to this model, we do have some reason to believe that, even if

firms are not fully aware of the specific shape of their own demand curve (including, for example, complete ignorance of the fact that the market demand curve is kinked), then it is still possible for firm pricing behavior to converge to the (profit-maximizing) kink price. Basically, the process specified by Day amounts to a numerical optimization process in the price space, where firms charge one price, observe the resulting profit, and then charge a higher (lower) price and observe if profit went up or down. If profit increases in response to the price change, then further price changes in the same direction are explored. If profit decreases in response to the price change, then price changes in the opposite direction are explored. The step size of the price change also depends upon the previously observed size of the change in profit. Firms settle on a final price once the changes in profit become “sufficiently” small, and so the firm’s price will eventually be within an arbitrarily small distance ε from the true profit-maximizing price (so that the minimum acceptable change in profit, ε , can itself be considered the satisficing criterion for the firm pricing process).

Again, nothing about this process requires that the firm have any more intricate knowledge of the demand curve than the consumers themselves do. In fact, firms can in principle have even *less* knowledge²⁷ of the demand curve for this result to hold. Furthermore, it is certainly not required that the firms have any knowledge of the fact that consumers are satisficers rather than optimizers. All that is required is that firms be pricing at (or arbitrarily near) the kink price. This pricing behavior would be part of an equilibrium, even if firms believed, mistakenly, that consumers were behaving in manner consistent with full optimization. That is, firms may be mistaken about the nature of the

²⁷ In reality, of course, firms are likely to have far more information about market demand than consumers are, but this observation is not essential to the result from Day (1967).

consumer strategy, and the nature of consumer preferences, but if they are correct about the location of the profit-maximizing price, then the strategy profile may persist indefinitely as part of an equilibrium.

Of course, this combination of pricing strategy and (mistaken) firm beliefs are entirely consistent with the concept of model equilibrium discussed earlier. Firms may have a profoundly inaccurate misconception of the nature of the consumers' preferences and strategies, but as long as the outcome of actual game play gives them no indication that they are mistaken (and there will be no such indication if the firm is pricing at its kink price), then there will be no incentive for that firm to alter either their pricing strategy, *or their beliefs* about the consumer. In other words, it is not, in principle, necessary to conclude that firms have any more intricate knowledge of consumers, or any higher levels of computational ability or rationality, than what the consumers themselves have.

In addition, it is also not necessarily the case that consumers who choose to employ a satisficing strategy are "stupid." Particularly in the case where there is some resource or utility cost required in order to employ a more involved optimizing strategy, then it is if individual consumers are all price-takers, then a simpler satisficing strategy will yield strictly higher utility than optimizing will²⁸. This feature of the model is discussed in greater detail in the following section, under the topic of business cycles.

²⁸ "Optimizing" used here in the sense of finding a consumption bundle which yields the highest possible direct utility for every price vector, and not necessarily the highest possible payoff for the payoff function which includes both direct utility from consumption as well as information/implementation costs.

CHAPTER V

IDENTIFICATION

So far, we have developed a model of price-competing markets in which consumers have the option of employing a satisficing demand strategy, in addition to the standard fully optimizing consumption strategy. We have demonstrated the existence of a fully stable Nash equilibrium in the overall market game, despite the fact that consumers may be playing a strategy which is *ex ante* only boundedly rational. A significant feature of this equilibrium, in a static analysis, is that it is possible for equilibrium prices to be supported at levels above marginal cost, even in the case where competing firms supply products which are perfect substitutes for each other²⁹. Obviously, the present model is not the only model for which this qualitative result holds.

If firms compete in quantities, rather than prices, for instance, then it is certainly well-known that price above marginal cost can be supported in equilibrium. Likewise, if there is some sort of transportation cost or information cost involved with the process of buying from different sellers, then it will also be the case that prices can be supported above marginal cost, even for perfect substitute goods. Finally, it should also be noted that, since preferences cannot be observed directly, it is also always a possibility that the consumer(s) perceive some relevant difference between products that we might otherwise expect to be close or perfect substitutes. Therefore, even if we were to observe price competition in markets with close substitutes leading to pricing behavior which is well in

²⁹ As mentioned earlier, this result is not in conflict with the uniqueness proof for the single Nash equilibrium in standard Bertrand competition, since the setup of the current model does not satisfy the hypothesis of the proof: instead of a single demand curve which is taken as given in a particular market, there are a collection of *possible* demand curves which may result from differing choice of consumer strategies, even in the face of a given fundamental consumer profile of wealth and preferences.

excess of marginal cost (say, in the market for sugary carbonated sodas, for instance), we would still not be able to consider that observation as definitive empirical evidence of satisficing consumer behavior. Such a market would be qualitatively consistent with both satisficing consumers, as well as with Bertrand competition with product differentiation.

We should certainly point out that, with respect to transportation costs, we hope that the current model is easily differentiable from the literal interpretation of standard transportation cost model, since we have in mind that the satisficing model can explain, among other things, how a consumer might make a purchase decision when faced with an array of similar choices which are all for sale at the same physical location³⁰. A literal interpretation of a travel cost model obviously would not be a valid model for this situation, although interpreting the concepts of distance and travel cost more figuratively, as in the common interpretation of Hotelling's law, might make it harder to differentiate between travel cost models and our satisficing model.

Nonetheless, as previously discussed examples demonstrate, for a *given* utility function which specifies the precise relationship between the various goods available to the consumer (so that, in a hypothetical example, we do not need to admit of any uncertainty whatsoever concerning the psychological "distance" between any of the products), the price profile which prevails in satisficing equilibrium can be strictly higher than the price profile which is possible in optimizing equilibrium.

³⁰ For example, one of my earliest inspirations for this satisficing model involves my own personal experience grocery shopping (specifically, for potato chips). My wife and I had a particular style of potato chips that we would buy on a regular basis. When I went down the potato chip aisle, there were probably dozens of different bags of chips for sale on the same few shelves. However, I found myself repeatedly buying my customary chips, and in doing so, rather automatically ignoring any of the bags of chips – brand, style, and most importantly, *price*. This behavior is absolutely inconsistent with optimizing behavior, and is clearly also inconsistent with the *literal* version of a typical travel cost model (linear city, etc). Furthermore, since I essentially never made any attempt to verify that there might be a low enough price that could coax me into trying a different brand of potato chip, it is probably true that this behavior was not even consistent with the "psychic distance" interpretation of travel cost models, either.

In other words, if we observe an actual market for which the prevailing price is greater than marginal cost, we might not necessarily know whether $P > MC$ because the consumers are satisficing or because they perceive an importance difference between the products supplied by various firms in the market. But we *can* be confident in concluding that, even if $P > MC$ because optimizing consumers perceive a differentiation in the products supplied by competing firms, prices would possibly be even higher still if consumers were to satisfice rather than optimize (holding preferences unchanged).

Features Unique to Satisficing Demand Curves

While looking only at the static qualitative features of a market equilibrium (primarily, the level of price in relation to marginal cost) might give us only a limited ability to differentiate between two hypotheses regarding consumer behavior, we can begin to see a wider range of variation in the predictions of the two models (optimizing demand, satisficing demand) if we begin to consider more dynamic aspects of how price competing markets behave.

First, we can point to a particular characteristic of the “kinked” demand curve which was pointed out long ago by the original authors of a kinked demand curve model (Sweezy, and Hall & Hitch). As previously mentioned, if the demand curve for a product is in fact kinked, so that there is a point of non-differentiability somewhere along it, then marginal revenue for that product will exhibit a jump discontinuity at the kink. This implies that the “kink price” is profit-maximizing for the seller not just for a single value of marginal cost, but for a *range* of values of marginal cost. In turn, this means that the

profit-maximizing price will remain constant even in the face of cost shocks (assuming the shocks are small enough in magnitude).

In light of this, we would expect the prices for *consumer* goods to be somewhat less volatile than prices for inputs (because we would expect consumers to be more likely to satisfice than producers are...more on this later). And in fact, when we compare CPI and PPI for the last 20 years or so, we do see considerably higher price volatility for consumer goods than producer goods (fig. 47).




Figure 47. CPI vs. PPI volatility (Source: St. Louis Fed).

Furthermore, the particular derivation of the satisficing demand curve has additional interesting features not predicted by the original kinked demand curve models. Specifically, satisficing demand curves in particular may exhibit the property that prices may be upwardly flexible but downwardly rigid.

Recall that the satisficing demand curve is composed of two sections, one representing the locus of points in price-quantity space for which the consumer's budget constraint is just binding (under the particular satisficing strategy used to derive said demand curve), and another representing the locus of points for the consumer's worth constraint is just binding for the good corresponding to the particular demand curve in question. Only when the price vector is such that the actual price and the kink price correspond for all goods will it be the case that the consumer is playing a best response (as both conditions for utility maximization are met: condition 1 being that the marginal rate of transformation equals the marginal rate of substitution across all goods, and condition 2 being that all available income spent). In section III, we demonstrated that, for the satisficing demand curves we have described, it can never be profit-maximizing for any firm (with positive marginal costs) to price below the kink price of the market demand curve. It can be profit-maximizing for a firm to price *at* the kink price, assuming that marginal cost falls within the range of marginal revenue at the kink. However, if marginal cost should ever rise *above* that range, then the profit-maximizing price will lie strictly above the kink price.

Assume that a market is initially in satisficing Nash equilibrium, with all firms pricing precisely at the kink price. If the production process experiences relatively small cost shocks, price and output can remain unchanged indefinitely. But if there should happen to be a large *enough* positive cost shock, then profit-maximization will dictate that prices rise in this market, and that quantity falls. However, though the higher price is consistent with profit-maximization among firms, this situation cannot continue to be a Nash equilibrium in the overall market game, because once prices move above the kink

price, consumers are no longer playing a best response strategy, relative to the new price vector.

