Surplus Maximization and Optimality

Edward E Schlee*
Department of Economics
Arizona State University
ed.schlee@asu.edu

February 23, 2012

Abstract

Expected consumer's surplus rarely represents a consumer's preferences over price-income lotteries. Still, I find that policies which maximize expected surplus are *interim* Pareto Optimal under four assumptions. Two are strong but standard partial equilibrium assumptions: the policy affects only one price; and income changes do not affect demand. The two others are that every consumer's indirect utility function satisfies increasing differences in price and income; and policies order prices by a single-crossing property. I argue that increasing differences is plausible. The single-crossing property appears strong, but holds in important applications. I use the result to extend some well-known welfare results beyond the knife-edge case of quasilinear utility.
JEL Classification: D61, D8.
Keywords: Welfare under Uncertainty, Consumer's Surplus, Risk Aversion, Optimal Regulation, Single-Crossing Property, Increasing Differences, Quasilinear Utility.

1 Introduction

Consumer's surplus cannot represent a consumer's preferences over *all* price-income pairs (Chipman and Moore (1976), p. 80, building on Samuelson (1942) and others). Perhaps worse, Chipman and Moore (1976 and 1980) show that even if we severely restrict allowable price-income pairs, consumer's surplus only represents a knife-edge class of preferences. Yet if there is no uncertainty, policies which maximize aggregate surplus just equate prices and marginal costs, a necessary condition for Pareto Optimality of a production plan, absent distortions. For example, per-unit taxes should be set to zero, and a monopolist should be subsidized until price equals marginal cost. This fact surely helps explain why applied economists often use aggregate surplus to evaluate policies. Another surely-important fact is that one only needs to know aggregate demand to calculate it, not individual preferences and endowments.

It is easy to see how surplus-maximizing policies can be Pareto Optimal if there is no uncertainty: if prices and incomes vary smoothly with the policy (for example, taxes) then the

*I thank Hector Chade, Kerry Smith, and conference participants at the 2010 Society for the Advancement of Economic Theory meetings in Singapore and the 2010 Public Economic Theory conference in Istanbul for comments.

change in consumer's surplus for small policy changes is proportional to the change in utility: the change in consumer's surplus *locally* represents a consumer's preferences over policies, no matter what the preferences. It follows that the first-order conditions for maximizing aggregate surplus are the same as those for maximizing a particular weighted sum of consumer utilities.

With uncertainty and incomplete markets, the outlook for a positive result seems bleak: even if incomes and the prices of all but one good are fixed, the change in expected consumer's surplus is not proportional to change in expected utility unless the marginal utility of income is constant in the one price (Rogerson (1980) and Turnovsky, Shalit, and Schmitz (1980)).[1] And policies which maximize expected surplus often lack the immediate intuitive appeal of "price = marginal cost." For example, Deneckere, Marvel, and Peck (1997) find that letting a monopolist manufacturer impose a price floor for retailers can raise expected aggregate *and* expected consumers' surplus. And Baron and Myerson (1982) find that optimal regulation of a monopolist with private cost information sometimes puts positive probability on prices that are higher than the *highest* possible unregulated monopoly price. Curiously, defenses of aggregate surplus in Industrial Organization texts only consider economies with complete markets and no uncertainty—e.g. Tirole (1988), Introduction, and Vives (1999), Chapter 3—although those texts, and the applied microeconomics literature more generally, are filled with models with uncertainty and incomplete markets.

Still, I show that expected-surplus-maximizing policies are (*interim*) Pareto Optimal under four conditions (Theorem 1). Two are standard if strong partial equilibrium conditions: the policy only affects one price, and policy-induced income changes do not affect demand. The other two are more novel in this literature, one on preferences, the other on how policies affect prices. The preference condition is that each consumer's indirect utility function satisfies *increasing differences* in the price of the good and money, or equivalently the marginal utility of money is increasing in the price. It turns out that increasing differences holds if and only if relative risk aversion for money lotteries is at least as large as the budget elasticity of demand for the good. Since budget elasticities must on average equal 1, I argue that this condition is empirically plausible. The policy condition is that an expected surplus maximizing policy generates a random price which is "less variable" in a sense to be made precise than any other policy. The condition seems strong, but I show it often holds in applications. I argue that, roughly, it is more likely to hold if uncertainty is about consumers rather than producers (Section 3.2). I illustrate my results with some important I.O. applications (Section 4). For example I show that if the indirect utilities satisfy increasing differences, minimum resale price maintenance in Deneckere, Marvel, and Peck (1997) remains Pareto Optimal, but the policy identified by Baron and Myerson (1982) is not—indeed all consumers and the firm can prefer no regulation at all to the Baron-Myerson policy.

The models of Baron and Myerson (1982) and Deneckere, Marvel, and Peck (1997) share two things with many in I.O. and applied microeconomics. First, they are *partial equilibrium* models: the policy only affects one relative price, and income changes do not affect demand. Second, consumers' attitudes towards the uncertainty in the model are irrelevant: *by the time consumers act, all the uncertainty in the model is resolved*, so their choices reveal nothing about attitudes towards this uncertainty. These attitudes *are* needed to evaluate policies chosen before the uncertainty is resolved.

[1] I exploit this fact (Schlee, 2008) to show that the percentage error from using expected consumer's surplus to approximate the willingness to pay for a price change is unbounded, in contrast to Willig (1976)'s famous approximation result in the certainty case (Vives (1987) and Hayashi (2008) substantially extend Willig's result). Here I am not concerned with *measuring* the welfare change, but with finding optimal policies.

