Statement on Research and Teaching
Nolan Miller
October 4, 2008
Long Version[1]

I. Introduction

My work falls broadly into the area of applied microeconomic theory and focuses on strategic approaches to questions of importance to businesses and policymakers. It has addressed an array of practical areas, including industrial organization and business strategy, health economics, environmental regulation, international trade, design of reputation systems, and others. In recent years, the areas where I have concentrated the most are models of contracting, incentives, and competition in the health care sector and on the behavior of consumers living on the boundary of subsistence.

While a number of my papers are methodological in nature, all of them are driven by the need to provide a method to answer important real-world questions. I consider both developing methods to attack real-world problems and direct assaults on such problems to be important endeavors, and my work has tended to move back and forth between the two. I can describe this work using vertical and horizontal as metaphors. "Vertically," I usually complement a more technical paper with one that, using similar tools, closely considers a particular policy problem. "Horizontally," I have applied tools and models I developed to policy problems ranging from the environment, to regulation, to international trade, to poverty, gender, and household welfare. Given this strategy, and my tendency to be strongly attracted to problems stimulated, and at times proposed, by my colleagues and students, my research agenda is quite broad.

While at the Kennedy School, my teaching assignments at KSG have focused on teaching applied microeconomic theory to students in the Masters program in international development (MPA/ID) and PhD programs. Over the years these courses have developed into highly successful blends of serious theory with real attention to applications. My attempts to connect theory with the real world has led to several research projects as well as a textbook, *Applied Microeconomic Theory* (in collaboration with Chris Avery). The text, which is aimed at teaching advanced microeconomic theory to students such as professional school PhD students and advanced undergraduates and masters students, who are more interested in using economic theory in applied models/studies of practical problems than in proving technical properties, is currently in draft form, and Harvard University Press has agreed to publish it.

Section II discusses my research work to date and my ongoing research agenda. Section III discusses my teaching experience, and Section IV discusses my textbook project.

II. Research: Summary of Work to Date

Broadly speaking, my research to date has been in the area of **applied microeconomic theory.** Although my doctoral training was in theoretical microeconomics, my research reflects my training at business schools (Wharton undergraduate and Kellogg graduate) and the professional-school environment at the Kennedy School. Hence, my research tends to be driven by real-world questions relevant to businesses and policy-makers. Even most of my early work was motivated by real-world concerns and often focused on incorporating real-world features into economic models to show how seemingly strange predictions had more to do with unrealistic modeling choices than with the real world. More recently, much of my work has directly addressed business/policy questions, and my current work tends to move from practical question to theory and back again.

[1] A shorter, summary version of this document is also available at http://www.hks.harvard.edu/nhm/contact.htm .

In the past, my research has focused on three broad themes. The first is the interplay between information, incentives, and welfare in organizations. The second is concerned with industrial organization and game theoretic models of strategic interactions, with a particular emphasis on studying the roles of delegation and commitment. This line of research includes a number of papers on health economics, the practical area on which I have concentrated the most. The third theme grows out of my teaching experience, and is concerned with exploring links between the theory of consumer behavior and decision making and the actual behavior of extremely low-income consumers.

When I approach a topic it usually results in a cluster of papers that range from very theoretical to very applied. This is nicely illustrated by the pair **"Mechanism Design with Multidimensional, Continuous Types and Interdependent Valuations,"**[2] and **"Eliciting Informative Feedback: The Peer Prediction Method."**[3] As I discuss below, the former lays out very general principles for analyzing the problem of inducing people to truthfully reveal information for use in public decision making, while the latter considers practically implementing such a system for solving the very applied problem of eliciting informative feedback in Internet transactions. My papers on price versus quantity competition with delegation have a similar pattern. **"The Equivalence of Price and Quantity Competition with Delegation,"**[4] illustrates a general solution to the difficulties that arise when seemingly innocuous situations about how to model a strategic interaction yield vastly different predictions, while **"Strategic Trade and Delegated Competition"**[5] applies these tools to the strategic trade problem (i.e., the question of whether a nation benefits from subsidizing its domestic industries), where the literature had been stymied by the fact that optimal trade policies are highly sensitive to the details of the way trade is modeled.

I will now expand on each of the general areas of my research mentioned above. Although much of my work is coauthored, for expositional purposes I will use the first person singular in describing my work. Bibliographic information is contained in footnotes when the paper is first mentioned and in the bibliography at the end of this document.

IIA. Information, Incentives, and Welfare.

This group of papers deals with one of the oldest and most persistent problems in economics: How can organizations and institutions be designed to ensure that resources are allocated efficiently and effectively, and that agents take actions appropriate actions to maximize social welfare? At the highest level, this can be seen as the problem of designing effective governance mechanisms for organizations. Under ideal conditions, markets do an excellent job of allocating resources efficiently (although they are indifferent to questions of equity). However, once imperfections are introduced in the form of information asymmetries or externalities, markets generally do not allocate resources efficiently, and may in fact perform very poorly even on this standard, leaving aside concerns of equity. This group of papers studies the interaction between market imperfections (especially in the form of asymmetric information), incentives, and welfare.

"Mechanism Design with Multidimensional, Continuous Types and Interdependent Valuations" and **"Eliciting Informative Feedback: The Peer Prediction Method,"** are two papers that consider the question of how people can be induced to reveal private information in order to guide public decision

[2] with John Pratt, Richard Zeckhauser, and Scott Johnson. *Journal of Economic Theory* (2007).
[3] with Paul Resnick and Richard Zeckhauser. *Management Science* (2005).
[4] with Amit Pazgal. *RAND Journal of Economics*, (2001).
[5] with Amit Pazgal. *Journal of International Economics* (2005).

makers in circumstances where truthful revelation may lead to public decisions that, while in the overall public interest, are not in a particular agent's interest. **"Mechanism Design …"** considers a very abstract, theoretical version of the problem. Put simply, the mechanism design problem asks the question posed above: when is it possible to induce agents to truthfully reveal their information (which is needed in order to make informed social decisions) when the decision that is ultimately made may not be in their own interest. This question is relevant in a wide variety of circumstances, from deciding whether to build and where to cite a public facility, to deciding which of a variety of potential public programs should be enacted, to selecting public officials, to deciding which firms should be allocated electromagnetic spectrum rights. The mechanism design problem has, to date, proven intractable when agents' information is multidimensional (e.g., they care both about the size and location of a publicly-funded bridge) and agents' payoffs are interrelated (e.g., others' use of the bridge affects their own enjoyment), which is likely to be the case in real-world decisions. In this paper, I bring together tools from economics, statistical decision theory, and functional analysis to prove a powerful general result: if agents' signals are correlated (as is likely to be the case in the real world), then it is possible to induce them to truthfully reveal their information (or more accurately, to reveal something arbitrarily close to the truth) using a system of "proper scoring rules".