Specifically, consumers will be consuming along the interior of the upper portion of their demand curve for at least one good, which implies that consumers will have a positive amount of income which is going unspent (since consumption is no longer also taking place along the lower portion of the demand curve). They could, therefore, increase their overall utility by consuming more (i.e., they have a profitable deviation). However, more consumption is not possible under their *existing* satisficing strategy, since the fact that all goods are being consumed along the upper portion of the demand curve means that the worth constraint binds for all goods at quantities less than what would be needed to exhaust all income.

Of course, one way that consumers could alter their strategy so as to begin playing a best response to the new (higher) price vector would be to switch from satisficing to optimizing. But consumers could *also* play a best response to the new price vector simply by altering their choice of μ . If the higher prices make the old worth constraint so restrictive that not all income can be spent under the old decision rule, then by relaxing the worth constraint, consumers can improve their total utility under the new price vector. Of course, by downwardly revising μ , and effectively selecting a new satisficing demand strategy, consumers will also end up changing the location of the kink on their demand curves for all goods. Specifically, we could imagine that consumers, realizing that each dollar does not go as far as it used to due to higher prices, have lowered expectations for how much utility each dollar ought to bring them, and so keep revising their estimate of μ until their new demand curves have kinks located precisely at

the prevailing price level. In this case, a new Nash equilibrium can be reached, where the higher prices are still profit-maximizing, but satisficing consumers are still playing a best response to those prices. If and once this happens, then not even an equally large *negative* cost shock would be sufficient to return prices to their initial equilibrium levels. Again, assuming that the negative shock occurs *after* the adjustment of consumer satisficing strategies, then the lower (original) prices would then be located along the lower portion of the new satisficing demand curves. As a result, even if costs were to return to normal, prices would be expected to remain elevated for some time, since consumers have taken the chance to reevaluate their assessment of μ , or of how much they expect each dollar of income should be able to buy for them, in utility terms.

Here, we present a numerical example of this “upwardly flexible, downwardly rigid” price characteristic. This shows the initial satisficing equilibrium for the utility function

$$U(x, y) = 2000 \ln(.2 x + 1) + 2000 \ln(.2 y + 1)$$

where consumer wealth is $W = 100$ and constant marginal cost is $c = 3$, and the particular satisficing strategy being used by the consumer is $\mu = 10$ and $\Pi = (x, y)$. As previously shown, the Nash equilibrium price vector for this market is $(30, 30)$, which does correspond to the kink price on each demand curve (fig. 48).




Figure 48. Initial Equilibrium at price vector (30, 30).

Assume that there is a cost shock which increases the marginal cost of both firms, and leads to a situation where the higher price vector $(P_x, P_y) = (35, 35)$ now becomes profit-maximizing for each firm. Until the consumer change her strategy, it must be the case that both firms' price of 35 will be above the consumer's kink price. As we have mentioned, although this situation is consistent with profit-maximization for each firm, this situation could not be a Nash equilibrium in the market game. The consumer's unspent income implies that that condition 2 for utility maximization is not met, and the consumer's satisficing strategy cannot be a best response to the price vector.

If we do in fact assume that the price vector (35, 35) is profit-maximizing, then obviously neither firm will have an incentive to alter price. The only avenue remaining for adjustment to equilibrium is for the consumer to alter her demand strategy.

Specifically, she can lower the value of μ that she is using, which will lower the marginal utility per dollar threshold that potential purchase need to satisfy, and therefore will make her worth constraints for all goods less restrictive. This will enable her to spend all of her income, while continuing to employ a satisficing strategy (fig. 49).




Figure 49. By decreasing μ from 10 to 9, the consumer can maximize utility under the price vector (35, 35), and doing so moves the kink price up to meet the price vector. (Demand curve for Y is symmetrical).

The consumer can exercise a profitable deviation from her original satisficing strategy by revising her estimate μ downward. This lowers the threshold for potential purchases to satisfy, and enables her to exchange 100% of her available wealth for goods.

Furthermore, this alters the demand curve for each firm, and once again aligns the profit-maximizing price vector with the kink price. Left-hand marginal revenue at a price of 35 will be higher under the new demand curve than marginal revenue was at the same price on the original demand curve. Therefore, if a price of 35 was profit-maximizing before the revision in μ , it will continue to be profit-maximizing after the change in μ . And, importantly, even if costs were to subsequently return to their original levels, neither firm will have any incentive to lower its price.

Through this example, we can begin to see some hints of a fundamental asymmetry between upward and downward price movements. That is, using a given satisficing strategy as a starting point, prices below the kink level will likely never persist

for long, since firms can always strictly increase profit by raising price. But prices above the kink might persist for a while, if costs shocks are large enough. In this case, the satisficing consumer's best response is to alter her strategy in such a way as to "lock in" these higher prices, by altering μ , and effectively shifting the demand curve kink upward to meet the high prices. Therefore, if there is some temporary cost shock which initially raises prices above the kink, when the consumer revises μ in response, there may be something of a "ratchet up" effect. Once the location of the demand kink has changed, then, as long as the consumer is playing the equilibrium satisficing strategy, there will be no immediate incentive for firms to lower prices, even if costs should subsequently fall back to their original levels. This is true despite the fact that the firms are assumed to be competing in prices. In other words, it is at least possible that there may exist a structural bias in markets with satisficing consumers, in which prices are upwardly flexible, but downwardly rigid.

This example provides us with another clear prediction about consumer behavior which cannot easily be matched by existing models, and therefore gives us another potential way in which the satisficing demand model can be differentiated from the predictions of existing models.

With respect to this predicted possibility of upwardly flexible, downwardly rigid output prices in response to cost shocks, we can look to the literature which investigates the relationship between movements in the price of crude oil and resulting movements in the price of gasoline. In various studies, Bacon (1991), Karrenbock (1991), Borenstein, et al. (1997), Balke, et al. (1998), and Brown and Yucel (2000) all find empirical evidence that gasoline prices do react asymmetrically to increases and decreases in the price of

crude oil. In particular, Brown and Yucel note, conspicuously, that “In addition, no formal theory relating market power to asymmetry has been tested (to our knowledge).” It is encouraging to see that the present model is potentially consistent with the observed empirical evidence in this area, and also a candidate to provide a theoretical structure which may begin to explain this previously unexplained market behavior.

Thoughts on Monetary Policy

Also worth emphasizing is the fact, when consumers behave in the manner described by this satisficing demand model, it will always be the case that their beliefs about the utility value of a dollar (i.e., μ) are correct, in equilibrium. Most of the time, this implies that consumer beliefs about μ are self-fulfilling prophecies. That is, consumers each have their own utility function, and their own estimate of μ , and then firms select their own prices in such a way as to maximize profit. This then implies that consumers buy the amount of each good for which it will be case that the marginal utility of their final dollar of expenditures (for each good) is precisely equal to μ . Since there is nothing which requires any consumer to select the “correct” μ for *every* price vector, so that $\mu=\lambda$ (where λ is the Lagrangian multiplier of the fully rational constrained optimization problem, or the shadow value of wealth in utility terms) for *every* price vector, it is at least worth considering the question of to what extent changes in the money supply effect real output. Since we assume that consumers may find it reasonable to satisfice rather than truly optimize in the static setting of their single-period budget problems, it is not a forgone conclusion that these consumers would react fully rationally to changes in the money supply by altering their μ 's in precisely the “correct” value

based on the size of the change in the money supply. Although the current model does not include the necessary structure for assessing the impact of monetary policy (we have no financial assets such as bonds which may be bought or sold by the central bank, no structural equation for how aggregate consumption and prices respond to changes in the money supply, and no interest rates) the fact that consumers *always* behave in a way which confirms their own μ ³¹ is at least suggestive of the possibility that there will be rational expectations equilibria in which money is non-neutral in the real economy.

Business Cycles

A final area where we expect to be able to find implications of the satisficing demand model which are distinct from those of existing models concerns the nature of business cycles. As in Adam (2005), we hypothesize that one effect of boundedly rational consumption behavior might be cyclical behavior of macro level variables such as output, employment, and inflation³². Although we have not yet been able to develop a satisfactory specific mathematical or numerical example, we can at least try to sketch the intuition underlying our hypothesis that the satisficing demand model is consistent with the observed stylized facts of the typical United States business cycle.

The main point here is that the satisficing demand model illustrates that all price-competing markets exhibit multiple Nash equilibria: there is (usually) one Nash equilibrium where all consumers play an optimizing consumption strategy, and then there

³¹ Or, more precisely, the fact that firms never find it profit-maximizing to price below the market kink price.

³² I wish to acknowledge a debt of gratitude to Dr. Klaus Adam, whose excellent paper, cited here, was perhaps the earliest conceptual inspiration for this dissertation.

is an entire class of additional Nash equilibria in which consumers play a satisficing demand strategy. The specific realization of the set of satisficing Nash equilibria depends, in part, upon the group of consumers' estimate of the utility value of a dollar at any point in time.

We hypothesize that many of the features of the US business cycle can be explained in terms of the process of endogenous adjustment or switching back and forth between these two types of consumption behaviors. The gist of the story we are trying to tell here, and the fundamental feature of the model which causes us to suspect that this model may be a good explanation for the existence of business cycles, is that when all consumers are optimizing, then individual consumers have an incentive to satisfice instead, and when all consumers are satisficing, then individual consumers have an incentive to optimize.

This story, as we are in the process of trying to tell it, depends upon at least two assumptions. First, all consumers are price takers. And second, that consumers have a lexicographic preference for strategy-payoff combinations which values utility and simplicity. That is, they place the foremost importance on selecting a strategy which provides the highest level of direct utility. But, if two distinct strategies were to offer the same amount of direct utility in a given situation, then they strictly prefer to use the simpler strategy. Specifically, we will assume that, for a given price vector, every consumer would prefer to employ a satisficing strategy rather than an optimizing strategy, if that strategy could be guaranteed to select a bundle which yields the same amount of direct utility as optimizing.