The closest paper is Rogerson (1980), who determines whether the expected-surplus ranking of policies which affect a good's price is *Pareto Consistent*, namely, whether expected surplus rises whenever *all* consumers prefer a change. He finds it is if and only if each consumer's indirect utility is additively separable in the price of the good (and any other variables that are random), a much stronger condition than increasing differences. He also concludes that expected surplus is apt to perform better if the uncertainty is about producers rather than consumers. I defer a fuller discussion to Sections 3.1.1 and 5, but briefly the two main differences are that Rogerson (1980) considers *ex ante* rather than *interim* welfare; and he puts *no* restriction on how policies affect the price distribution. My contribution is to find a restriction on policies so that expected-surplus-maximizing policies are Pareto Optimal for a plausible class of preferences. Since the class puts no upper bound on consumer risk aversion, the result is especially useful for applications in which risk aversion is important.

2 Certainty: Complete Markets Benchmark

I begin with a benchmark economy with no uncertainty and complete markets. My goal here is to write down sufficient conditions for surplus-maximizing polices to be Pareto Optimal (Proposition 2), a result I extend to economies with uncertainty and incomplete markets. I also point out that other simple sufficient conditions cannot easily be extended.

The model is close to the one Radner (1993) uses for local cost-benefit analysis. The economy has I consumers, L goods ($L \geq 2$), and a finite number of firms. Good L is the *numeraire*. All consumers face the same prices, which they take as given. My first restriction on how policies affect prices is in the spirit of partial equilibrium: policies only affect the price of good 1 (relative to other goods) and incomes. I maintain the restriction throughout. Consumer i's income $m_i > 0$ equals the value of i's endowment plus i's profit income. Each consumer's preferences over goods are represented by a continuous, strongly monotone, and strictly quasiconcave utility function, so demand is single-valued and satisfies budget balance. Consumer i's demand for good 1 is $D(p, m_i, \theta_i)$, where θ_i is a preference parameter and p the price of good 1. (Since the prices of other goods do not depend on the policy I suppress them as arguments.) Consumer i's preferences over price-income pairs (p, m_i) are represented by an indirect utility function $V(p, m_i, \theta_i)$.[2] I assume that each consumer's indirect utility function is C^2 and use subscripts on V to denote partial derivatives: V_p is the derivative of V with respect to the price of good 1, and V_m is the derivative of V with respect to money (income or wealth, depending on the application). I assume that $V_m > 0$ globally.

I index policies by a real number α in a set $\mathcal{A}$ of available polices; in this section I set $\mathcal{A} = [0, 1]$. As mentioned, consumers are price-takers, but I do not explicitly model producers or define equilibrium other than to say that excess demand is nonnegative, and that equilibrium consumption equals demand for each consumer (no consumer is rationed, though firms might be). The broad equilibrium concept allows me to cover a large literature with a range of assumptions about market structure, from competition to monopoly and regulated monopoly. I assume that for each policy $\alpha \in \mathcal{A}$, there is a unique equilibrium price of good 1, $\widetilde{p}(\alpha)$, and distribution of income, $(\widetilde{m}_1(\alpha), ..., \widetilde{m}_I(\alpha))$. The uniqueness assumption is strong, but consistent with the partial equilibrium applications my results cover. In this section I also assume that the equilibrium price, incomes, and consumer utility functions are all C^1. Since

[2] The indirect utility gives the maximum utility from the consumer's budget. Section 3.D of Mas-colell, Whinston, and Green (1995) briefly review properties of indirect utility functions and demands.

consumers face the same prices, the allocation of the aggregate consumption is Pareto Optimal, so the only issue is optimality of the aggregate consumption plan. Consumer i's preferences over policies in $\mathcal{A}$ are represented by the function $\alpha \mapsto V(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)$. As is common, I define the change in consumer i's surplus between policy 0 and α to be the line integral[3]

$$S_i(\alpha) = -\int_{[0,\alpha]} D(\widetilde{p}(v), \widetilde{m}_i(v), \theta_i) \cdot d\widetilde{p}(v) + \widetilde{m}_i(\alpha) - \widetilde{m}_i(0). \tag{1}$$

If consumer i's income does not depend on the policy, then (1) collapses to the familiar "area to the left of the demand curve." To see why (1) might be an appropriate extension of the area definition when policies affect income, differentiate it to find $S_i'(\alpha) = \widetilde{m}_i' - D^i\,\widetilde{p}\,'$. The expression $-S_i'$ is the local *Slutsky compensation* for a small policy change, the change in money income that keeps the original choice just affordable.[4] (Up to an additive constant, (1) is the *unique* function with the derivative property.) In that sense S_i' is the change in the consumer's *real income* from a small policy change.

Since I only consider a single path of prices and income—the one generated by the policy—I bypass the usual problem that the change in consumer's surplus between two price-income pairs depends on the path chosen. And since I only consider price-income pairs along this path, I don't need to recover preferences over *all* price-income pairs to find an optimal policy.

2.1 Local Policy Preferences: Certainty vs. Uncertainty

Differentiate consumer i's equilibrium indirect utility with respect to α and use Roy's Identify ($D = -V_p/V_m$) and (1) to find[5]

$$\frac{d}{d\alpha}V(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i) = V_p^i\,\widetilde{p}\,'(\alpha) + V_m^i \widetilde{m}_i'(\alpha) = V_m^i\left[-D^i\,\widetilde{p}\,'(\alpha) + \widetilde{m}_i'(\alpha)\right] = V_m^i \times \frac{d}{d\alpha}S_i, \tag{2}$$

where $V^i = V(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)$ and $D^i = D(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)$; or in short,

$$\frac{dS_i}{d\alpha} = \frac{dV^i/d\alpha}{V_m^i}. \tag{3}$$

Equation (2) just says that the change in utility from a small policy change equals the change in the marginal utility of money times the change in real income (in the sense of a Slutsky compensation). Equations (2) and (3) formalize the fact already mentioned, that the change in consumer i's surplus *locally* represents i's preferences over policies: $dV^i/d\alpha > 0$ if and only if $dS_i/d\alpha > 0$. It also illustrates a potential problem with *expected* surplus: with $E[\cdot]$ the expectations operator, $E\left[dS_i/d\alpha\right] > 0$ does not imply that $E[dV^i/d\alpha] > 0$ unless the marginal utility of money is constant in any variables that are random, as Rogerson (1980) points out. In other words, with uncertainty and incomplete markets, if *expected* real income locally increases with the policy it does not follow that *expected utility* does.