While **"Mechanism Design …"** considers a very general version of the problem, **"Eliciting Informative Feedback …"** considers a particular, important, applied one. In recent years, the volume of transactions that takes place over the internet has grown tremendously. At the same time, so has the potential for fraud because of the internet's anonymity and lack of repeated interactions. Methods of enforcing honesty in internet transactions include reputation and rating systems such as are used on eBay or amazon.com. However, such systems often fall short because of limited participation and because reviews are uninformative. In this paper, I lay out how a transfer scheme based on the ideas developed in **"Mechanism Design …"** can be used to induce people to truthfully reveal their experiences with internet transactions in order to facilitate fair and honest trade. While **"Mechanism Design …"** was a general theory paper aimed at pushing forward the technical frontier in the mechanism design literature, **"Eliciting Informative Feedback …,"** was aimed at moving from an abstract scheme to one that could be implemented in practice. As such, the paper pays particular attention to implementation questions such as the types of questions that raters may be asked, the timing of how that information would be revealed, and how to structure incentives so as to induce raters to put effort into producing informative feedback.

Another group of papers focuses on issues related to the provision of incentives within organizations. **"Efficiency in Partnerships with Joint Monitoring"**[6] studied the partnership problem. That is, how can incentives to be designed to induce all members of a team to work hard when only the total output, not individual effort levels, is observable? In real world contexts, it is often difficult or impossible to measure the contributions of single individuals to a team or group output. Needless to say, this is a general challenge to managers, and of higher-level decision makers within any bureaucracy. Previous results in the area had shown that, if agents have no information about their teammates' actions, there is no way to induce all agents to work hard without threatening to destroy some of the output (i.e., imposing a collective punishment) if the realized output ends up smaller than expected. My paper shows that if agents can observe other agents' actions (as seems reasonable in real-world team situations), an incentive scheme can be created that induces all agents to put forth efficient effort that is balanced (i.e., always distributes the full output) and respects individuals' limited liability constraints. Thus, using agents to monitor other agents solves the partnership problem.

This paper, which was my first, illustrates a general theme that runs through many of my projects: sometimes when the theory predicts something that we do not see in the real world, it is because the

[6] *Journal of Economic Theory* (1997).

theory has forgotten to include an important feature of the real-world interaction. In this case, we know that despite moral hazard problems, teams often work reasonably well, because the individual team members can observe each others' contributions, even if the managers cannot. This paper shows that including this real-world feature in the model reconciles the theory with reality.

"Screening Budgets"[7] considers the problem of dividing an organization's resources between multiple units. In any organization, some units may be more productive than others, and it is in the organization's interest to identify these and give them a greater share of the available resources. The problem is that all units prefer more resources to less, and so even unproductive units have an interest in trying to convince the center that they are highly productive. In this paper, I show how correlation in the units' productivity across different tasks or over time can be used to distinguish between highly productive and unproductive units. I present data from a large company that is consistent with the predictions of the model. In ongoing (unpublished) work, **"Screening through Risk Aversion,"**[8] I consider the same question of how to identify productive units in order to give them greater access to resources, but use the fact that highly productive units are often less risk averse than their less productive counterparts to show how the center can use risky rewards to induce units to reveal their productivity. The results are validated using data from a survey of managers.

The theme of using risk to distinguish between different types of agents also appears in **"Possibly Final Offers,"**[9] where I investigate optimal selling mechanisms. I show that if buyers are risk averse it is optimal for a seller to use a "possibly final offer" strategy where he begins by offering to sell to the buyer at a high price. If the buyer rejects this initial offer, the seller employs a risky strategy, sometimes following with a lower price, but other times walking away from the transaction. I show that when buyers are risk averse such strategies improve the seller's expected profit, and that when buyers are extremely risk averse the seller's expected profit approaches the first-best profit, i.e., what he would earn if he could identify each buyer's valuation and make a take-it-or-leave-it offer to each.

IIB-1. Industrial Organization/Representing Strategic Interactions

The motive behind using theory-based modeling to answer practical questions is that models help us isolate the basic factors underlying human behavior and, in the case of game theoretic analysis, strategic interactions. It is therefore quite troubling when small variations in the way a situation is modeled lead to large changes in the model's predicted outcomes, especially if there are no compelling reasons why one of the candidate specifications is clearly superior to the others. One of the most striking examples of this type of pitfall is the so-called Bertrand paradox in duopolistic (two-firm) competition. Consider two identical firms that produce a homogeneous product. The Bertrand paradox concerns the fact that if firms are modeled as competing by setting quantities and letting the market determine the price (i.e., a la Cournot), then the model's equilibrium involves the equilibrium price exceeding the firms' marginal cost and both firms earning a positive profit. However, if the firms compete by setting prices, then the equilibrium involves firms' prices equaling marginal cost and their profits being driven to zero. Thus, the two approaches to modeling the problem (Cournot vs. Bertrand) lead to starkly different outcomes.

I have written a number of papers on the relationship between the way strategic interactions are modeled and the resulting predictions. One group has been concerned with showing how differences such as the Bertrand paradox evaporate once the game is seen in a sufficiently general light, while another shows that a firm can gain a competitive advantage by controlling how its opponent sees its approach to the game (e.g., in the Cournot-Bertrand case, whether the other firm thinks I am setting price or setting quantity).

[7] with Alex Wagner and Richard Zeckhauser. *Journal of Economic Behavior and Organization* (2006).
[8] with Alex Wagner and Richard Zeckhauser (2007).
[9] with Nikita Piankov and Richard Zeckhauser. *Journal of Economics and Management Strategy* (2006).

"The Equivalence of Price and Quantity Competition with Delegation," considers the issue of price vs. quantity competition in a differentiated-products environment. The paper connects the price vs. quantity literature with the strategic delegation literature, which studies how the owners of firms can manipulate the incentives of their managers in order to gain a strategic advantage in oligopolistic (few-firm) competition. In this paper, I show that in quite general environments, if owners have sufficient ability to manipulate their managers' incentives, then the equilibrium outcomes under price and quantity competition must coincide. Thus, the discrepancy in the outcomes found in the original Bertrand paradox (and its differentiated-products counterpart) arises from the fact that owners are limited to only two possible types of behavior: price-setting or quantity-setting. Once the full realm of possible behaviors is made available to owners it ceases to matter whether the product-market competition is modeled as taking place in prices or quantities.

"Strategic Trade and Delegated Competition," builds on the main result of **"Equivalence of Price and Quantity Competition …"** and applies it to one of the main theoretical difficulties in the theory of strategic international trade. In recent years, strategic trade has become a salient topic in the policy sphere as well as economics journals. One of the seminal results in the theory of strategic international trade is that if firms produce substitute products and compete by setting quantities, then the optimal strategic trade policy is for countries to subsidize their domestic industries (Brander and Spencer, 1985, *Journal of International Economics*). However, the policy recommendation is exactly reversed if firms compete by setting prices. In this case, countries should tax their domestic industries (Eaton and Grossman, 1986, *Quarterly Journal of Economics*). From the point of view of setting trade policy this is especially troubling, since it is unreasonable to think that governments will ever be able to observe whether firms in a market compete by setting prices or quantities, and indeed to some extent they may do both, e.g., adjusting on both dimensions. Further, in light of the analysis of my previous paper, even this information would not be sufficient to guide in setting trade policy, since the mode of competition combines with internal incentive mechanisms within the firms (some of which operate in other countries) to determine whether it is beneficial to subsidize the domestic industry.