Recalling the definition of a Calibrated Satisficing Strategy (CSS), we know that consumers will always have at least one such strategy for any price vector/income combination for which they have previously optimized. Here, the distinction between the two involves the manner in which consumers utilize their previous information. True optimization, arguably, treats all previous information about prices as irrelevant; to find a truly optimal bundle yesterday, it was necessary to compare every possible margin of consumption against every other margin of consumption. To find a truly optimal bundle today, it is also necessary to compare every possible margin of consumption against every other. Even if no prices change from period to period, true optimization requires, at a minimum, that consumers verify that the level of *every* price has remained unchanged. Satisficing, on the other hand, allows the consumer to ignore the prices of goods which are never actually purchased (along the principal consumption path), and to ignore the comparison of marginal utility levels for quantities higher than the consumption path selects.

In other words, optimizing is more complicated because it requires that the consumer consider, at least implicitly, every bundle in the budget set, while satisficing only requires purchases along the consumption path to be considered. If consumers have a preference for simplicity, then it can never be a best response for a price-taking consumer to optimize under a given price vector, *if* that consumer has the information available necessary to employ a CSS for that vector. Based on this, it can be said that the employment of a satisficing strategy can be thought of as a type of “memory” on the part of consumers, since the parameters of a CSS are derived from a previous period where the consumer actually optimized.

Roughly, we can begin to describe the intuition of the cyclical behavior of markets as a dynamic switching between satisficing and optimizing consumption regimes. If we start in a situation where consumers all optimize, we get one price vector and corresponding level of output. But, if this price vector were to persist through a certain period of time, price-taking consumers each individually have an incentive to “play from memory” rather than to truly optimize, and thereby utilize the costly information which was obtained in previous periods of optimizing. As more and more consumers switch from an optimizing to a satisficing (CSS) consumption regime, firms have more and more incentive to increase price, as a higher proportion of satisficing consumers means that their demand curve becomes less elastic, and they lose fewer sales at higher prices than they would under a fully optimizing set of consumers. Competition amongst firms therefore leads to higher prices once consumers have begun to satisfice.

If we assumed that more complicated strategies involved an explicit resource cost C to employ, this would strengthen our results, but the lexicographic preference for simplicity is the weakest assumption which supports the following conclusion. In addition, we shall here assume that there are two firms which supply the market with perfect substitutes, but many individual consumers who are each price-takers (though their joint behavior may be described by a positive representative consumer). We will try to motivate our story using myopic Nash equilibrium in each stage of the game as our equilibrium concept (ignoring for now any more complicated multi-period effects on consumer or firm strategies, and ignoring more dynamic notions of what an equilibrium is in this multi-period market game).

Expansionary/Inflationary Phase

We begin our description of a generic business cycle at period 0, where all consumers are playing an optimizing demand strategy, and both price-competing firms charge the Nash equilibrium price which is equal to their (identical) constant marginal cost of production. The prevailing price vector is \mathbf{P}^{MC} .

In period 1, this same strategy profile cannot continue to be a Nash equilibrium. Consumers, having optimized under \mathbf{P}^{MC} in period zero, now have all the information that they need to play a CSS (with respect to \mathbf{P}^{MC}) in period 1, and given that the CSS does not require them verify that every price has remained unchanged (and doesn't require them to perform a complete optimization over again in period 1 which they have already performed), CSS is lexicographically preferable to optimizing, for *every* individual consumer (and the CSS is even more preferable if there is an information cost C which must be paid in order for consumers to fully optimize). As a result, no consumer can be playing a best response in period 1, under these assumptions, if the price vector is still \mathbf{P}^{MC} , and the consumer is optimizing. Each consumer in period 1 would prefer instead to “play from memory,” and, through the employment of a CSS, just buy the same things in period 1 that they bought in period 0.

However, even though all consumers would be playing a best response in each stage if they optimized in period zero, and then played $\text{CSS}(\mathbf{P}^{\text{MC}})$ in period 1, this situation cannot be a Nash equilibrium either. If all consumers were playing their preferred satisficing strategy in period 1, then the price vector \mathbf{P}^{MC} would not be profit-maximizing for either firm. Instead, once all consumers (or even just a positive fraction

of all consumers) have begun to satisfice rather than optimize, then both firms have an incentive to raise their price above marginal cost.

We hypothesize that what might actually happen in this case is that, having exhibited a standard Bertrand equilibrium outcome in period 0, the market will slowly, over time, evolve from that Bertrand equilibrium to a satisficing equilibrium. Each subsequent period after period 0, a fraction of consumers will satisfice, allowing firms to increase their prices a little bit. Higher prices lead to two reinforcing effects. The first effect results from the fact that slightly higher prices mean, holding income constant, the worth constraints will begin to bind for both goods before the budget constraint does. As in the example above concerning upwardly flexible but downwardly rigid prices, this means that each consumer has an incentive to revise her estimate of μ downward, thereby giving herself a less restrictive decision rule which allows her to spend all of her income. The second reinforcing effect is that, as prices rise, so does marginal revenue product³³, and so, ultimately, so must wages. As wages increase, each consumer has more income to spend on each of the two available goods. This also must eventually lead to a downward revision of μ , all else equal.

The reason that both of these effects are “reinforcing” is that, as μ gets lower and lower, then the ability of the price-competing firms to raise their prices becomes greater and greater. As a result, we hope to be able to show that this inflation cycle can continue in the economy for some time. Notably, this seems like it will occur even in the absence of any sort of monetary authority or even an explicitly modeled money supply.

³³ My current attempts at mathematically modeling this have been troublesome. I’m having some trouble justifying how or why quantity would increase as price increases, which I suspect must be a key part of this “expansionary phase” story.

An important difference between the process described here and a standard model in which information costs C allow for prices to exceed MC in what would otherwise be standard Bertrand competition is that, in the standard approach, prices can exceed MC by at most C . But in this setup, if there were an information cost C which was required of any consumer who wished to optimize in any given period, then C would constrain only the amount by which prices could increase in any *one* period. So long as prices increase by less than C in any period, no price-taking consumer has an incentive to optimize instead of satisfice. But the cumulative effect of several consecutive periods in which satisficing consumer behavior has allowed prices to slowly rise may, in principle, be much larger than C . The ultimate equilibrium price level will depend upon consumer wealth, and on consumers' estimates of μ , and not on the size of C . Information costs, therefore, can only place a constraint on the *rate* of growth in prices, not, ultimately, on the level of prices. As a result, we hope to be able to show that satisficing consumption behavior may lead to long periods where feedback between rising prices, rising marginal revenue product, and rising wages, allow mild—but persistent— inflation to take place in an economy.

Contractionary Phase

While we expect to be able to show that the inflationary phase of the business cycle can persist for some time in this model, it wouldn't be a very satisfactory model of business cycles if that expansionary phase never actually ended. We hope to motivate the narrative of how economies can find themselves in period recessions as being basically the story of how the economy switches from a satisficing to an optimizing equilibrium.

Having exhibited a period of sustained inflation, one way³⁴ in which an economy might find itself in recession is through the actions of some sort of central monetary authority (which we have not needed to consider until now). If that central monetary authority notices that inflation is proceeding too rapidly, it may feel a policy intervention designed at breaking the wage-price feedback alluded to in the preceding section is warranted. By acting to increase interest rates and affect the real economy through the investment channel, for instance, the monetary authority can indeed achieve their goal of reducing inflation. But at the same time, they may also end up leading to the unwinding of the process described above. As real demand falls in interest-rate sensitive sectors, so too will employment in those sectors. As employment falls, income falls. Both this fall in income, as well as pessimistic expectations about the future can probably be expected to lead to an increase in optimizing behavior, as unemployed and employed-but-skittish workers become more careful with their money across the board. But even among those consumers who continue to satisfice, falling income and falling employment will lead to upward revisions in μ . Higher μ 's are not able to support prices as high as they what prevailed prior to the start of the recession, so firms that do not lower their prices will see their sales fall. This will lead to these firms being unable to justify keeping their own employment at pre-recession levels, and so will lead to further layoffs, less income, and a combination of more consumer optimizing and higher μ 's from those consumers who continue to satisfice. Both of these last effects will mean that some workers whose wage

³⁴ There are other recession triggers that I've considered, including some unforeseen shock with darkens expectations about the future across the board, as well as the idea that wages do not quite keep pace with inflation for the all of the expansionary phase, so that eventually consumers find themselves in a situation where their real wage is low enough, or their debt levels are high enough, that they no longer find it a best response to satisfice. The central bank story just seems like the easiest story to actually use to sketch out this portion of the business cycle at the moment.

was justifiable to the employing firm during a time when every consumer was satisficing, and when prices were rising, will no longer pay for themselves, and so will be laid off, furthering the downward cycle.

The recession reaches its trough when all consumers are once again optimizing, prices are low, but the economy is once again in Bertrand equilibrium.

Once again, this was merely an attempt to sketch out the intuition underlying our attempt to describe business cycles in light of the satisficing demand model. There is obviously a great deal of work yet to do on developing a mathematical example of such a business cycle, but the central hypothesis is that the transition back and forth between satisficing and optimizing equilibria can be a compelling way to explain why it is that economies do experience booms and recessions.