[3] Chipman and Moore (1980), equation (9); Silberberg (1972), pages 943-4; Just, Hueth, and Schmitz (2004), page 116, equation (5.7); and Slesnick (1998), p. 2111, equation (4). When the policy affects incomes, some scholars define surplus differently, specifically in what to insert into the income argument of demand in the integrand. For example, Blackorby (1999) uses the *final* income, $\widetilde{m}(\alpha)$, while Hammond (1990) uses the *initial* income, $\widetilde{m}(0)$. On either of these definitions both equation (3) and the conclusion of Proposition 1 fail—an argument in favor of (1).

[4] See Mas-colell, Whinston, and Green (1995), pp. 29-30, 72-3 for discussion of the Slutsky compensation.

[5] Roy's Identity holds under my assumptions that preferences are monotone and strictly convex and that the indirect utility function is C^1 with $V_m > 0$. See for example Mas-colell, Whinston, and Green (1995), p. 73.

2.2 Optimality of Surplus-Maximizing Policies

Let $S(\alpha) = \sum S_i(\alpha)$ denote aggregate surplus. Sum equation (3) over all consumers to find

$$S'(\alpha) = \sum \frac{\frac{d}{d\alpha} V(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)}{V_m(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)}. \tag{4}$$

It is immediate from (4) that if $S'(\widehat{\alpha}) > 0$, then at least *one* consumer's preferences are strictly increasing in α in a neighborhood of $\widehat{\alpha}$. The next result strengthens the conclusion to a *potential Pareto improvement*: after a small-enough increase in α from $\widehat{\alpha}$, the new supply can be allocated so that *every* consumer prefers the higher policy. Recall that excess demand is the difference between demand and supply (the sum of production and the endowment).

Proposition 1. *Suppose that for every $\alpha \in \mathcal{A}$, the equilibrium production plan is unique and differentiable, equilibrium excess demand is 0, and the equilibrium allocation is interior. If $S'(\widehat{\alpha}) > 0$ for $\widehat{\alpha} \in [0, 1)$, then a small-enough increase in α from $\widehat{\alpha}$ is a potential Pareto improvement.*

The proof (in Appendix A) uses a result in Radner (1993) which gives a sufficient condition for a small policy change to be a potential Pareto improvement.[6]

A policy $\alpha' \in \mathcal{A}$ is *Pareto Optimal* if there is no other policy in $\mathcal{A}$ that every consumer weakly prefers to α' and some consumer strictly prefers to α'. Recall that if a policy maximizes a positive weighted sum of utilities, then it is Pareto Optimal. Suppose that an interior point α^* maximizes aggregate surplus S on $\mathcal{A}$, so that $S'(\alpha^*) = 0$. Consider a weighted sum of indirect utilities, with weight on consumer i's indirect utility equal to the reciprocal of the marginal utility of income, $1/V_m$, evaluated at $(\widetilde{p}(\alpha^*), \widetilde{m}_i(\alpha^*), \theta_i)$. By (4), the derivative of this weighted sum with respect to α equals 0 at $\alpha = \alpha^*$: In words, if an interior policy maximizes surplus, then the *first-order condition* for maximizing a particular weighted sum of indirect utilities holds. The coincidence between first-order conditions partly explains why surplus-maximizing policies are often Pareto Optimal in complete markets. Of course comparing first-order conditions is not enough to *prove* that surplus-maximizing policies are Pareto Optimal: to pass rigorously from local to global optimality, the economy or the policies must be restricted. With uncertainty and incomplete markets, this line of reasoning seems hopeless, since as noted the analogue of equation (3) which lies behind the coincidence fails.

If each S_i is *strongly quasiconcave* ($S_i'(\alpha) = 0$ implies that $S_i''(\alpha) < 0$) then any maximizer of aggregate surplus S is Pareto Optimal (Appendix B, Section 8.1). Unfortunately, there is also little hope for extending this result to the case of uncertainty and incomplete markets (unless the same policy maximizes surplus in all states). The result I extend imposes more than quasiconcavity or even monotonicity of surplus. The next assumption, again in the spirit of partial equilibrium, restricts policies and the preferences of some consumers so that any policy-induced income changes do not affect demand.

Assumption 1. *(**No Income Effects**) For some nonnegative integer $I_C \leq I$ we have*

(i) for $i = 1, ..., I_C$, consumer i's income does not depend on the policy;

[6] It also follows that the change in surplus is locally proportional to the *coefficient of resource utilization* (Debreu (1954)), the welfare loss from a policy expressed as a fraction (ρ) of the economy's endowment: apply (5) in Section 3 of Debreu (1954) to find that $\rho'(\alpha) = \sum \frac{dV^i/d\alpha}{V^i_m} \times (\sum \widetilde{m}_i(\alpha))^{-1}$, which by equation (4) is proportional to S'.

(ii) for any $i = I_C + 1, ..., I$, consumer i's demand for good 1 does not depend on income (equivalently, i's indirect utility function is affine in income).

Part (i) holds if consumers $i \leq I_C$ have no endowment of good 1 and own no shares of any firm whose profit is affected by the policy. The assumption is strong, but is weaker than the common partial-equilibrium practice of simply dividing agents into "consumers" and "producers." Mas-colell (1982) for example splits the economy into two groups (see his Section 3.2), producers, who own the firms and consume *only* the *numeraire* good; and consumers who do not own any share of any firm. The No Income Effects assumption allows 'producer-households' to care about the consumption of all goods, and it allows 'consumer-households' to have own shares in some firms. The next result barely requires a proof since Assumption 1 ensures that the change in consumer's surplus represents each consumer's policy preferences. (It follows immediately from Theorems 1 and 2 in Chipman and Moore (1976)).