In this paper, I apply the equivalence result from **"The Equivalence …"** to the strategic international trade problem. I argue that if owners have sufficient power to control their managers' behavior, then the outcome of the subsidy-/tax-setting game does not depend on whether firms compete by setting prices or quantities. Consequently, the optimal strategic trade policy is also independent of the form of product market competition. In addition, because any observed product-market behavior is consistent with either a delegated price-setting or a delegated quantity-setting model, I argue that the analysis of strategic trade interactions should not begin by looking for the "right model." Rather, they should take product-market *behavior* as primitive and use observed behavior to determine the direction of beneficial trade policy directly. I end the paper by outlining how, using readily available data, a government could estimate the parameters necessary to determine the direction of the optimal trade policy (i.e., whether it is a tax or subsidy).

In another group of papers, I look at other situations in which there are alternative ways of thinking about how to represent competition. My first paper in this area, **"Relative Performance as a Strategic Commitment Mechanism,"**[10] grew out of my interest in price vs. quantity competition. It explores how, in delegation-game environments, the owners of firms can use the fact that different owners have different attitudes toward how they perform relative to their rivals in order to gain a strategic advantage in competition with other firms.

[10] with Amit Pazgal. *Managerial and Decision Economics* (2002).

More recently, I have become interested in more general versions of the strategic implications of how players approach games. That is, how adopting one or another high-level approach to the game affects opponents' behavior and the likelihood of success. To illustrate, consider two firms that compete with each other in a particular market. If each firm chooses an advertising budget, these determine sales volumes and overall profit. On the other hand, if each firm were to choose a sales target and commit to such advertising expenditures as are necessary to achieve the target, then sales targets would determine advertising levels and overall profit. Thus, just as in the price vs. quantity case, we can think of firms as competing either by setting input budgets and letting output targets follow from the market, or setting output targets and letting input budgets follow from the market. And, as in the price vs. quantity case, the equilibrium outcome prediction will depend on which of these alternatives is chosen and firms will have definite preferences over whether they prefer to be seen by their opponents as "budget-setters" or "target-setters."

In **"Budget or Target: The Choice Between Input and Output Strategies,"**[11] I develop the general theoretical analysis of the budget or target question, and show how our analysis informs our understanding of strategic interactions ranging from industrial competition to international relations (the U.S.-U.S.S.R moon race) to negotiations. I show that players are less aggressive when facing an output-setting opponent who has made an open-ended commitment to achieve a particular goal than when facing an input-setting opponent who commits only to expend a certain amount of resources. This basic insight is robust to dynamic considerations as well as allowing for moderate levels of uncertainty and the possibility that output commitments are imperfect and/or costly. Thus, in situations where a player wants his opponent to be aggressive (i.e., firms producing complementary products), it is best to set an output strategy. In contrast, in situations where a player wants his opponent to be non-aggressive, it is better to set an output strategy. So, for example, President Kennedy's approach to the space race, declaring the bold output goal of being first to the moon and doing it by the end of the 1960's, was, from a strategic point of view, superior to declaring an input strategy (such as announcing how much money the US would spend on the moon program). Of course, this statement is only true to the extent that Kennedy could make the commitment credible. The strategic effects of output-setting are realized only if the other player believes the open-ended commitment to be credible.

"Advertising Budgets in Competitive Environments,"[12] is a companion paper to **"Budget or Target"** that takes the more general analysis to a particular applied problem. In this case, I investigate the question of whether it is better for firms to set advertising budgets (as in the "percentage of previous year's sales" approach used by many retailers and service providers) or to set sales targets (as in the "objective and task approach", also widely used by actual companies). Although the analysis accords with our analysis of the general case in "Budget or Target …," relevance and realism are gained by including much more market detail. In particular, both the nature of advertising and the substitutability of the firms' products play roles in determining whether it is better to follow the percentage of sales or objective and task approach.

A final paper on applied industrial organization theory combines my interests in industrial organization and in the role of information provision in influencing economic performance. **"The Effects of Environmental Regulation on Technology Diffusion: The Case of Chlorine Manufacturing"**[13] considers the question of whether regulation leads to the adoption of environmentally superior technologies, i.e., whether environmental regulations promote efficiency. Included in the regulations we consider is the Toxics Release Inventory, which required chlorine manufacturers to publicly report their use of certain hazardous chemicals and several indirect regulations on end users of chlorine. The paper

[11] with Amit Pazgal. *RAND Journal of Economics* (2006).
[12] with Amit Pazgal. *Quantitative Marketing and Economics* (2007).
[13] with Lori Snyder and Robert Stavins. *American Economic Review (Papers and Proceedings),* (2003).

found that regulation did not appear to speed the adoption of cleaner technologies to manufacture chlorine, but it did appear to hasten the closure of facilities that used the most-polluting technologies.

IIB-2. Industrial Organization/Incentives in Health Care

Although my background is not in health economics, in recent years I have become increasingly interested in the area. The questions that arise in the study health economics and health care markets are real, practical, and important, and issues of incentives arise in practically every aspect of the health care sector. Indeed, incentive issues are of even greater importance in health care problems, because there are so many places where price signals have been distorted. Thus, insured consumers consume too much health care because they are charged only a small fraction of the cost of care at the time of delivery, hospitals that receive a fixed payment for caring for a patient may stint on quality, etc. Because of this, I have found it to be a fascinating and fruitful area of research.

The health care arena also provides an excellent applied area in which to pursue my more general interests into the interplay between information and incentives, and in the role of strategic commitments in determining welfare.

In **"Pricing Health Benefits: A Cost Minimization Approach,"**[14] I study the question of how a cost-minimizing employer should design and price the health benefits packages it offers to its employees. This study is motivated by two real-world concerns. First, one of the policy concerns people have with employer-provided health benefits is that employers will use their compensation policies to make employment unattractive to unhealthy workers. Second, in recent years there have been a number of high-profile "premium death spirals" in which high employee contributions lead the healthiest employees to opt out of high-end insurance plans. Since the remaining members of the plan are now, on average, sicker, this leads to premium increases and further drop-outs. The result is a vicious circle in which premiums rise and enrollment declines to such an extent that the plan becomes unviable. (For example, Cutler and Reber, *Quarterly Journal of Economics,* 1998, document the plan death spiral that occurred at Harvard in the mid-1990's.) Given the importance of employer-provided private health insurance in the US health care system, instability of these markets is of great concern. This concern appears to be growing in light of new health reform proposals such as the one president Bush made during the 2007 State of the Union address, which extended tax deductibility of health insurance contributions to those who purchase insurance from the individual (non-group) market. At this time, health economists and health policy experts are actively debating the question of whether implementing this proposal would increase the likelihood of death spirals and/or other instabilities in the employer-provided (group) health insurance market.