CHAPTER VI

CONCLUDING REMARKS

We have begun to develop a rigorous model of consumer behavior which is based, not on a standard optimization algorithm consistent with full rationality and subgame perfection, but based instead upon a simplifying rule of thumb, which we model as satisficing behavior. We have begun to explore the consequences of this type of boundedly rational consumer behavior, within the context of a specific, somewhat stylized model. We have demonstrated that it is possible for ex ante boundedly rational behavior to nonetheless form part of a Nash equilibrium strategy profile. In other words, the particular strategy choices of satisficing consumers may end up being rationalized by the strategic behavior of firms which compete in prices to supply the market with output. In strategic settings, what is “rational” and what is “irrational” is, in some sense, a murkier proposition than the same question posed in a static setting. Our main hope is that we have, at a minimum, made a reasonable case that we, as economists, ought to consider broadening the scope of models which we consider to provide a compelling description of consumer behavior, and attempt to explore this and related models in the future as thoroughly as possible. We further hope that such models may ultimately prove to be a useful and appealing source of new insight into the way in which consumers and firms behave within the context of the strategic interaction of markets, and new insight into the types of outcomes for those markets that can be supported in equilibrium.

To the extent that this model does provide novel insight, we do wish to make clear that our view is that this insight ought to be properly viewed as a complement, rather than a substitute, to the many well-known and valuable lessons of standard models, which rightly form the foundation of consumer theory in particular, and microeconomics in general.

We also recognize that this work also potentially opens up several new areas which appear promising for future research. First and foremost, we must admit that there is still plenty of room for exploring the implications of the most basic form of this model. Furthermore, there is similar room for continued exploration of the implications of each line of discussion laid out in Chapter V, Identification. Undoubtedly, the investigation of the implications of satisficing consumption behavior is quite preliminary in this document, but we are hopeful that future work will be able to make more progress along these lines.

We also suspect that there is future opportunity for successful and compelling investigation along several lines of testing, including a more thorough investigation of the relationship of gasoline and crude oil prices, the implications for monetary policy, and perhaps even direct testing, through an appropriate designed experimental study, of the underlying structures of the model.

We close with two final possibilities for future work, which we have not previously touched on. First, since the behavior of consumers, and consequently the location of market Nash equilibria, depends crucially on the estimate μ of the utility value of a dollar, we have reason to believe that there may be further insight to be gained by extending the model to investigate the manner in which income, or more specifically, the

distribution of income, might impact consumer and market behavior, in a positive sense rather than in a normative one.

Finally, we also suspect that some version of a satisficing model might do well in explaining why we seem to observe persistent cost increases in certain sectors, such as health care or higher education, which consistently outpace the rate of income growth or inflation. We are especially hopefully of developing future fruitful work in these areas, since it we suspect that these two sectors in particular are likely to have a close correspondence with the setup of our current model; specifically, consumers in those markets potentially do not compare the perceived utility of those goods against their true opportunity cost, as calculated by considering every other available margin of consumption. In other words, if consumers essentially take it as given that they are going to purchase health care or higher education, we expect that we may find a close correspondence between their actual behavior in those particular markets, and the behavior predicted by our current model.

APPENDIX

SAMPLE MATHEMATICA CODE

The code appearing on the following pages is a sample of the actual program used to generate most of the graphics included in this work, as well as to calculate the shape and location of the relevant functions, loci, equilibria, and vector fields. The program itself was written and executed using Wolfram *Mathematica*® 7 for Students.

Sample programming related to the calculation of the relevant constituent parts of a satisficing demand system for a single consumer in a two-good economy can be found on pages 204 through 210. Sample code used for the more complicated problem of finding aggregated market demand curves in settings where there are multiple heterogeneous consumers, and for calculating Calibrated Satisficing Strategies (CSS), can be found on pages 211 through 228.

Single Consumer Case

```
"Graph Slider Names in [Brackets] Denote Slidable Display
  Parameters (horizontal range, vertical range, etc.)";

Quiet[ClearAll["`*"]]

"User-Defined Functions and Commands";

ColumnSwap[list_] :=
  Drop[Transpose[Insert[Transpose[list], list[[All, 1]], 3]], {}, {1}];

Default[PlotTranspose, 2] = Null;
Default[PlotTranspose, 3] = Null;
Default[PlotTranspose, 4] = Null;
Default[PlotTranspose, 5] = None;
Default[PlotTranspose, 6] = Automatic;
PlotTranspose[plot_, plabel_.,
  xlabel_., ylabel_., gridlines_., plotstyle_.] := (
  Points = Cases[plot, {_?NumericQ, _?NumericQ}, {6}];
  Pause[TransposeDelay];
  ListLinePlot[ColumnSwap[Points], AxesLabel -> {xlabel, ylabel},
    PlotLabel -> plabel, GridLines -> gridlines, PlotStyle -> plotstyle]

"Structural Parameter Values";
a = 3000; b = 1 / 10; cc = 1 / 10; d = 1;

Wlinit = 100;
mulinit = 10;
xCost = 3;
yCost = 3;

"Display Parameter Values, non-structural variables";

TransposeDelay = 3;

Group0PMax = 40;
Group1PMax = 40;
Group2PMax = 40;
```

```

Group3PMax = 40;
Group4PMax = 40;

Group0QMax = 25;
Group1QMax = 25;
Group2QMax = 25;
Group3QMax = 25;
Group4QMax = 25;

PMaxMarket = 40;
QMaxMarket = 50;

PmaxBR = 40;

EpilogPx = 35;
EpilogPy = 35;

"Utility Functions and Partial Derivatives";

Utility1[x_, y_] = a Log[b x + cc y + d];

Ux1[x_, y_] = Derivative[1, 0][Utility1][x, y];
Uy1[x_, y_] = Derivative[0, 1][Utility1][x, y];
Print["Ux1[x,y] = ", FullSimplify[Ux1[x, y]]]
Print["Uy1[x,y] = ", FullSimplify[Uy1[x, y]]]

"Principal Path of  $\Pi$  Field
LHS=RHS -> {LHS,RHS}";

piPP1 = {y, x};

"Satisficing Critical Values";

"Consumer 1";

clxcrit0[px_, mul_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y]},

```

```

Max[0, x /. First@Solve[Ux1[x, y1] == mul px, x]];

clycrit0[py_, mul_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x]},
    Max[0, y /. First@Solve[Uy1[x1, y] == mul py, y]];

clxcrit1[px_, py_, mul_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y], y2 =
    y /. First@Solve[Uy1[x, y] == py * mul && piPP1[[1]] == piPP1[[2]], y, x]},
    Max[0, x /. First@Solve[Ux1[x, y1] == mul px, x],
    x /. First@Solve[Ux1[x, Max[0, y2]] == mul px, x]];

clycrit1[px_, py_, mul_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x], x2 =
    x /. First@Solve[Ux1[x, y] == px * mul && piPP1[[1]] == piPP1[[2]], x, y]},
    Max[0, y /. First@Solve[Uy1[x1, y] == mul py, y],
    y /. First@Solve[Uy1[Max[0, x2], y] == mul py, y]];

clxcrit2[px_, py_, mul_, W1_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y],
    y2 = y /. {y -> clycrit0[py, mul]}},
    Max[0, x /. First@Solve[W1 == px x + py y1, x],
    x /. First@Solve[W1 == px x + py y2, x]];

clycrit2[px_, py_, mul_, W1_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x],
    x2 = x /. {x -> clxcrit0[px, mul]}},
    Max[0, y /. First@Solve[W1 == px x1 + py y, y],
    y /. First@Solve[W1 == px x2 + py y, y]];

"Graph of Consumer 1 Constraints in Consumption Space";

Manipulate[Plot[{Flatten[Solve[W1 == py y + px x, y]][[1]][[2]],
  Flatten[Solve[Ux1[x, y] == mul px, y]][[1]][[2]], Min[clycrit0[py, mul],
  Flatten[Solve[piPP1[[1]] == piPP1[[2]], y]][[1]][[2]]]},
{x, 0, 50}, PlotRange -> {{0, 50}, {0, 50}}, AxesLabel -> {"X", "Y"},
PlotLabel -> "Consumption Space (X Constraints)",
{{px, 10, Subscript[P, x]}, 0, 50}, {{py, 10, Subscript[P, y]}, 0, 50},

```

```

{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{mul, mulinit, Subscript[μ, 1]}, 0, 100}}

Manipulate[
Plot[{Flatten[Solve[W1 == py y + px x, x]][[1]][[2]], Min[c1xcrit0[px, mul],
  Flatten[Solve[piPP1[[1]] == piPP1[[2]], x]][[1]][[2]]],
  Flatten[Solve[Uy1[x, y] == mul py, x]][[1]][[2]]], {y, 0, 50},
PlotRange → {{0, 50}, {0, 50}}, AxesLabel → {"Y", "X"},
PlotLabel → "Consumption Space (Y Constraints)",
{{px, 10, Subscript[P, x]}, 0, 50}, {{py, 10, Subscript[P, y]}, 0, 50},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{mul, mulinit, Subscript[μ, 1]}, 0, 100}}

"Demand Functions";

      "Individual Demand Functions";

Default[c1Dx, 3] = mulinit;
Default[c1Dx, 4] = W1init;
c1Dx[px_, py_, mul_., W1_.] :=
  Max[0, Min[c1xcrit1[px, py, mul], c1xcrit2[px, py, mul, W1]]]

Default[c1Dy, 3] = mulinit;
Default[c1Dy, 4] = W1init;
c1Dy[px_, py_, mul_., W1_.] :=
  Max[0, Min[clycrit1[px, py, mul], clycrit2[px, py, mul, W1]]]

"Plotting Individual Demand Curves and Constituent Loci";

Manipulate[Plot[c1xcrit0[px, mul], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 First Critical X",
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
  {{PxMax, Group0PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group0QMax, "[QxMax]"}, 0, 100}}

Manipulate[Plot[clycrit0[py, mul], {py, 0, PyMax},