Proposition 2. *If the No Income Effects assumptions holds then any surplus-maximizing policy is Pareto Optimal.*

My main task is to find conditions on preferences and policies which extend the conclusion of Proposition 2 to economies with uncertainty and incomplete markets.

2.3 Related Results

Facchini, Hammond, and Nakata (2001) consider a competitive production economy with a tax on a single good, and show that the Pareto Optimal tax of zero can locally *minimize* surplus—but not that a surplus-maximizing policy is not Pareto Optimal. Blackorby (1999) considers a single consumer economy, and finds that aggregate surplus is (locally) maximized at a Pareto Optimal policy if and only if the consumer has preferences which are quasilinear in a *numeraire* good (Theorem 1). There are several differences between his Theorem 1 and Propositions 1 and 2. First, he considers a single-consumer economy. Second, he allows *all* relative prices to change, but defines surplus using a demand function of one good only (any non-*numeraire* good, say good 1). Third, as already mentioned (note 3), he defines the change in surplus to be $\int_0^\alpha D(\widetilde{p}(v), \widetilde{m}(\alpha))dv + \widetilde{m}(\alpha) - \widetilde{m}(0)$, while I define it to be (1): specifically $\widetilde{m}(\alpha)$ appears in his integrand, rather than $\widetilde{m}(v)$ in mine. (The appearance of $\widetilde{m}(\alpha)$ in his integrand forces the consumer to have quasilinear preferences.) Of course the two definitions coincide under the No Income Effects Assumption. I escape the quasilinear implication in Proposition 2 by relaxing the single consumer assumption and imposing the No Income Effects Assumption.

3 Uncertainty: Incomplete Markets

I now turn to the main issue, economies with uncertainty and incomplete markets. Incompleteness is important: With complete markets the analysis of Section 2 applies, with goods distinguished according to the state of the world as in Debreu (1959), Chapter 7.[7] From now on I allow tastes, incomes, and technologies to depend on the realization of an unknown state

[7] The assumption in Section 2 that the policy only affects one relative price seems especially limiting if goods are classified by state of the world. But equation (3) and Proposition 1 both hold if the policy affects all relative prices: just replace the good-1 price and demand in the derivations by the *vector* of prices and demands and replace scalar multiplication of prices and demands by the inner product. The crucial assumption is that markets are complete, not that the policy only affects one relative price.

of the world $\omega \in \Omega$ where $\Omega \subset \mathbb{R}$ is a nonempty, bounded set of states. Let F be a cumulative distribution function (c.d.f.) on Ω. (Some applications require an infinite number of states; otherwise there is no harm of thinking of Ω as a finite set.) I assume that the policy is chosen before the state is known. Spot markets open only after the state is realized, hence markets are incomplete. As before, consumers take prices as given and are not rationed in any state, though firms might not be price takers and might be rationed. Again in the spirit of partial equilibrium—and to match the positive predictions of the partial equilibrium Applications I consider in Section 4—I assume that only the relative price of good 1 is affected by either the policy or the state of the world. I write $\widetilde{p}(\alpha, \omega)$ for the equilibrium price of good 1 in state ω under policy α. From now on I assume that, with probability one, the equilibrium price of good 1 lies in an interval $P = [\underline{p}, \overline{p}]$ for every $\alpha \in \mathcal{A}$, where $0 < \underline{p} < \overline{p}$.

I assume that each consumer has expected utility preferences, and interpret the indirect utility functions as von Neumann-Morgenstern utilities over price-income lotteries. I impose the maintained assumptions of Section 2 in every state of the world: each vN-M indirect utility is C^2 with a positive income derivative globally. I also impose an extension of the No Income Effects Assumption to economies with uncertainty. The only difference is that I assume that consumers $i > I_C$ are also risk neutral over money lotteries.

Assumption 2. ***(No Income Effects-Uncertainty)*** *For some nonnegative integer $I_C \leq I$:*

(i) for $i = 1, ..., I_C$ consumer i's income is constant in the policy α for each state ω;

(ii) for $i = I_C + 1, ..., I$, consumer i's demand for good 1 is constant in income and *i is risk neutral over money lotteries (equivalently, i's vN-M indirect utility function is quasilinear in income).*

To repeat, the No Income Effects Assumption is most plausible if consumers $i \leq I_C$ do not own any shares of firms who produce good 1. Kihlstrom and Laffont (1979) propose a model of endogenous entrepreneurship which gives a foundation for the risk neutrality part of Assumption 2. In their model everyone has a labor endowment and access to a risky technology. In equilibrium only the least risk averse people become entrepreneurs (use the risky technology); the equilibrium is Pareto Optimal only if all entrepreneurs are risk neutral.

To model demand uncertainty I introduce consumer types. Anticipating my imposition of Assumption 2, I first consider consumers $i = 1, ..., I_C$. For these consumers a type is a pair $(m, \theta) \in T = M \times \Theta$, where as before m is income and θ is a taste parameter. M and Θ are each finite subsets of the real line. The type of consumer $i \leq I_C$ depends on the state of the world through a function $\tau^i : \Omega \to T$ from the set of states into the set of types. The common prior belief about the state is F and I assume throughout that each consumer's preference parameter and income are statistically independent. Once a consumer i learns its type, i updates beliefs. I denote a type-t consumer i's posterior c.d.f. by $F(\cdot \,|\, i, t)$.[8] The policy preferences of a consumer i who is type t are represented by $\alpha \mapsto \int V(\widetilde{p}(\alpha, \omega), t) dF(\omega \,|\, i, t)$. I should stress that these preferences are *interim* preferences, *after* i has learned its type t. A particular type of consumer i is uncertain about price because that type does not know either the types of other consumers or the technologies of firms.