The main results of this paper are that, first, because unhealthy employees value health benefits more than healthy ones, an employer who compensates its employees with wages and health benefits cannot use its compensation scheme to exclude the least healthy employees (except possibly by placing an upper limit on total insurance benefits). Second, an employer who follows the cost-minimization approach will offer the more generous plan whenever it is efficient to do so (although possibly to too few people). Thus, death spirals are both unnecessary and unprofitable. As a result, employers who follow the cost-minimization approach to pricing their health plan offerings will be less prone to premium death spirals than those who follow the "equal contribution" rule (where the employer contributes a fixed dollar amount to the employee's insurance, regardless of which plan the employee chooses) as Harvard did when it experienced its death spiral.

[14] *Journal of Health Economics,* (2005).

In **"Insurer-Provider Integration, Credible Commitment, and Managed Care Backlash,"**[15] I consider the increasing dissatisfaction with managed care organizations such as HMOs that became apparent during the late 1990's. Earlier, managed care had been touted as a means of reducing the cost of health care while increasing its effectiveness. However, it eventually became apparent that one of the factors hindering the growth of managed care was the fact that people simply did not trust their HMOs to provide them with the right kind of medical care (as opposed to the cheapest type of care). In this paper, I study this phenomenon in the context of a model of the interaction between an insurer, physician and patient, focusing on the role of the structure of the health insurance market (i.e., whether it is separate from providers as in traditional insurance or integrated with them, as in managed care). I argue that the arms-length relationship between insurer and provider present in traditional insurance allows the insurer to credibly commit to providing high-quality care in a way that integrated insurer-providers such as HMOs are unable to do. In light of this, consumers' distrust of managed care organizations may be entirely justified. Managed care's promises to provide high-quality care are simply not credible. Having laid out the basic model, I then turn toward how the insights of the model inform our thinking about health care quality. To be viable, it is not enough for managed care organizations to provide high-quality care. They may also find ways to credibly signal their intention to provide high-quality care to consumers. Innovations such as quality report cards and pay for performance are examples of mechanism that may aid in doing so.

"Provider Choice of Quality and Surplus"[16] is the first part of an ongoing project with Karen Eggleston and Richard Zeckhauser. In this paper, we show that if a health care provider cares about both profit and the quality it provides to its patients (as would be the case in, for example, a nonprofit hospital), this may give rise to counter-intuitive incentives on the part of the provider. In particular, it is possible that the provider may respond to an increase in the rate of its profits that it retains by decreasing quality. For example, lowering taxes (either explicit or implicit) taxes on a provider who balances profit against quality turns the provider into a more efficient "profit machine," and thus a reasonable response to this may be to focus more on profit, reducing quality.

IIC. Individual Decisions and Policies

A final branch of my research focuses on the interface between the theory of consumer behavior, real world behavior, and policies. My interest in this area grows out of my teaching over the past eight years. My teaching consists of courses in advanced microeconomic theory intended for Kennedy School masters students and doctoral students at KSG (Public Policy, Political Economy and Government), HBS (Doctor of Business Administration), the Health Policy PhD program, and students in related programs (e.g., the PhD program in the Government Department and various graduate programs at MIT's Sloan School). These courses have in common that they focus on PhD level microeconomic theory. Thus the material addresses the standard subjects that would be taught to students in economics PhD programs. However, the audiences for my courses are not economics PhD students. The students in my courses need to be exposed to advanced microeconomic theory and to the language of economists, but their ultimate goal is not to be economic theorists. Rather, they will be applied economists or sophisticated policy analysts who are using advanced economic reasoning to guide their thinking about applied problems. I have found that one of the places where the microeconomic theory curriculum as taught to PhD students in economics departments falls short is in making connections between the theory as taught in the classroom and actual applications. My efforts to make these connections in the classroom have translated into a research interest in connecting economic theory with applications, especially in the area of consumer behavior.

---

[15] *Journal of Health Economics* (2006).

[16] with Karen Eggleston and Richard Zeckhauser. *International Journal of Health Care Finance and Economics.* (2006).

My first major research project in this area is concerned with the consumption behavior of very poor consumers in China. The poor in China are extremely poor – over 30% of consumers live below the World Bank's extreme poverty line of spending one (US) dollar per person per day. These consumers eat a very simple diet consisting of rice, vegetables (mainly cabbage) and meat when they can afford it in southern China, and wheat-based buns and breads, vegetables and meat when they can afford it in northern China. Despite their prevalence, little is known about the consumption decisions made by such consumers. This project is aimed at providing insight into the behavior of these consumers, and we expect it to yield several papers.

The data for the project comes from a study conducted in two provinces of China, Hunan (south) and Gansu (north) during 2006, during which we subsidized the staple commodity (rice in Hunan, wheat in Gansu) for a randomly selected group of consumers, collecting detailed household and consumption data before, during and after the intervention.

The first project using the data aims at resolving one of the oldest mysteries in economic theory. This mystery that dates back to Marshall (1895), who proposed the theoretical possibility of "Giffen behavior," a situation in which a consumer responds to an increase in the price of a good by consuming more of the good, a violation of the "Law of Demand." Despite this idea appearing in virtually every microeconomics textbook, no convincing evidence has been found. In this paper, **"Giffen Behavior and Subsistence Consumption" (formerly known as "Giffen Behavior: Theory and Evidence,")**[17] we argue that Giffen behavior is theoretically most likely in situations where consumers face pressing subsistence concerns, as they do in our Chinese data set. We then show that our experimental evidence provides compelling evidence of the existence of Giffen behavior among the very poor, as well as broader evidence in favor of a model of consumption in the face of subsistence concerns. Beyond its pedagogical significance and significance for economic theory in general, the study also provides insights into the behavior of very poor consumers and the effects of policies aimed at improving their nutrition by providing income or price subsidies.

Other papers using this data are discussed under current research below.

III: Research -- Current Agenda

My current research agenda continues to develop a number of the themes from my earlier work. In particular, I have three major lines of research at this point, although there is some overlap between them.

The first combines my interest in information, incentives and welfare with my interest in health economics. In recent years, information provision and competition have been hailed as a potential cure for many of society's ills, ranging from poor quality in health care to poor compliance with environmental health mandates. This group of projects seeks to understand the connection between information provision, quality and welfare.