```

```

PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
PlotLabel → "Consumer 1 First Critical Y",
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{PyMax, Group0PMax, "[PyMax]"}, 0, 100},
{{Ymax, Group0QMax, "[QyMax]"}, 0, 100}}

Manipulate[Plot[clxcrit1[px, py, mul], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 Worth Locus (X)",
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100}, {{py, 5, Subscript[P, y]}, 0, 25},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[Plot[clxcrit2[px, py, mul, W1], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 Budget Locus (X)",
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[
Plot[{clxcrit1[px, py, mul], clxcrit2[px, py, mul, W1]}, {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 Combined Loci (X)",
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[PYG = py; Dx = Plot[c1Dx[px, py, mul, W1], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 Demand Curve X (transposed) for Py = "[py]],
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, ]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50},

```

```

{{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[Plot[clycrit1[px, py, mu1], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 Worth Locus (Y)",
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

Manipulate[Plot[clycrit2[px, py, mu1, W1], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 Budget Locus (Y)",
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

Manipulate[
Plot[{clycrit1[px, py, mu1], clycrit2[px, py, mu1, W1]}, {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 Combined Loci (Y)",
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

Manipulate[PXG = px; Dy = Plot[c1Dy[px, py, mu1, W1], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 Demand Curve Y (transposed) for Px = "[px]],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50},
  {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

```

```

"Standard Presentation of Demand Curves";

PlotTranspose[Dx,
  "Good X Demand Curve for Py = " [PYG], "X", Subscript[P, x]]
PlotTranspose[Dy, "Good Y Demand Curve for Px = " [PXG],
  "Y", Subscript[P, y]]

"Firm Profit Functions";

Manipulate[
  Quiet[ProfitX1 = Quiet[Plot[(px - xc) c1Dx[px, py, mu1, W1], {px, 0, PxMax},
    PlotRange -> {{0, PxMax}, {0, xProfitmax}},
    AxesLabel -> {Subscript[P, x], "ProfitX"},
    PlotLabel -> "Firm X Profit Function for Py =" [py],
    WorkingPrecision -> MachinePrecision]],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100}, {{py, yCost, Subscript[P, y]},
  0, 50}, {{xc, xCost, "Firm X Cost"}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{PxMax, PMaxMarket, "[PxMax]"}, 0, 100},
  {{xProfitmax, (W1init), "[X Profit Max]"}, 0, 2 (W1init)},
  ContinuousAction -> None]

Manipulate[Quiet[
  ProfitY1 = Plot[(py - yc) c1Dy[px, py, mu1, W1], {py, 0, PyMax}, PlotRange ->
    {{0, PyMax}, {0, yProfitmax}}, AxesLabel -> {Subscript[P, y], "ProfitY"},
  PlotLabel -> "Firm Y Profit Function for Px = " [px],
  WorkingPrecision -> MachinePrecision]],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{px, xCost, Subscript[P, x]}, 0, 50}, {{yc, yCost, "Firm Y Cost"}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{PyMax, PMaxMarket, "[PyMax]"}, 0, 100},
  {{yProfitmax, (W1init), "[Y Profit Max]"}, 0, 2 (W1init)},
  ContinuousAction -> None]

"Best Response Functions";

Default[BestResponseX, 2] = xCost;
Default[BestResponseX, 3] = mu1init;
Default[BestResponseX, 4] = mu2init;

```

Multiple Consumer Case

```
"Run each cell in order once to calibrate, disregarding
  output. Then run each cell (in order from 1st) a second time"

"Graph Slider Names in [Brackets] Denote Slidable Display
  Parameters (horizontal range, vertical range, etc.)";

Quiet@ClearAll["`*"]

"User-Defined Functions and Commands";

ColumnSwap[list_] :=
  Drop[Transpose[Insert[Transpose[list], list[[All, 1]], 3]], {}, {1}];

Default[PlotTranspose, 2] = Null;
Default[PlotTranspose, 3] = Null;
Default[PlotTranspose, 4] = Null;
Default[PlotTranspose, 5] = None;
Default[PlotTranspose, 6] = Automatic;
Default[PlotTranspose, 7] = 1;
PlotTranspose[plot_, plabel_., xlabel_.,
  ylabel_., gridlines_., plotstyle_., delayfactor_.] := (
  Points = Cases[plot, {_?NumericQ, _?NumericQ}, {6}];
  Pause[delayfactor * TransposeDelay];
  ListLinePlot[ColumnSwap[Points], AxesLabel -> {xlabel, ylabel},
    PlotLabel -> plabel, GridLines -> gridlines, PlotStyle -> plotstyle]

Default[IsEqual, 3] = IEtolerance;
IsEqual[q1_, q2_, tolerance_.] := If[Abs[q1 - q2] <= tolerance, True, False]

CSS1Slope[px_, py_] :=
  If[OCB1[px, py][[1]] ≠ 0, c1A = OCB1[px, py][[2]] / OCB1[px, py][[1]],
    Vertical1 = True; && Print["Consumer 1 Principal Path is Vertical"];]

CSS2Slope[px_, py_] :=
  If[OCB2[px, py][[1]] ≠ 0, c2A = OCB2[px, py][[2]] / OCB2[px, py][[1]],
    Vertical2 = True; && Print["Consumer 2 Principal Path is Vertical"];]

CSS[px_, py_] := (
```



```

Unprotect[c1A, c2A, lambda1, lambda2];

CSS1Slope[px, py];
CSS2Slope[px, py];

c1CSSmu1 = Ux1[OCB1[px, py][[1]], OCB1[px, py][[2]]] / px;
c1CSSmu2 = Uy1[OCB1[px, py][[1]], OCB1[px, py][[2]]] / py;
If[c1CSSmu1 == c1CSSmu2, Print["Lambda 1 = ", lambda1 = c1CSSmu1],
  Print["Error: Lambda mismatch"]];

c2CSSmu1 = Ux2[OCB2[px, py][[1]], OCB2[px, py][[2]]] / px;
c2CSSmu2 = Uy2[OCB2[px, py][[1]], OCB2[px, py][[2]]] / py;
If[TrueQ[c2CSSmu1 == c2CSSmu2], Print["Lambda 2 = ", lambda2 = c2CSSmu1],
  Print["Error: Lambda mismatch"]];

Protect[c1A, c2A, lambda1, lambda2];

)

"Principal Path of Each  $\Pi$  Field
LHS=RHS -> {LHS,RHS}";

If[Vertical1 == True, piPP1 = {0, x};, piPP1 = {y, c1A * x};];
If[Vertical2 == True, piPP2 = {0, x};, piPP2 = {y, c2A * x};];

"Structural Parameter Values";
a = 3000; b = 1 / 10; c = 1 / 10; d = 1;

Quiet[lambda1 = 10];
Quiet[lambda2 = 10];
Quiet[c1A = 1];
Quiet[c2A = 1];
Quiet[Vertical1 = False];
Quiet[Vertical2 = False];

Wlinit = 150;

```

```

W2init = 150;
mulinit = lambda1;
mu2init = lambda2;
xCost = 3;
yCost = 3;

"Display Parameter Values, non-structural variables";

IEtolerance = 0.00001;
TransposeDelay = 1;
BRDelayFactor = 1;

Group0PMax = 40;
Group1PMax = 40;
Group2PMax = 40;
Group3PMax = 40;
Group4PMax = 40;

Group0QMax = 25;
Group1QMax = 25;
Group2QMax = 25;
Group3QMax = 25;
Group4QMax = 25;

PMaxMarket = 40;
QMaxMarket = 50;

PmaxBR = 40;

EpilogPx = 35;
EpilogPy = 35;

"Utility Functions and Partial Derivatives";

Utility1[x_, y_] = a Log[b x + d] + 2 a Log[2 c y + d];

Ux1[x_, y_] = Derivative[1, 0][Utility1][x, y];
Uy1[x_, y_] = Derivative[0, 1][Utility1][x, y];
Print["Ux1[x,y] = ", FullSimplify[Ux1[x, y]]]
Print["Uy1[x,y] = ", FullSimplify[Uy1[x, y]]]

```

```

Utility2[x_, y_] = 2 a Log[2 b x + d] + a Log[c y + d];

Ux2[x_, y_] = Derivative[1, 0][Utility2][x, y];
Uy2[x_, y_] = Derivative[0, 1][Utility2][x, y];
Print["Ux2[x,y] = ", FullSimplify[Ux2[x, y]]]
Print["Uy2[x,y] = ", FullSimplify[Uy2[x, y]]]

"Optimizing Demand Functions";

Print["Consumer 1 Optimizing Consumption Bundle = ",
OCB1[px_, py_] = With[{xd =
  x /. With[{y1 = y /. First@Solve[(Ux1[x, y]) / px == (Uy1[x, y]) / py, y]},
    First@Solve[Wlinit == x px + y1 py, x]}, yd =
  y /. With[{x1 = x /. First@Solve[(Ux1[x, y]) / px == (Uy1[x, y]) / py, x]},
    First@Solve[Wlinit == x1 px + y py, y]}],
  {Min[Wlinit / px, Max[0, xd]], Min[Wlinit / py, Max[0, yd]]}] //
FullSimplify]

Print[]

If[Length@Solve[(Ux1[x, y]) / px == (Uy1[x, y]) / py, y] == 1 &&
  Length@Solve[Wlinit == x px + y1 py, x] == 1 &&
  Length@Solve[(Ux1[x, y]) / px == (Uy1[x, y]) / py, x] == 1 &&
  Length@Solve[Wlinit == x1 px + y py, y] == 1,
Print["No Multiple Solutions (Consumer 1)"],
Print["Check for Multiple Solutions! (Consumer 1)"]]

Print[]

Print["Consumer 2 Optimizing Consumption Bundle = ",
OCB2[px_, py_] = With[{xd =
  x /. With[{y1 = y /. First@Solve[(Ux2[x, y]) / px == (Uy2[x, y]) / py, y]},
    First@Solve[W2init == x px + y1 py, x]}, yd =
  y /. With[{x1 = x /. First@Solve[(Ux2[x, y]) / px == (Uy2[x, y]) / py, x]},
    First@Solve[W2init == x1 px + y py, y]}],
  {Min[W2init / px, Max[0, xd]], Min[W2init / py, Max[0, yd]]}] //
FullSimplify]