[8]By Bayes's rule, $F(\omega \,|\, i, t) = (1/\lambda^i(t)) \int_{\Lambda(\omega,i,t)} dF(\xi)$, where $\lambda^i(t)$ is the probability that consumer i is of type t and $\Lambda(\omega, i, t)$ is the set of states no larger than ω but consistent with consumer i being of type t. Formally, $\lambda^i(t) = \int_{\{\omega \in \Omega \,|\, \tau^i(\omega) = t\}} dF(\omega)$ and $\Lambda(\omega, i, t) = \{\xi \in \Omega \,|\, \tau^i(\omega) = t \text{ and } \xi \leq \omega\}$.

Under the No Income Effects–Uncertainty Assumption, the preferences of each consumer $i = I_C + 1, ..., I$ are represented by the expected consumer's surplus for i, so these consumers pose no problem for the optimality of policies which maximize expected surplus. Accordingly I don't dwell on the notation for types for these consumers. But formally a type for one of these consumers is a point $t = (\mathbf{e}, \theta)$ where $\mathbf{e}$ is a vector of the endowment of goods and ownership shares of firms, and θ is a preference parameter. Such a consumer's income is a function $f(\mathbf{e}, p)$ of the endowment and the price of good 1. Since the price of good 1 could affect the value of the endowment, income can depend on both the state and the policy. In what follows I assume that each consumer type occurs with positive probability for at least one consumer (otherwise that type can be deleted from the set of types).

My broad notion of equilibrium—simply that consumers are not rationed in any state—allows me to cover mechanism design problems as well as decentralized market clearing models. Since the types are determined by the state ω, I can take α to be an index of the mechanism that a planner implements after eliciting types. In this case $\widetilde{p}(\alpha, \omega)$ is then the price the planner chooses under mechanism α for the type profile at state ω. The optimal regulation papers I discuss in Section 4.3 fit this interpretation.

3.1 Optimality and Expected Surplus Maximization

Rogerson (1980) and Turnovsky, Shalit, and Schmitz (1980) show that expected surplus can be *Pareto inconsistent*: expected surplus can be higher for policy A than B, even though every consumer prefers B to A. To illustrate, recall that the consumer's surplus integral $\int_p D(v)dv$ is strictly convex in p if demand is strictly decreasing in price. This fact implies that expected consumer's surplus rises with a mean-preserving increase in price risk (Waugh (1944)). Turnovsky, Shalit, and Schmitz (1980) point out that a consumer's indirect utility function can be strictly concave in the price of a single good, implying that that consumer's expected utility *falls* with a mean-preserving increase in price risk. It is easy to use this fact to construct a single-consumer economy in which income does not depend on either the policy or the state of the world to illustrate the Pareto inconsistency of expected surplus.

My main theorem on the optimality of expected-surplus-maximizing policies imposes the No Income Effects-Uncertainty Assumption. I impose two others, one on preferences, the other on how the policy affects prices. The preference assumption is that each consumer's indirect utility function satisfies *increasing differences* in price and income: the marginal utility of income $V_m(p, m)$ is increasing in p.[9] A C^2 indirect utility satisfies increasing differences in (p, m) if and only if relative risk aversion for money lotteries is at least as large as the budget elasticity for good 1: Differentiate Roy's Identity $D = -V_p/V_m$ with respect to income, multiply both sides by m/D and rearrange to find

$$rra - \varepsilon_m = \frac{m}{D}\frac{V_{pm}}{V_m}, \tag{5}$$

where $rra = -mV_{mm}/V_m$ is relative risk aversion for money lotteries and ε_m is the income elasticity of demand for good 1. Intuitively, increasing differences implies that risk aversion is more important than income effects on demand. The average budget elasticity across all goods

[9] A function f on the real line is *increasing* if $f(y) \geq f(x)$ whenever $y \geq x$. A function g on the product $I \times J$ of two subsets of the real line satisfies *increasing differences* if for all $x'' \geq x'$ in I, the function $g(x'', \cdot) - g(x', \cdot)$ is increasing on J. If I and J are intervals and g is C^2, then g satisfies increasing differences if and only if the cross partial derivative g_{xy} is nonnegative on $I \times J$.

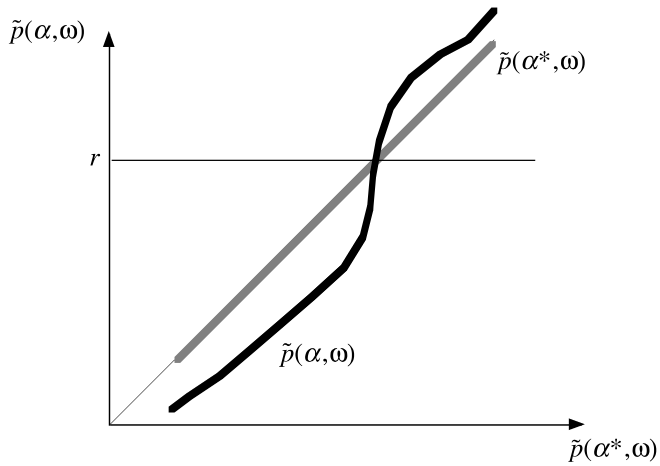

Figure 1: Single-crossing property: policy α^* leads to prices that are *less variable* than prices from policy α.

equals 1 by budget balance. Estimates of relative risk aversion over money lotteries vary, but they almost always exceed 1, with many much higher. Meyer and Meyer (2005) summarize the empirical evidence. The estimates vary partly because some take the outcome of the money lottery to be wealth, others to be income or consumption. For wealth, estimates mostly lie between 1 and 4; for consumption or income, they are usually *much* higher.[10] Although budget elasticities vary across goods, surely on average across goods they are less than relative risk aversion for money lotteries, so increasing differences is surely the typical case.[11]

The policy restriction is that expected-surplus-maximizing policies generate random prices that are *less variable* than other policies in the sense that, if α^* maximizes expected surplus, then $\widetilde{p}(\alpha^*, \cdot)$ crosses each random price from other policies at most once from above (Figure 1): formally, for each $\alpha \in \mathcal{A}$, there is an $r \in P$ such that

$$[\widetilde{p}(\alpha^*, \omega) - \widetilde{p}(\alpha, \omega)]\,[\widetilde{p}(\alpha^*, \omega) - r] \leq 0 \text{ for all } \omega \in \Omega. \tag{6}$$