The first project studies the relationship between information, in the form of health care report cards, and quality. It has long been argued that if consumers are given access to better information about the quality of their health care providers, they will flock to the best providers, and thus competition will give all providers an incentive to improve quality. However, despite the intuitive appeal of this argument, the evidence on the effectiveness of health quality report cards has been mixed. They have certainly not led to uniform improvements in quality, and some have even argued that report cards have reduced quality (e.g., Dranove et al. (2003)). In this paper (**"Report Cards, Incentives, and Quality Competition in**

[17] With Robert Jensen (2008), *American Economic Review*, 98(4), 1553-1577.

Health Care,") I consider the fundamental argument underlying the justification for report cards: will making more accurate information on quality available to consumers *necessarily* induce providers to increase quality? I find the answer to this question to be no. Better information need not improve overall quality. (Quick intuition: if low-quality providers know they can longer "get lucky" and look like high-quality providers, they will lose any incentive they might have had to try and keep up with the industry leaders and fall even further behind them.) Further, I find that the impact of quality competition can be particularly harmful to vulnerable populations such as the poor or elderly who may be less able to "vote with their feet" than typical consumers. Thus, I argue, while report cards may lead to improvements in quality for some consumers in some circumstances, they are not a panacea. Indeed, report cards may actually harm some populations that we are particularly concerned about such as the poor or elderly. Hence it may be important to couple a report card program with other policies targeted at improving quality for these vulnerable populations.

A second project continues to focus on information and health care quality, but instead focuses on the consumer side. Just as in the case of report cards, consumerism, where individual consumers acquire and use information from sources other than their health care providers to guide their health decision making. In **"Demanding Patients: Consumerism and Quality of Care"** (joint with Hai Fang, John Rizzo, and Richard Zeckhauser), I evaluate the impact of "consumerism" on quality. It has been argued that consumerism will improve health care quality by helping patients better understand their interactions with the medical care system. However, while this is true for some consumers, it is not necessarily true for all consumers, especially when not all patients are well-informed "consumerists." Once again, it is the poor, elderly, and uneducated who are most likely to suffer. The key to the argument is to note that consumerist patients will tend to take up more of their doctors' time, leaving less time for non-consumerist patients. Even taking into account that doctors will reallocate their time in response to some patients becoming more informed, it is still the case that consumerism can harm quality, sometimes harming quality for all patients. The theoretical predictions of the model are validated by looking at data from data from the 2000/2001 Community Tracking Survey (CTS) maintained at the Center for Studying Health System Change. The analysis supports the contention that consumerism may, in fact, lower quality, as some doctors report.

Finally, I am starting work on a project in collaboration with members of the Health Care Policy Department at Harvard Medical School, that studies Medicare Advantage (MA), the managed-care Medicare program that runs alongside traditional Medicare. Intended as an option that will provide high quality care more efficiently than traditional Medicare (TM), recently MA has come under fire for actually being more expensive than TM. This project (joint with Joe Newhouse, Tom McGuire, Jacob Glazer and others) starts from the premise that in MA the government has the opportunity to design the rules of competition between MA plans and asks how these rules should be designed to achieve the goal of inducing plans to provide high-quality care at low cost. We will focus on questions such as (i) what is the optimal allocation of beneficiaries to traditional Medicare and MA (i.e., what kinds of patients are better-served by managed care)? (ii) what is the relationship between market structure and plans' conduct (i.e., how does the level of competition affect quality)? and (iii) what is the optimal form of payment to private plans in Medicare (taking into account cost differentials, risk adjustment, and the effects of competition)? Although this project is currently in the planning stage, I expect that it will develop over the next several years and produce a number of peer reviewed papers.

The second major strand of my ongoing work focuses on individual decision making, and in particular on policy-relevant field experiments that examine how the theory of individual decision making matches up with actual decisions, and how this information can be used to inform the design and evaluation of policies aimed at improving individual welfare.

The first of these field experiments is the China experiment described above, where we are now focusing on further analyzing the data and extending the theory in new directions. The second major paper using the China data is entitled "Do Subsidies Improve Nutrition?" This paper focuses on nutrition as the outcome variable of interest, rather than demand, and explores the impact of our experimental subsidy on nutritional outcomes and the implications of these results for public policy. Extremely impoverished consumers in China live on the boundary of malnutrition, and thus they are representative of the type of people throughout the world who are the targets of programs aimed at improving nutrition such as subsidies for basic goods, food stamps, and public food distributions. In light of this, understanding how consumers will respond to such programs is of paramount importance. In this paper, we consider how lowering the price of the staple commodity affects nutritional outcomes (e.g., calorie and protein intake). In the paper, we find evidence that our experimental subsidy actually decreased caloric intake for households in Hunan, and had virtually no effect on households in Gansu (except perhaps for the poorest of the poor). Thus, to the extent that subsidy programs aimed at the poor are intended to improve nutrition, our results show that these programs may be, at best, ineffective, and at worst actually harm nutrition. Further, we find a good deal of heterogeneity in behavior even among the very poor. Even though in one of our provinces the subsidy reduced calories overall, it may have had a positive effect on the poorest of the poor. Thus, proper targeting and evaluation of programs such as food subsidies or other aid for the poor should be mindful that this heterogeneity can obscure real effects if it is not taken into account.

A third paper using this dataset takes advantage of fortuitous timing. At the time we were collecting our data in late 2006, the world was beginning to experience a widespread, significant increase in food prices. In this paper, **"The Impact of the World Food Price Crisis on Nutrition in China,"**[18] we use our nutritional data to analyze the impact of the change in food prices over the course of our survey on nutrition, and find little or no impact on nutrition.

The work on the first several papers on this project has suggested new avenues for investigation. In the Giffen goods paper, we use a new measure of household poverty: the share of calories they receive from the local staple good. In developing countries, staple goods are frequently the cheapest source of calories available. Very poor households get a very large share of their calories from the staple, while richer households diversify their diets and have a lower "staple calorie share." The maximum staple calorie share consistent with a health diet can be computed using a variation on the minimum-cost diet problem studied by Stigler and others, and it is remarkably stable across different types of people. For example, while an active, working man may need two or three times the calories of a sedentary, elderly woman, making it difficult to judge just from caloric intake whether a person is malnourished, the maximum staple calorie share for the two is actually quite close. Due to this and other features, staple calorie share may provide a valuable poverty measure. The project, joint with Rob Jensen, explores this idea.

Another project expands on the theoretical model used in the Giffen goods paper and extends it to other arenas. The theoretical model underlying our analysis is one where consumers pay a utility penalty if they consume less than subsistence calories. Thus, the subsistence calorie constraint does not operate like a normal constraint, since consumers are free to violate it and face the penalty. In this project, joint with Daniel Hojman, we explore how the idea of optimization in the face of soft constraints can be used to explain a wide variety of reference-dependent behaviors, including Giffen behavior, conspicuous consumption, mental accounting, and others.

The second field experiment project that I have under way is also on the theme of information, quality and welfare. I am currently working on a project (with Lori (Snyder) Bennear of the Nichols School at Duke University (a KSG public policy PhD)) that studies contamination in private well water in North

[18] with Robert Jensen, *Agricultural Economics*, 2008.