```

```

Print[]

If[Length@Solve[(Ux2[x, y]) / px == (Uy2[x, y]) / py, y] == 1 &&
  Length@Solve[W2init == x px + y1 py, x] == 1 &&
  Length@Solve[(Ux2[x, y]) / px == (Uy2[x, y]) / py, x] == 1 &&
  Length@Solve[W2init == x1 px + y py, y] == 1,
Print["No Multiple Solutions (Consumer 2)"],
Print["Check for Multiple Solutions! (Consumer 2)"]]

Print[]

Manipulate[Plot[OCB1[px, py][[1]], {px, 0, pxmax},
  PlotLabel -> "Optimizing Demand Curve (Consumer 1, Good X) for Py = " [py],
  PlotRange -> {{0, pxmax}, {-5, xmax}}, WorkingPrecision -> MachinePrecision],
{{py, 10, Subscript[P, y]}, 0, 1000}, {{pxmax, 50, "[Pxmax]"}, 0, 1000},
{{xmax, 50, "[Xmax]"}, 0, 100}]

Manipulate[Plot[OCB1[px, py][[2]], {py, 0, pymax},
  PlotLabel -> "Optimizing Demand Curve (Consumer 1, Good Y) for Px = " [px],
  PlotRange -> {{0, pymax}, {-5, ymax}}, WorkingPrecision -> MachinePrecision],
{{px, 10, Subscript[P, x]}, 0, 1000}, {{pymax, 50, "[Pymax]"}, 0, 100},
{{ymax, 50, "[Ymax]"}, 0, 1000}]

Manipulate[Plot[OCB2[px, py][[1]], {px, 0, pxmax},
  PlotLabel -> "Optimizing Demand Curve (Consumer 2, Good X) for Py = " [py],
  PlotRange -> {{0, pxmax}, {-5, xmax}}, WorkingPrecision -> MachinePrecision],
{{py, 10, Subscript[P, y]}, 0, 1000}, {{pxmax, 50, "[Pxmax]"}, 0, 1000},
{{xmax, 50, "[Xmax]"}, 0, 100}]

Manipulate[Plot[OCB2[px, py][[2]], {py, 0, pymax},
  PlotLabel -> "Optimizing Demand Curve (Consumer 2, Good Y) for Px = " [px],
  PlotRange -> {{0, pymax}, {-5, ymax}}, WorkingPrecision -> MachinePrecision],
{{px, 10, Subscript[P, x]}, 0, 1000}, {{pymax, 50, "[Pymax]"}, 0, 100},
{{ymax, 50, "[Ymax]"}, 0, 1000}]

"Define CSSs";

"Principal Path of Each  $\Pi$  Field
LHS=RHS -> {LHS,RHS}";

If[Vertical1 == True, piPP1 = {0, x};, piPP1 = {y, c1A * x};]
If[Vertical2 == True, piPP2 = {0, x};, piPP2 = {y, c2A * x};]

```

```

"Define CSSs";

CSS1[px_, py_] := (

  Unprotect[c1A, c2A, lambda1, lambda2];

  If[OCB1[px, py][[1]] ≠ 0, c1A = OCB1[px, py][[2]] / OCB1[px, py][[1]],
    Vertical1 = True; &&Print["Consumer 1 Principal Path is Vertical"];]

  If[OCB2[px, py][[1]] ≠ 0, c2A = OCB2[px, py][[2]] / OCB2[px, py][[1]],
    Vertical2 = True; &&Print["Consumer 2 Principal Path is Vertical"];]

  "Principal Path of Each  $\Pi$  Field
  LHS=RHS -> {LHS,RHS}";

  If[Vertical1 == True, piPP1 = {0, x};, piPP1 = {y, c1A * x};];
  If[Vertical2 == True, piPP2 = {0, x};, piPP2 = {y, c2A * x};];

  c1CSSmu1 = Ux1[OCB1[px, py][[1]], OCB1[px, py][[2]]] / px;
  c1CSSmu2 = Uy1[OCB1[px, py][[1]], OCB1[px, py][[2]]] / py;
  If[c1CSSmu1 == c1CSSmu2, Print["Lambda 1 = ", lambda1 = c1CSSmu1],
    Print["Error: Lambda mismatch"]];

  c2CSSmu1 = Ux2[OCB2[px, py][[1]], OCB2[px, py][[2]]] / px;
  c2CSSmu2 = Uy2[OCB2[px, py][[1]], OCB2[px, py][[2]]] / py;
  If[TrueQ[c2CSSmu1 == c2CSSmu2], Print["Lambda 2 = ", lambda2 = c2CSSmu1],
    Print["Error: Lambda mismatch"]];

  Protect[c1A, c2A, lambda1, lambda2];

)

"Satisficing Critical Values";

"Consumer 1";

```

```

clxcrit0[px_, mu1_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y]},
    Max[0, x /. First@Solve[Ux1[x, y1] == mu1 px, x]];

clycrit0[py_, mu1_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x]},
    Max[0, y /. First@Solve[Uy1[x1, y] == mu1 py, y]];

clxcrit1[px_, py_, mu1_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y], y2 =
    y /. First@Solve[Uy1[x, y] == py * mu1 && piPP1[[1]] == piPP1[[2]], y, x]},
    Max[0, x /. First@Solve[Ux1[x, y1] == mu1 px, x],
    x /. First@Solve[Ux1[x, Max[0, y2]] == mu1 px, x]];

clycrit1[px_, py_, mu1_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x], x2 =
    x /. First@Solve[Ux1[x, y] == px * mu1 && piPP1[[1]] == piPP1[[2]], x, y]},
    Max[0, y /. First@Solve[Uy1[x1, y] == mu1 py, y],
    y /. First@Solve[Uy1[Max[0, x2], y] == mu1 py, y]];

clxcrit2[px_, py_, mu1_, W1_] =
  With[{y1 = y /. First@Solve[piPP1[[1]] == piPP1[[2]], y],
    y2 = y /. {y -> clycrit0[py, mu1]}},
    Max[0, x /. First@Solve[W1 == px x + py y1, x],
    x /. First@Solve[W1 == px x + py y2, x]];

clycrit2[px_, py_, mu1_, W1_] =
  With[{x1 = x /. First@Solve[piPP1[[1]] == piPP1[[2]], x],
    x2 = x /. {x -> clxcrit0[px, mu1]}},
    Max[0, y /. First@Solve[W1 == px x1 + py y, y],
    y /. First@Solve[W1 == px x2 + py y, y]];

      "Consumer 2";

c2xcrit0[px_, mu2_] =
  With[{y1 = y /. First@Solve[piPP2[[1]] == piPP2[[2]], y]},
    Max[0, x /. First@Solve[Ux2[x, y1] == mu2 px, x]];

c2ycrit0[py_, mu2_] =
  With[{x1 = x /. First@Solve[piPP2[[1]] == piPP2[[2]], x]},

```

```

Max[0, y /. First@Solve[Uy2[x1, y] == mu2 py, y]];

c2xcrit1[px_, py_, mu2_] =
  With[{y1 = y /. First@Solve[piPP2[[1]] == piPP2[[2]], y], y2 =
    y /. First@Solve[Uy2[x, y] == py * mu2 && piPP2[[1]] == piPP2[[2]], y, x]},
  Max[0, x /. First@Solve[Ux2[x, y1] == mu2 px, x],
    x /. First@Solve[Ux2[x, Max[0, y2]] == mu2 px, x]];

c2ycrit1[px_, py_, mu2_] =
  With[{x1 = x /. First@Solve[piPP2[[1]] == piPP2[[2]], x], x2 =
    x /. First@Solve[Ux2[x, y] == px * mu2 && piPP2[[1]] == piPP2[[2]], x, y]},
  Max[0, y /. First@Solve[Uy2[x1, y] == mu2 py, y],
    y /. First@Solve[Uy2[Max[0, x2], y] == mu2 py, y]];

c2xcrit2[px_, py_, mu2_, W2_] =
  With[{y1 = y /. First@Solve[piPP2[[1]] == piPP2[[2]], y],
    y2 = y /. {y -> c2ycrit0[py, mu2]}},
  Max[0, x /. First@Solve[W2 == px x + py y1, x],
    x /. First@Solve[W2 == px x + py y2, x]];

c2ycrit2[px_, py_, mu2_, W2_] =
  With[{x1 = x /. First@Solve[piPP2[[1]] == piPP2[[2]], x],
    x2 = x /. {x -> c2xcrit0[px, mu2]}},
  Max[0, y /. First@Solve[W2 == px x1 + py y, y],
    y /. First@Solve[W2 == px x2 + py y, y]];

"Graph of Consumer 1 Constraints in Consumption Space";

Manipulate[
  Plot[{With[{y1 = y /. First@Solve[W1 == px x + py y, y]}, y1], Min[c1ycrit0[py,
    mul], Flatten[Solve[piPP1[[1]] == piPP1[[2]], y]][[1]][[2]]]},
  {x, 0, 50}, Epilog -> {With[{x1 = x /. First@Solve[Ux1[x, y] == mul px, x]},
    Line[{{x1, 0}, {x1, 50}}]}],
  PlotRange -> {{0, 50}, {0, 50}}, AxesLabel -> {"X", "Y"},
  PlotLabel -> "Consumption Space 1 (X Constraints)",
  {{px, 10, Subscript[P, x]}, 0, 50},
  {{py, 10, Subscript[P, y]}, 0, 50},
  {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000}]