If (6) holds, call r a *crossing price* for α^* and α. Suppose that the No Income Effects–Uncertainty holds. For $i = 1, ..., I_C$, a type $t = (m, \theta) \in T$ consumer i's *interim* expected surplus is (here $S(\alpha, \omega, t) = \int_{\widetilde{p}(\alpha,\omega)}^{\bar{p}} D(p, m, \theta)dp + m$ in an abuse of my previous notation)

$$E[S(\alpha, \omega, t) \,|\, i, t] = \int_{\Omega} \left[\int_{\widetilde{p}(\alpha,\omega)}^{\bar{p}} D(p, t)dp \right] dF(\omega \,|\, i, t) + m. \tag{7}$$

Lemma 1. *Consider two policies α^* and α in $\mathcal{A}$. Suppose that*

(i) the No Income Effects Assumption–Uncertainty holds;

[10] In more recent work, Cohen and Einav (2007) use information about insurance demand to estimate absolute and relative risk aversion with income as the lottery outcome. The estimates vary widely across individuals; the the mean individual has relative risk aversion equal to 97, the median, less than 0.5. As I explain later, the assumption that *every* consumer's indirect utility satisfies increasing differences can be substantially relaxed.

[11] Since it is a regulated good, electricity is a natural example for some of the applications I consider in Section 4. Studies agree that the short-run income elasticity is near 0 and the long-run elasticity is less than 1.

(ii) $\widetilde{p}(\cdot, \alpha^*)$ *crosses* $\widetilde{p}(\cdot, \alpha)$ *at most once from above (i.e. equation (6) holds); and*

(iii) $V_m(\cdot, t)$ *is increasing in the price of good 1 for* $t \in T$ *(increasing differences).*

Consider a consumer $i \leq I_C$ *of type* t. *If expected surplus at* α^* *is greater than or equal to expected surplus at* α *(* $E[S(\alpha^*, \omega, t) \mid i, t] \geq E[S(\alpha, \omega, t) \mid i, t]$ *), then a type-*t *consumer* i *weakly prefers* α^* *to* α. *If the inequality is strict, then the preference is strict.*

Proof: Appendix A.

To explain the proof informally, recall from equation (2) that the change in utility from the policy change locally equals the change in real income (= surplus) times the marginal utility of money: in short, and again informally, $\Delta V \approx \Delta S \times V_m$. (The exact expression is the integrand in the third line of equation (15)). Now suppose that the policy changes from α to α^*. By the single-crossing property, the price is lower under policy α^* in states of the world with prices above r, and higher in states with prices below r (Figure 1). Since the marginal utility of money is increasing in price, V_m is higher in states in which surplus rises and lower in states in which surplus falls. If $E[\Delta S] > 0$, then it follows that $E[\Delta S \times V_m] > 0$, so *expected* utility is higher at α^* than at α. It is worth noting that if the mean prices for all the policies are the same, then condition *(ii)* implies that the random price under α is a mean-preserving increase in risk of the random price under α^*. Condition *(iii)* however does not imply that V is concave in the price of good 1: the consumer need not be risk averse over price lotteries. Condition (ii) and the assumption that expected surplus is higher at α^* rule out the Turnovsky-Shalit-Schmitz (1980) example at the beginning of Section 3.1 of a mean-preserving increase in price risk that increases expected surplus but lowers expected utility.

Lemma 1 deals with the relationship between a single consumer type's *interim* policy preferences and *interim* expected surplus (after types are revealed). Of course a social planner does not know the types, and I assume that the planner maximizes *ex ante* expected surplus, calculated using the common prior F. My main result gives conditions under which a policy which maximizes *ex ante* aggregate expected surplus is *interim* Pareto Optimal.[12]

Theorem 1. *Suppose* α^* *maximizes* ex ante *aggregate expected surplus on* $\mathcal{A}$. *If*

(i) the No Income Effects Assumption–Uncertainty holds;

(ii) for every $\alpha \in \mathcal{A}$, $\widetilde{p}(\cdot, \alpha^*)$ *crosses* $\widetilde{p}(\cdot, \alpha)$ *at most once from above; and*

(iii) $V_m(\cdot, t)$ *is increasing in the price of good 1 for every* $t \in T$,

then α^* *is* interim *Pareto Optimal.*

Theorem 1 follows immediately from Lemma 1: If each type of each consumer (weakly) prefers $\alpha \in \mathcal{A}$ to $\alpha^* \in \mathcal{A}$, and the preference is strict for at least one consumer-type pair, then by Lemma 1 aggregate expected surplus is higher at α than at α^* and α^* cannot maximize expected surplus on $\mathcal{A}$.

Equation (5) makes clear that monotonicity of the marginal utility of money in the price is preserved under a strictly increasing, concave transformation ϕ: for $V = \phi(\widetilde{V})$, the income elasticity is the same and relative risk aversion larger for V than for $\widetilde{V}$. Indeed differentiate $V = \phi(\widetilde{V})$ with respect to m and then p and rearrange to find that

$$V_{mp} = \phi' \left[\frac{\phi''}{\phi'} \widetilde{V}_m \widetilde{V}_p - \widetilde{V}_{mp} \right]. \tag{8}$$

[12]A policy $\alpha^* \in \mathcal{A}$ is *interim Pareto Optimal* if there is no other policy $\alpha' \in \mathcal{A}$ such that each type of each consumer weakly prefers α' to α^*, and the inequality is strict for at least one consumer-type pair.