Carolina, looking at how by giving households information on the contaminants in their water (and, in particular, by manipulating the way this information is presented) the government can better induce households to take action to improve their water quality and thereby protect their health. By way of background, drinking water obtained from private wells is currently unregulated in the United States. Consequently, private well water may contain inorganic contaminants such as arsenic, radon, and radium, long-term exposure to which is known to increase cancer risk. In addition to their increased likelihood of exposure, the population that receives its water from private wells is likely to be more vulnerable to the harmful health effects of drinking water contaminants than the population at large, since these households tend to be poorer, have less access to health care, and may live in their homes for longer periods of time. One way in which local, state and federal agencies have addressed private well water quality is through increasing information provision. That is, the government enacts policies that help make households aware of the contaminants in their water and the risks associated with them, and they rely on the newly-informed households to take action (when appropriate) to reduce their health risks. However, the efficacy of such policies depends on whether the information the government gives households is successful in motivating them to take action to reduce the risks to which they are exposed. It is not enough merely to inform households of the risks they face. Thus a critical task for policymakers is to determine how to present information to citizens in ways that effectively prompt appropriate actions.

Through a field experiment, this project seeks to determine whether changing the way information on contaminants is presented to households affects their decision to adopt health-promoting technologies such as water filters. The study focuses on households with private wells in Durham, Orange and Wake counties in North Carolina. Participants' wells will be tested for arsenic, radon, and radium and they will receive a detailed report of the level of contaminants detected in their well. In addition, participants will answer a series of surveys about their knowledge of drinking water hazards, their drinking water sources, and their use of treatment technologies. Data from these surveys will be used to estimate how different methods of presenting the risk information to households affect the likelihood that the households take action to mitigate their exposure to these contaminants.

In the experiment, we will focus on two types of informational manipulations. The first considers framing effects and considers how the way in which identical information induces different behavioral responses based on how it is framed. In our experiment, one group of participants will receive their contamination reports in a "gain" frame (e.g., using filters will reduce your cancer risk) while other will receive a "loss" frame (e.g., continuing to drink unfiltered water increases your cancer risk). Laboratory studies (starting with Kahneman and Tversky and continuing with a host of others) have shown that decision makers tend to be risk averse over gains and risk loving over losses. In our experiment, treating your water is a risk-reducing action, and so we would expect that those receiving the gain-framed information will be more likely to take action in response to their results.

While there is ample evidence of the effect of framing on decisions in the laboratory, little is known about the extent to which these laboratory phenomena translate into the real-world, and thus whether framing effects are a lever that policymakers can use to spur people to action. The results of this experiment, which track both the short-term effects of the manipulation (e.g., whether gain-framed messages induce people to say they are more concerned) as well as the long-term effects (e.g., whether gain-framed messages more effectively induce people to take costly actions to reduce their exposure to pollution), will yield insight into this important question.

The second type of informational manipulation we are interested in is the use of relative performance information. Initial pilot testing of the intervention revealed that people tended to be very interested in how their results compared to those of others in the study. In fact, they were more interested in this comparison than in how their results compared to state or national standards. A second branch of the

study will look at the impact of giving households information on relative performance on the likelihood of taking action to remove contaminants.

This project is currently in the data collection phase.

The third major branch of my ongoing research continues my work on game theoretic models of conflict. The first of my current papers, **Outcome Commitments in Third Party Intervention: Theory and Application to U.S. Policy in Iraq,** extends the ideas developed in my earlier work (see "Budget or Target: The Choice Between Input and Output Strategies" above) and applies it to the role of the U.S. in the conflict in Iraq. Much of the debate over U.S. policy in Iraq hinges on the level of U.S. commitment to the conflict. However, embedded in the debate is a question over how the conflict should be approached. Pro-war parties tend to emphasize goals and/or outcomes that will be attained ("We'll stand down when they stand up."), while anti-war forces tend to focus on input-based approaches (e.g. timetables for withdrawal). In this paper, I develop a model of strategic interaction in which a third party intervenes on behalf of a government in its conflict with insurgents. It examines whether it is better for the intervenor to adopt an input-based strategy (i.e., specify the total resources it will spend) or an outcome-based strategy (i.e., specify the goal that it will achieve), and it shows that outcome-based strategies are better for the intervenor than input-based ones if and only if in the absence of intervention the insurgents are stronger than the government. A system of benchmarks that are tied to the efforts of both parties outperforms both input-based and outcome-based strategies. Lessons from the theory are then applied to U.S. strategy in Iraq.

The Iraq paper raises a number of additional issues that I am not beginning to explore. For example, one of the critical ingredients in implementing an outcome-based strategy is credibly committing to that strategy. Tying in with my previous work on delegation and commitment, I am working on a model that brings together the political process, and in particular the fact that politicians in democracies are held accountable for the truth of their statements (even when the truth may not be in the constituents' best interests), and explores how this phenomenon can explain how politicians are able to credibly commit to strategies discussed in the Iraq paper.

In addition to the papers described here, I have several others in progress that fit into this broad area. These include a paper that looks at how a firm should optimally allocate its sales force to multiple markets (i.e., is it better to match strong salespeople against rival firms' strong people or against its weak people) and a paper that looks at how risk aversion by managers can be used to screen better managers from worse ones. Finally, some of the work on health care, especially the work on Medicare Advantage, also falls into this line of research.

III. Summary of Teaching Experience

IIIA. Degree Programs Teaching at KSG

My teaching assignments at the Kennedy School include three required courses in advanced microeconomic theory. API 109, Advanced Microeconomic Analysis I, is the first semester of the required microeconomics sequence for students in the Masters program in International Development (MPA/ID). This program is aimed at giving doctoral-level training in economics to students who will ultimately go on to be development practitioners (rather than academics), as part of an interdisciplinary approach to international development. The other courses I teach comprise the two-semester sequence in PhD-level microeconomics for Harvard doctoral students other than those enrolled in the PhD program in economics. API 111, Microeconomic Theory I, is the first semester of this sequence. The course is jointly listed as Economics 2020a and HBS 4010, and is required for students in KSG's Public Policy PhD program, the Political Economy and Government PhD program, HBS's DBA program, and some students in the Health Policy PhD program. It is also frequently taken by students in the Kennedy School's masters programs (especially the MPP and MPA/2 programs) who are interested in a course in advanced microeconomics, the government department's PhD program, advanced undergraduates in economics and/or applied math, and students in PhD programs at MIT's Sloan School. API 112 is the second semester of this sequence (jointly listed as Economics 2020b and HBS 4011).

All of these courses present the challenge of being taught at the doctoral level, but the students for whom the course is required often lack strong training in economics and/or mathematics, and, more importantly, they differ from most students in an economics department PhD program in that they are primarily interested in an applied area (e.g., international development, public policy, business, health policy) rather than in economic theory for its own sake. Nevertheless, most of these students will be doing professional-level work using the tools of advanced economics, and all of them will be using advanced economic reasoning in their professional lives. Thus we face the dual challenges of giving students the tools they need to do applied research or deal with professional economists, but at the same time helping them to see the connection between the technical tools and the policy questions that really interest them.