```

```

Manipulate[
Plot[{With[{x1 = x /. First@Solve[W1 == px x + py y, x]}, x1], Min[c1xcrit0[px,
    mul], Flatten[Solve[piPP1[[1]] == piPP1[[2]], x]][[1]][[2]]]},
  {y, 0, 50}, Epilog -> {With[{y1 = y /. First@Solve[Uy1[x, y] == mul py, y]},
    Line[{y1, 0}, {y1, 50}]}]},
PlotRange -> {{0, 50}, {0, 50}}, AxesLabel -> {"Y", "X"},
PlotLabel -> "Consumption Space 1 (Y Constraints)",
{{px, 10, Subscript[P, x]}, 0, 50},
{{py, 10, Subscript[P, y]}, 0, 50},
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000}}

```

"Graph of Consumer 2 Constraints in Consumption Space";

```

Manipulate[
Plot[{With[{y1 = y /. First@Solve[W2 == px x + py y, y]}, y1], Min[c2ycrit0[py,
    mu2], Flatten[Solve[piPP2[[1]] == piPP2[[2]], y]][[1]][[2]]]},
  {x, 0, 50}, Epilog -> {With[{x1 = x /. First@Solve[Ux2[x, y] == mu2 px, x]},
    Line[{x1, 0}, {x1, 50}]}]},
PlotRange -> {{0, 50}, {0, 50}}, AxesLabel -> {"X", "Y"},
PlotLabel -> "Consumption Space 2 (X Constraints)",
{{px, 10, Subscript[P, x]}, 0, 50},
{{py, 10, Subscript[P, y]}, 0, 50},
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{W2, W2init, Subscript[Wealth, 2]}, 0, 1000}}

```

```

Manipulate[
Plot[{With[{x1 = x /. First@Solve[W2 == px x + py y, x]}, x1], Min[c2xcrit0[px,
    mu2], Flatten[Solve[piPP2[[1]] == piPP2[[2]], x]][[1]][[2]]]},
  {y, 0, 50}, Epilog -> {With[{y1 = y /. First@Solve[Uy2[x, y] == mu2 py, y]},
    Line[{y1, 0}, {y1, 50}]}]},
PlotRange -> {{0, 50}, {0, 50}}, AxesLabel -> {"Y", "X"},
PlotLabel -> "Consumption Space 2 (Y Constraints)",
{{px, 10, Subscript[P, x]}, 0, 50},
{{py, 10, Subscript[P, y]}, 0, 50},
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{W2, W2init, Subscript[Wealth, 2]}, 0, 1000}}

```

"Demand Functions";


```

    "Individual Demand Functions";

Default[c1Dx, 3] = mulinit;
Default[c1Dx, 4] = W1init;
c1Dx[px_, py_, mu1_., W1_.] =
    Max[0, Min[c1xcrit1[px, py, mu1], c1xcrit2[px, py, mu1, W1]]];

Default[c1Dy, 3] = mulinit;
Default[c1Dy, 4] = W1init;
c1Dy[px_, py_, mu1_., W1_.] =
    Max[0, Min[c1ycrit1[px, py, mu1], c1ycrit2[px, py, mu1, W1]]];

Default[c2Dx, 3] = mu2init;
Default[c2Dx, 4] = W2init;
c2Dx[px_, py_, mu2_., W2_.] =
    Max[0, Min[c2xcrit1[px, py, mu2], c2xcrit2[px, py, mu2, W2]]];

Default[c2Dy, 3] = mu2init;
Default[c2Dy, 4] = W2init;
c2Dy[px_, py_, mu2_., W2_.] =
    Max[0, Min[c2ycrit1[px, py, mu2], c2ycrit2[px, py, mu2, W2]]];

    "Market Demand Functions";

Default[MDx, 3] = mulinit;
Default[MDx, 4] = mu2init;
Default[MDx, 5] = W1init;
Default[MDx, 6] = W2init;
MDx[px_, py_, mu1_., mu2_., W1_., W2_.] =
    c1Dx[px, py, mu1, W1] + c2Dx[px, py, mu2, W2];

Default[MDy, 3] = mulinit;
Default[MDy, 4] = mu2init;
Default[MDy, 5] = W1init;
Default[MDy, 6] = W2init;
MDy[px_, py_, mu1_., mu2_., W1_., W2_.] =
    c1Dy[px, py, mu1, W1] + c2Dy[px, py, mu2, W2];

"Plotting Individual Demand Curves and Constituent Loci";

```

```

Manipulate[Plot[c1xcrit0[px, mu1], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 First Critical X", PlotPoints → 2000],
{{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
{{PxMax, Group0PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group0QMax, "[QxMax]"}, 0, 100}]

Manipulate[Plot[c1ycrit0[py, mu1], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 First Critical Y", PlotPoints → 2000],
{{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
{{PyMax, Group0PMax, "[PyMax]"}, 0, 100},
{{Ymax, Group0QMax, "[QyMax]"}, 0, 100}]

Manipulate[Plot[c2xcrit0[px, mu2], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 2 First Critical X", PlotPoints → 2000],
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{PxMax, Group0PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}]

Manipulate[Plot[c2ycrit0[py, mu2], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 2 First Critical Y", PlotPoints → 2000],
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{PyMax, Group0PMax, "[PyMax]"}, 0, 100},
{{Ymax, Group0QMax, "[QyMax]"}, 0, 100}]

Manipulate[Plot[c1xcrit1[px, py, mu1], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 1 Worth Locus (X)", PlotPoints → 2000],
{{mu1, mu1init, Subscript[μ, 1]}, 0, 100}, {{py, 5, Subscript[P, y]}, 0, 25},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}]

Manipulate[Plot[c1xcrit2[px, py, mu1, W1], {px, 0, PxMax},

```

```

PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
PlotLabel → "Consumer 1 Budget Locus (X)", PlotPoints → 2000],
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[
Plot[{c1xcrit1[px, py, mul], c1xcrit2[px, py, mul, W1]}, {px, 0, PxMax},
PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
PlotLabel → "Consumer 1 Combined Loci (X)", PlotPoints → 2000],
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[Plot[c1Dx[px, py, mul, W1], {px, 0, PxMax},
PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
PlotLabel → "Consumer 1 Demand Curve X (transposed) for Py = "[py],
PlotPoints → 2000], {{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, ]}, 0, 1000},
{{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group1PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group1QMax, "[QxMax]"}, 0, 100}}

Manipulate[Plot[c1ycrit1[px, py, mul], {py, 0, PyMax},
PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
PlotLabel → "Consumer 1 Worth Locus (Y)", PlotPoints → 2000],
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
{{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

Manipulate[Plot[c1ycrit2[px, py, mul, W1], {py, 0, PyMax},
PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
PlotLabel → "Consumer 1 Budget Locus (Y)", PlotPoints → 2000],
{{mul, mulinit, Subscript[μ, 1]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
{{Ymax, Group2QMax, "[QyMax]"}, 0, 100}}

```

```

Manipulate[
  Plot[{clycrit1[px, py, mu1], clycrit2[px, py, mu1, W1]}, {py, 0, PyMax},
    PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
    PlotLabel → "Consumer 1 Combined Loci (Y)", PlotPoints → 2000],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}]

Manipulate[Plot[c1Dy[px, py, mu1, W1], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 1 Demand Curve Y (transposed) for Px = "[px],
  PlotPoints → 2000], {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group2PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group2QMax, "[QyMax]"}, 0, 100}]

Manipulate[Plot[c2xcrit1[px, py, mu2], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 2 Worth Locus (X)", PlotPoints → 2000],
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group3PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group3QMax, "[QxMax]"}, 0, 100}]

Manipulate[Plot[c2xcrit2[px, py, mu2, W2], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 2 Budget Locus (X)", PlotPoints → 2000],
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group3PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group3QMax, "[QxMax]"}, 0, 100}]

Manipulate[
  Plot[{c2xcrit1[px, py, mu2], c2xcrit2[px, py, mu2, W2]}, {px, 0, PxMax},
    PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
    PlotLabel → "Consumer 2 Combined Loci (X)", PlotPoints → 2000],

```

```

{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
{{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group3PMax, "[PxMax]"}, 0, 100},
{{Xmax, Group3QMax, "[QxMax]"}, 0, 100}}

```

```

Manipulate[Plot[c2Dx[px, py, mu2, W2], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Consumer 2 Demand Curve X (transposed) for Py = "[py],
  PlotPoints → 2000], {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50}, {{PxMax, Group3PMax, "[PxMax]"}, 0, 100},
  {{Xmax, Group3QMax, "[QxMax]"}, 0, 100}}

```

```

Manipulate[Plot[c2ycrit1[px, py, mu2], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 2 Worth Locus (Y)", PlotPoints → 2000],
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group4PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group4QMax, "[QyMax]"}, 0, 100}}

```

```

Manipulate[Plot[c2ycrit2[px, py, mu2, W2], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 2 Budget Locus (Y)", PlotPoints → 2000],
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group4PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group4QMax, "[QyMax]"}, 0, 100}}

```

```

Manipulate[
  Plot[{c2ycrit1[px, py, mu2], c2ycrit2[px, py, mu2, W2]}, {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 2 Combined Loci (Y)", PlotPoints → 2000],
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group4PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group4QMax, "[QyMax]"}, 0, 100}}

```

```

Manipulate[Plot[c2Dy[px, py, mu2], {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Consumer 2 Demand Curve Y (transposed) for Px = "[px],
  PlotPoints → 2000], {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{px, 5, Subscript[P, x]}, 0, 50}, {{PyMax, Group4PMax, "[PyMax]"}, 0, 100},
  {{Ymax, Group4QMax, "[QyMax]"}, 0, 100}]