Recall that the quantity $-\phi''/\phi'$ is the Arrow-Pratt measure of (absolute) risk aversion of the function ϕ. Since $\widetilde{V}_p < 0$, it follows that if Arrow-Pratt risk aversion is uniformly high enough, then V_m must be increasing in p on P for every $m \in M$. The condition that V_m is increasing in price for *every* consumer type can be replaced by the condition that *some* consumer types are sufficiently risk averse if $\mathcal{A}$ is finite: as Arrow-Pratt risk aversion uniformly increases without bound, preferences converge to "max-min" preferences; that is, in the limit, the consumer ranks lotteries by comparing the *worst* possible outcome in each. Since the highest possible price under α^* is lower than the highest possible price under α when (6) holds, a sufficiently risk averse consumer would prefer α^* to α for every $\alpha \in \mathcal{A}$.

3.1.1 Related Results

Blackorby (1999) considers a related question. Suppose that some consumer prices are exogenously related to producer prices—for example, because of distortionary taxes—so that a competitive equilibrium is not first-best Pareto Optimal. Do second-best Pareto Optimal policy at least locally maximize surplus? He finds that they do not (Theorem 2). There are several differences between his Theorem 2 and my Theorem 1. First, the form of market incompleteness differs: in his there is no uncertainty, and some relative prices are exogenously fixed, whereas in mine there is uncertainty but only spot markets for goods. Second, although his result is local, he makes no other restrictions on how policies affect prices: all policy directions are considered. Third, as mentioned, his definition of surplus is different.[13]

Closer still is Rogerson (1980), who considers the same setting of uncertainty and incomplete markets as I do. He asks when expected consumers' surplus is *Pareto Consistent*—expected consumers' surplus rises whenever all consumers prefer a change—when the price of a single good is affected by a policy. He finds (Theorem 2) that it is if and only if the marginal utility of money for each consumer does not depend on the price of the good (or any other variable that is random). As mentioned, there are at least two important differences between his Theorem 2 and my Theorem 1: he considers *ex ante* welfare—expected utility is calculated before any uncertainty is resolved—rather than *interim* welfare; and he makes no restrictions on how policies affect the price of good 1. I consider *ex ante* optimality in Section 5. But I should point out if I allow arbitrary policies, the conclusion of Theorem 1 fails: it is easy to construct a single-consumer economy with an indirect utility function that satisfies strictly increasing differences in price and income and a pair of price policies so that expected utility and expected surplus rank the polices differently.[14]

3.2 Price Variability and Surplus Maximization

Theorem 1 imposes what seem to be two unlikely conditions: one policy generates a random price that crosses the price from any other policy at most once from above; and that policy

[13]Duhamel (2006) shows that second-best polices can maximize surplus if surplus is defined as a line integral over *all* prices, included those exogenously distorted, rather than over just the price of a single good, as in Blackorby (1999). Since Duhamel (2006) treats income the same way in the definition of surplus as Blackorby (1999), he too concludes that preferences must be quasilinear.

[14]This fact follows from Theorem 1 in Rogerson (1980) which implies that expected surplus represents a consumer's preferences over all lotteries on a subset of prices if and only if the marginal utility of money is constant in any variable that is random. For interim welfare, the only uncertainty a consumer faces is the price of good 1. A counterexample can be constructed with two policies with equilibrium price distributions such that expected consumer's surplus and expected utility rank them differently.

happens to maximize aggregate expected surplus. Often the conditions are easy to check directly. But two natural questions arise: first, why should policies generate prices that are ordered by the single-crossing property; and, second, when does the least variable of these random prices maximize aggregate expected surplus? In short, when does Theorem 1 apply?

The answer to the first question—why should policies generate random prices which are ordered by the single-crossing property—depends both on the model and the set of policies, and I have no general theorems on how scholars choose either. Still, something can be said about the relationship between the two, since the set of policies a modeler considers sometimes depends on the source of uncertainty in the model. If the state only affects demand, then the price of the good will be high in states that consumers value it highly; it is natural in that case to consider policy changes which increase output in high demand states, possibly by decreasing output in low-demand states—a change that would generate less variable prices. If the state only affects the technology, then price is high in states with high marginal cost; it is natural in that case to consider policy changes which increase output in states with low marginal cost, possibly by decreasing output in states with high marginal cost—a change which would generate more variable prices. So a modeler might naturally be led to consider policies ordered by a single-crossing property (as the Applications in Section 4 amply illustrate).

The next result partly answers the second question, when do policies with the least-variable prices maximize expected surplus. Let $Q(\cdot, \omega)$ be the aggregate demand for good 1 in state ω. I impose the condition that, at each policy, equilibrium output and price move in the *same* direction in response to different state realizations, that is, equilibrium output and price are *comonotonic* in the state (see the right panel in Figure 2): Formally, for each $\alpha \in \mathcal{A}$,

$$\left[Q\left(\widetilde{p}(\alpha, \omega'),\ \omega'\right) - Q\left(\widetilde{p}(\alpha, \omega),\ \omega\right)\right] \times \left[\widetilde{p}(\alpha,\ \omega') - \widetilde{p}(\alpha,\ \omega)\right] \geq 0 \tag{9}$$

for any ω, ω' in Ω. (If the inequality is strict whenever *either* of the terms in brackets on the left is nonzero, then the output and price are *strictly* comonotonic.) If aggregate demand is decreasing in price for each state, it follows that demand cannot be constant in ω for every fixed price: if the equilibrium price is higher in state ω' than state ω and aggregate demand is strictly decreasing in price, then aggregate demand cannot be constant in ω. Equation (9) and the assumption that aggregate demand is decreasing in price imply that aggregate demand is uncertain, so the case of pure supply uncertainty is precluded. The conclusion is that the policy with the least variable price maximizes what I call expected *Marshallian consumers' surplus*: expected surplus *minus* expected income, namely, $\int_\Omega \int_{[p(\alpha,\omega),\overline{p}]} Q(p,\omega) dp dF$.