I arrived at the Kennedy School in the fall of 1999, along with the first cohort of MPA/ID students. Thus, I have been involved with the course and the MPA/ID program since its start, and I served on the admissions committee for several years. The challenges mentioned above are particularly salient with the MPA/ID course. The program is marketed to students as providing training at the level of economics PhD programs but aimed at future development practitioners. I, and the rest of the MPA/ID core faculty, take both of these points seriously. To paraphrase Dani Rodrik, we want to get the students to the point where they never have to start a policy debate by saying "I'm not an economist, but … ." At the time the program started, I believe that there was really no precedent for a course such as API 109. Over the past ten years, I have made a lot of progress on balancing the "advanced" and "applied" mandates, and at this point the course is quite successful because the students find it is relevant to their real world concerns. It is one of the highest rated economics courses at the Kennedy School.

An additional challenge in teaching the MPA/ID students is that many of the students entering the program have very limited exposure to economics and/or mathematics, and certainly nothing like the preparation of students entering good PhD programs in economics. So, the main goal of this course is getting students up to the point where they can understand professional-level economics while exposing them to the way the tools are used in international development. My approach to the course has been to stress intuition and applications. Mathematics is used as a tool to help understand the material, and I feel it is important to present the material mathematically since one of the purposes of the course is to introduce students to the language of economists, which is quite mathematical. Nevertheless, I try to use math to the extent that it aids understanding, avoiding mathematical generality for its own sake and

presenting mathematical arguments only when they aid in understanding or illustrate a general technique. I also try to include questions on exams that encourage students to consider the policy implications of the theory, to make connections between the theoretical material and real-world applications, and to think about how the concepts can be communicated to actual policy-makers.

To help bridge the gap between students' preparation and the level of the course, several years ago I wrote an extensive set of notes (about 230 pages). The notes as written served as a companion to the required text for the course, Mas-Collel, Whinston, and Green's *Microeconomic Theory* (aka MWG). This text, published in 1995, quickly became the standard text in PhD programs in economics, and its terminology and notation form much of the "language" of economists trained in the last fifteen years. The result is that any course that claims to be taught at the PhD level has to use this text. However, it is just too hard for many students, and certainly too abstract and light on applications for students like ours. The point of the notes, originally, was to help students navigate the MWG text and make connections between the math in the text and the real world. Over time, these notes have proven to be quite successful, and the have formed the basis for the textbook Chris Avery and I are currently working on (discussed below).

I have continued to develop the course over the past few years, revising the content and the accompanying materials in response to formal and informal student feedback. Most recently, I have begun introducing a series of applied, case-style questions to the problem sets. These problems are aimed at walking students through actual economic analyses (published papers) of issues such as food stamps, demand for nutrition, sin taxes, and pharmaceuticals in developing countries. Although this course is, first and foremost, a "tools" course, I have found that these applications really tend to drive home how the sometimes-abstract tools we are developing in the lectures are used in rigorous analyses of policy-relevant problems. Also, during the past year, I have been working closely with Michael Walton on integrating the tools we develop in the basic microeconomics course in to the MPA/ID applications and cases workshop. Often, I will present a general tool and the economists' view of a policy problem using that tool, and Michael will then take off from there, discussing the problem more generally.

API 111 – 112 is the required sequence in microeconomic theory for KSG PhD students and others. I have also been involved with this course since Fall 1999. Over time, I have co-taught it with Chris Avery and Jerry Green (of the Harvard economics department). The challenges we face in this course are similar to those in the MPA/ID course. However, the students in this course are, for the most part, people who will be doing applied research using the tools of microeconomics. Because of this, we tend to focus the course a bit more on "how could this come up in your research" than I do in the MPA/ID course. While API 109 and API 111 cover similar material, the material in API 112 is quite different.

The following table contains my teaching reviews for these courses. With the exception of my first semester here, my "instructor overall" ratings have always been over 4.0, and I have been over 4.5 in 14 out of 24 offerings, including the last 6 offerings and 11 out of the last 13 offerings. (A rating of 4.5 or higher qualifies for the "Dinner on the Dean" award and is considered excellent.) API 111 and 112 are generally co-taught with other instructors, and the "Course Overall" rating is for all instructors combined. The Instructor Overall rating is for me in particular.

Term	Course	Enrollment	Course Overall Rating	Instructor Overall Rating	Over 4.5
F99	API 109	65	3.33	3.61	
F99	API 111	33	3.5	3.56	
S00	ASI 112	17	4.29	4.29	
F00	API 109	58	4.06	4.39	
F00	API 111	33	4.06	4.33	
S01	API 112	25	4.18	4.5	*
F01	API 109	42	4.27	4.58	*
F01	API 109	41	4.5	4.67	*
F01	API 111	57	4.07	4.19	
S02	API 112	31	4	4.16	
S03	API 112	42	4.19	4	
F03	API 109	69	4.56	4.77	*
F03	API 111	59	4.23	4.61	*
S04	API 112	43	4.06	4.53	*
F04	API 109	70	4.28	4.63	*
F04	API 111	72	4.15	4.35	
S05	API 112	45	4.15	4.58	*
F05	API 109	71	4.15	4.42	
F05	API 111	56	3.5	4.56	*
F06	API 109	75	4.33	4.59	*
S07	API 112	34	4.26	4.6	*
F07	API 109	58	4.43	4.66	*
F07	API 111	59	4.60	4.70	*
S08	API 112	50	4.43	4.71	*

IIIB. Other Teaching Experience at Harvard

In addition to my regular teaching assignment, I also fulfill several ad hoc assignments. For example, for the past few years I have taught a session of Joe Newhouse's second-year reading course for the health-policy economics-track students. When required, I also write and administer the Microeconomics Qualifying Exam for the public policy PhD program, and I frequently serve on examination committees for public policy PhD students.

I have also recently begun teaching in KSG's Executive Programs. For the last two years I have taught sessions on game theory and strategy in one of our executive programs. The ratings for 2008 are net yet available. In 2007, my rating for these sessions was 4.5/5, which was at the average of all sessions for the program. Given that this was my first time teaching in executive programs and the inherent complexity of the material (most of the other sessions are on non-technical topics), I believe the sessions went very well. Further, I found them to be quite enjoyable, and I hope to continue to be involved in executive program teaching in the future.

Informally, I do a lot of work with our PhD students. Since I teach one of the few required courses for all public policy PhD students, I am one of the few faculty members who actually meets all of them, and I usually teach them for an entire year. Because of this, I frequently consult with students as they begin their research. As I mentioned above, our students' research tends to use concepts from economic theory, but they are not theorists. As a result, I frequently meet with students to discuss how they could integrate

a model into their analysis. I also attend practice job market talks and conduct mock job market interviews.