```

"Plotting Market Demand Curves";

```

Manipulate[Quiet[XDIndividuals =
  Plot[{c1Dx[px, py, mu1, W1], c2Dx[px, py, mu2, W2]}, {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Individual Demand Curves X (transposed)",
  WorkingPrecision → MachinePrecision, PlotPoints → 2000]],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50},
  {{PxMax, PMaxMarket, "[PxMax]"}, 0, 2 PmaxMarket},
  {{Xmax, QMaxMarket, "[QxMax]"}, 0, 2 QMaxMarket}]

```

```

Manipulate[PYG = py;
  Quiet[XDMarket = Plot[MDx[px, py, mu1, mu2, W1, W2], {px, 0, PxMax},
  PlotRange → {{0, PxMax}, {0, Xmax}}, AxesLabel → {Subscript[P, x], "X"},
  PlotLabel → "Market Demand Curve X (transposed) for Py = "[py],
  WorkingPrecision → MachinePrecision, PlotPoints → 2000]],
  {{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
  {{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
  {{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
  {{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
  {{py, 5, Subscript[P, y]}, 0, 50},
  {{PxMax, PMaxMarket, "[PxMax]"}, 0, 2 PmaxMarket},
  {{Xmax, QMaxMarket, "[QxMax]"}, 0, 2 QMaxMarket}]

```

```

Manipulate[Quiet[YDIndividuals =
  Plot[{c1Dy[px, py, mu1, W1], c2Dy[px, py, mu2, W2]}, {py, 0, PyMax},
  PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
  PlotLabel → "Individual Demand Curves Y (transposed)",

```

```

WorkingPrecision → MachinePrecision, PlotPoints → 2000]],
{{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
{{px, 5, Subscript[P, x]}, 0, 50},
{{PyMax, PMaxMarket, "[PyMax]"}, 0, 2 PmaxMarket},
{{Ymax, QMaxMarket, "[QyMax]"}, 0, 2 QMaxMarket}}

Manipulate[PXG = px;
Quiet[YDMarket = Plot[MDy[px, py, mu1, mu2, W1, W2], {py, 0, PyMax},
PlotRange → {{0, PyMax}, {0, Ymax}}, AxesLabel → {Subscript[P, y], "Y"},
PlotLabel → "Market Demand Curve Y (transposed) for PX = " [px],
WorkingPrecision → MachinePrecision, PlotPoints → 2000]],
{{mu1, mu1init, Subscript[μ, 1]}, 0, 100},
{{mu2, mu2init, Subscript[μ, 2]}, 0, 100},
{{W1, W1init, Subscript[Wealth, 1]}, 0, 1000},
{{W2, W2init, Subscript[Wealth, 2]}, 0, 1000},
{{px, 5, Subscript[P, x]}, 0, 50},
{{PyMax, PMaxMarket, "[PyMax]"}, 0, 2 PmaxMarket},
{{Ymax, QMaxMarket, "[QyMax]"}, 0, 2 QMaxMarket}}

p1 = 7;
p2 = 6;
CSS[p1, p2]
OCB1[p1, p2]
OCB2[p1, p2]

c1A
c2A

```

```

Plot[c1Dx[px, p2, lambda1, Wlinit], {px, 0, 50},
  PlotRange -> {{0, 25}, {0, 50}}, Epilog -> {Dashed, Line[{{p1, 0}, {p1, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]
Plot[c2Dx[px, p2, lambda2, W2init], {px, 0, 50},
  PlotRange -> {{0, 25}, {0, 50}}, Epilog -> {Dashed, Line[{{p1, 0}, {p1, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]

Plot[MDx[px, p2, lambda1, lambda2, Wlinit, W2init], {px, 0, 25},
  PlotRange -> {{0, 25}, {0, 25}}, PlotLabel -> "Market X",
  Epilog -> {Dashed, Line[{{p1, 0}, {p1, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]
Plot[c1Dy[p1, py, lambda1, Wlinit], {py, 0, 50},
  PlotRange -> {{0, 25}, {0, 50}}, Epilog -> {Dashed, Line[{{p2, 0}, {p2, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]
Plot[c2Dy[p1, py, lambda2, W2init], {py, 0, 50},
  PlotRange -> {{0, 25}, {0, 25}}, Epilog -> {Dashed, Line[{{p2, 0}, {p2, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]
Plot[MDy[p1, py, lambda1, lambda2, Wlinit, W2init], {py, 0, 50},
  PlotRange -> {{0, 25}, {0, 50}}, PlotLabel -> "Market Y",
  Epilog -> {Dashed, Line[{{p2, 0}, {p2, 50}}]},
  WorkingPrecision -> MachinePrecision, PlotPoints -> 2000]

c1xcrit1[p1, p2, lambda1] ==
  c1xcrit2[p1, p2, lambda1, Wlinit] == c1Dx[p1, p2, lambda1, Wlinit]

c1ycrit1[p1, p2, lambda1] ==
  c1ycrit2[p1, p2, lambda1, Wlinit] == c1Dy[p1, p2, lambda1, Wlinit]

c2xcrit1[p1, p2, lambda2] ==
  c2xcrit2[p1, p2, lambda2, W2init] == c2Dx[p1, p2, lambda2, W2init]

c2ycrit1[p1, p2, lambda2] ==
  c2ycrit2[p1, p2, lambda2, W2init] == c2Dy[p1, p2, lambda2, W2init]

{c2xcrit1[p1, p2, lambda2],
  c2xcrit2[p1, p2, lambda2, W2init], c2Dx[p1, p2, lambda2, W2init]}

{c2ycrit1[p1, p2, lambda2],
  c2ycrit2[p1, p2, lambda2, W2init], c2Dy[p1, p2, lambda2, W2init]}

{c1xcrit1[p1, p2, lambda1],
  c1xcrit2[p1, p2, lambda1, Wlinit], c1Dx[p1, p2, lambda1, Wlinit]}

{c1ycrit1[p1, p2, lambda1],
  c1ycrit2[p1, p2, lambda1, Wlinit], c1Dy[p1, p2, lambda1, Wlinit]}

```



```

test1 = Plot[c2Dx[px, p2, lambda2, W2init], {px, 0, 50},
  PlotRange -> {{0, 25}, {0, 25}}, Epilog -> {Dashed, Line[{{p1, 0}, {p1, 50}}]},
  WorkingPrecision -> MachinePrecision,
  PerformanceGoal -> "Quality", PlotPoints -> 2000]
Plot[{Evaluate@c2xcrit1[px, p2, lambda2],
  Evaluate@c2xcrit2[px, p2, lambda2, W2init]}, {px, 0, 50},
  PlotRange -> {{0, 25}, {0, 25}}, Epilog -> {Dashed, Line[{{p1, 0}, {p1, 50}}]},
  WorkingPrecision -> MachinePrecision,
  PerformanceGoal -> "Quality", PlotPoints -> 2000]

```

REFERENCES CITED

- Abreu, D., & Rubinstein, A. (1988). The Structure of Nash Equilibrium in Repeated Games with Finite Automata. *Econometrica*, 56, 1259-1281.
- Adam, K. (2005). Learning to Forecast and Cyclical Behavior of Output and Inflation. *Macroeconomic Dynamics*, 9, 1-27.
- Bacon, R.W. (1991). Rockets and Feathers: The Asymmetric Adjustment of UK Retail Gasoline Prices to Cost Changes. *Energy Economics*, 13, 211-218.
- Balke, N.S., Brown, S.P.A., & Yücel, M.K. (1998). Crude Oil and Gasoline Prices: An Asymmetric Relationship? *Federal Reserve Bank of Dallas Economic Review*, First Quarter, 2-11.
- Borenstein, S., Cameron, A.C., & Gilbert, R. (1997). Do Gas Prices Respond Asymmetrically to Crude Oil Prices? *Quarterly Journal of Economics*, 112, 305-339.
- Brown, S.P.A., & Yücel, M.K. (2000). Gasoline and Crude Oil Prices: Why the Asymmetry? *Federal Reserve Bank of Dallas Economic and Financial Review*, Third Quarter, 23-29.
- Day, R.H. (1967). Profits, Learning and the Convergence of Satisficing to Marginalism. *Quarterly Journal of Economics*, 81, 302-311.
- Eaton, J. & Kortum, S. (1999). International Technology Diffusion: Theory and Measurement. *International Economic Review*, 40, 537-570.
- Federal Reserve Bank of St. Louis (2011). FRED, Federal Reserve Economic Data, from the Federal Reserve Bank of St. Louis. Retrieved from <http://research.stlouisfed.org/fred2/>
- Hall, R.L., & Hitch, C.J. (1939). Price Theory and Business Behavior. *Oxford Economic Papers*, 2, 12-45.
- Hogarth, R.M. (1975). Cognitive Processes and the Assessment of Subjective Probability Distributions. *Journal of the American Statistical Association*, 70, 271-289.
- Kapetyn, A., Wansbeek, T., & Buyze, J. (1979). Maximizing or Satisficing. *Review of Economic Statistics*, 61, 549-563.
- Mas-Colell, A., Whinston, M.D., & Green, J.R. (1995). *Microeconomic Theory*. Oxford, UK: Oxford University Press.

- Simon, H.A. (1956). Rational Choice and the Structure of the Environment. *Psychological Review*, 63, 129-138.
- Stigler, G.J. (1947). The Kinky Oligopoly Demand Curve and Rigid Prices. *Journal of Political Economy*, 55, 432-449.
- Stigler, G.J. (1961). The Economics of Information. *Journal of Political Economy*, 69, 213-225.
- Sweezy, P.M. (1939). Demand Under Conditions of Oligopoly. *Journal of Political Economy*, 47, 568-573.