Since most of our students are not doing explicitly theoretical research, I tend to be more of an informal advisor than a formal one. Nevertheless, I have served on the doctoral committees of three students (Lori Snyder (PPOL), Steve Anderson (PPOL), and Alex Wagner (PEG)). I have also co-authored and published papers with three PhD students (Lori Snyder, Alex Wagner, and Nikita Piankov (Economics)).

I have found that one of the ways in which our PhD students' training falls short is in accustoming them to the academic culture. I have also felt that our students are relatively unprepared to enter the academic job market – not in terms of the research, but in terms of understanding how the process works. To address this, over the past two years I have held a series of "post-game" discussion sessions with our PhD students (PEG and PPOL) after the junior analytics search seminars. During the discussions, we talk about what went right and what went wrong during the particular seminar, but we also discuss more general job market strategies (i.e., what kinds of papers tend to go over better, what happens during fly-outs, etc.). My impression is that these meetings, usually attended by 6 – 12 students, have been very useful for them.

IIIC. Teaching Experience Prior to Harvard

Before coming to Harvard, I taught at Kellogg, offering two sessions of managerial economics (one for the full-time MBA program, one for the part-time MBA program), as well as teaching the math review course for students in Kellogg's Executive MBA program.

IV. Book Project: Applied Microeconomic Theory

My interest in making connections between the standard tools of microeconomic theory and policy-relevant applications has led to the production of the textbook in *Applied Microeconomic Theory* that I have been writing with Chris Avery. Too often the standard, doctoral-level courses in microeconomic theory are intended, or taught as if they were intended, to train future economic theorists. Even in economics PhD programs, most students will not be theorists but rather applied researchers who use the tools of economic theory to guide their thinking in (often) empirical studies, which are the predominant form of doctoral dissertation. While this is true in economics departments, it is even more true for KSG PhD students, virtually all of whom go on to study applied topics such as health, education, development, and environment, and for our masters students interested in advanced microeconomic analysis (e.g., MPA/IDs, MPA2s, and MPPs). These students are going to be focused on applied problems in public policy and/or international development. As such, they need a course in microeconomic theory that is going to help them make connections between the theory and real problems. The goal of this textbook is to meet this need.

The project began as a series of notes intended to help students in API 109 (MPA/ID microeconomic theory) and API 111 bridge the gap between the technical level of PhD level textbooks and their background (not usually economics majors) and interests (not interested in theoretical generalities for their own sake). Over time, the notes have incorporated applied material on topics such as the economic evaluation of welfare, and the impact of market imperfections such as those that are likely to be found in less-developed economies. These notes (which comprise the first half of the textbook) have been used in my fall classes for the last few years, and they have been available on my web page for the last two. They are widely used throughout the world as a guide to understanding economic theory. The web page has received approximately 15,000 hits over the past year, and I have received (voluntary) feedback from students and faculty at over 85 countries (over 300 foreign and 90 domestic institutions) since I started collecting this information in November of 2005.

Over the past year, Chris Avery and I have labored to turn these notes into the first half of a textbook, which we are calling *Applied Microeconomic Theory*, and he has joined me in writing new material based on what we cover in API 112. We have a nearly complete draft (530 pages) of the textbook at this point. Harvard University Press has agreed to publish the book, and we are in the process of signing a book contract. The intended audience of the book include students in applied doctoral programs and advanced undergraduate or masters programs.

Bibliography

Section IIA: Information, Incentives, and Welfare.

"Efficiency in Partnerships with Joint Monitoring," *Journal of Economic Theory*, **77** (2), December 1997, 285-299.

"Eliciting Informative Feedback: The Peer Prediction Method," with Paul Resnick and Richard Zeckhauser (2005). *Management Science*, **51**(9), September 2005, 1359-1373.

"Mechanism Design with Multidimensional, Continuous Types and Interdependent Valuations," with Scott Johnson, John Pratt, and Richard Zeckhauser (2006). Forthcoming, *Journal of Economic Theory*.

"Possibly-Final Offers," with Nikita Piankov and Richard Zeckhauser (2006). *Journal of Economics and Management Strategy*, **15(**3), Fall 2006, 789-819.

"Screening Budgets," with Alex Wagner and Richard Zeckhauser (2006). *Journal of Economic Behavior and Organization*, **61**(3), 351-374.

"Screening Through Risk Aversion," with Alex Wagner and Richard Zeckhauser (In Progress).

Section IIB: Industrial Organization/Representing Strategic Interactions

"Advertising Budgets in Competitive Environments," with Amit Pazgal (2007). *Quantitative Marketing and Economics* Vol. 5:2 pp. 131–161.

"Budget or Target: The Choice Between Input and Output Strategies," with Amit Pazgal (2006). R*AND Journal of Economics*, **37**(2), Summer 2006, 391-415.

"The Effects of Environmental Regulation on Technology Diffusion: The Case of Chlorine Manufacturing," *American Economic Review*, **93**(2), May 2003 (Papers and Proceedings), 431-435, with Lori D. Snyder and Robert N. Stavins.

"The Equivalence of Price and Quantity Competition with Delegation," *RAND Journal of Economics*, **32**(2), Summer 2001, 284-301, with Amit Pazgal.

"Relative Performance as a Strategic Commitment Mechanism," *Managerial and Decision Economics*, **23**(2), March 2002, 51-68, with Amit Pazgal.

"Strategic Trade and Delegated Competition," with Amit Pazgal. *Journal of International Economics*, **66**(1), 2005, 215-231.

Section IIC: Industrial Organization/Incentives in Health Care

"Demanding Patients: Consumerism and the Quality of Care," with Hai Fang, John Rizzo (Stonybrook University) and Richard Zeckhauser (In Progress).

"Insurer-Provider Integration, Credible Commitment, and Managed Care Backlash" (2006). *Journal of Health Economics,* **25**(5), 861-876.

"Pricing Health Benefits: A Cost Minimization Approach," (2005). *Journal of Health Economics,* **24**(5), 2005, 931-949.

"Provider Choice of Quality and Surplus," with Karen Eggleston and Richard Zeckhauser (2006). *International Journal of Health Care Finance and Economics,* **6**, 103-117.

"Report Cards, Incentives, and Quality Competition in Health Care," (Working Paper).

Section IID: Applications of Microeconomic Theory

"Giffen Behavior and Subsistence Consumption," with Robert Jensen (2008), *American Economic Review*, 98(4), 1553-1577.

"Do Consumer Price Subsidies Really Improve Nutrition?" with Robert Jensen (2008). Revision requested at *The Review of Economics and Statistics.*

"The Impact of the World Food Price Crisis on Nutrition in China," with Robert Jensen (2008). Forthcoming, *Agricultural Economics*.

"The Role of Information Framing in Environmental Health Policy: A Field Experiment," with Lori Bennear (Duke University) (In Progress).