Theorem 2. *Suppose that*

- *(i) aggregate demand $Q(\cdot, \omega)$ is nonincreasing in price for every $\omega \in \Omega$;*
- *(ii) equilibrium output $Q(\widetilde{p}(\alpha,\ \cdot),\ \cdot)$ and price $\widetilde{p}(\alpha,\ \cdot)$ are comonotonic for every $\alpha \in \mathcal{A}$; and*
- *(iii) there is a $\beta \in \mathcal{A}$ such that for any $\alpha \in \mathcal{A}$,*
  - *(a) the price $\widetilde{p}(\beta, \cdot)$ crosses $\widetilde{p}(\alpha, \cdot)$ at most once from above, and*
  - *(b) the mean price under β is no larger than the mean price from any other $\alpha \in \mathcal{A}$.*

Then β maximizes expected **Marshallian** *consumers' surplus on $\mathcal{A}$.*

Proof: Appendix A.

That the least variable price maximizes expected Marshallian consumers' surplus might seem surprising, since as mentioned at the beginning of Section 3.1, expected Marshallian

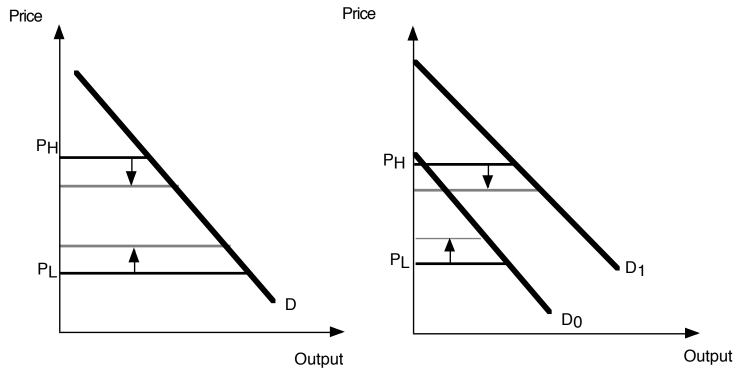

Figure 2: Both panels depict a mean-preserving contraction in state prices. The left shows a fixed demand, the right a variable demand in which demand and equilibrium price are comonotonic. With a mean preserving contraction in the price distribution, expected consumers' surplus falls in the left panel, but rises in the right.

surplus rises with a mean-preserving increase in price risk. The important assumptions are that the demand and equilibrium price are comonotonic, and that the mean price is lowest for the policy with the least variable price. To simplify, suppose that the mean price is the same for all policies. Intuitively, when equilibrium price and quantity demanded are comonotonic, expected consumers' surplus rises with a mean-preserving contraction in the price distribution, since price falls when quantity demanded is high and rises when quantity demanded is low. The right panel of Figure 2 helps illustrate the point: when price and demand are comonotonic, the gain in surplus when the high price p_H falls exceeds the loss of surplus when the low price p_L rises. On the other hand if the state only affects supply, then then the policy with the most variable price maximizes expected consumers' surplus (since consumer surplus is convex in price). Intuitively, when equilibrium price and *minus* quantity demanded are comonotonic, then expected consumers' surplus rises with a mean-preserving (or mean-reducing) spread, since price rises when quantity demanded is high and (possibly) falls when quantity demanded is low. The left panel of Figure 2 illustrates this more familiar point.

By Theorem 2, if the No Income Effects–Uncertainty assumption holds and the mean total income of consumers $i > I_C$ is the same for all policies, then the policy with the least variable price maximizes aggregate expected surplus. Some applications maximize a weighted average of expected Marshallian consumers' surplus and expected profit, usually with more weight on expected consumers' surplus. If the weight on consumers is high enough and the set of policies is finite, Theorem 2 implies that the policy with the least variable price maximizes this weighted average. If aggregate expected income is not constant and the planner does not weight expected Marshallian consumers' surplus heavily enough, then Theorem 2 does not help us determine what policy maximizes aggregate surplus. *If* it is possible to set price equal to marginal cost in every state of the world, then that policy would maximize aggregate expected surplus, but such a policy is typically not feasible with uncertainty and incomplete markets. I have not been able to prove a useful general result on what policies maximize expected *aggregate* surplus, since applications differ in what policies are feasible.

I instead consider the question in the context of an example, though an important one, a variant of the model Weitzman (1974) uses to answer his "Prices vs. Quantities" question: If there is uncertainty about both cost and demand and a planner can regulate *either* price *or* quantity, then which gives higher expected surplus? I use Weitzman's notion of a fixed quantity —the price simply adjusts to clear the market—but I use a different notion of a fixed price: he requires that firms are price takers and choose output to maximize profit; I require that firms meet all the demand at the fixed price. My notion of a fixed price is in the spirit of price regulation of firms. The change in the definition of a fixed price dramatically changes the conclusion: as in Weitzman (1974), what determines whether it is better to set prices or quantities is the *sum* of the slopes of the inverse demand and the marginal cost, but here the sign is reversed: it's better to fix prices rather than quantities if and only if the sum is *negative.* Example 1 shows that if cost is not too convex, aggregate expected surplus is higher with a fixed price than with a fixed quantity.

Example 1 (Prices vs. Quantities, after Weitzman (1974))**.** *Let the cost of good 1 in state* ω *be* $A(\omega)q + \frac{1}{2}Bq^2$ *and the inverse demand for good 1 in state* ω *be* $a(\omega) + bq$*, for all* $q \geq 0$ *where for all* $\omega \in \Omega$, $a(\omega) > A(\omega) > 0 > b$ *and* $B > b$*. Following Weitzman, I assume that* $a(\cdot)$ *and* $A(\cdot)$ *are statistically independent and allow marginal cost and marginal benefit (here the price) to be negative. Aggregate surplus in state* ω *is*

$$(a(\omega) - A(\omega))\, q + \frac{1}{2}\,(b - B)\, q^2. \tag{10}$$

Consider two polices: fix a price and require the firm(s) to sell whatever is demanded at that price; or fix a quantity which firms must produce and sell, with the price adjusting so that the market clears. The policy is chosen before the state of the world is realized.

The fixed quantity which maximizes expected surplus is $q^* = E[a - A]/(B - b)$ *and the maximized expected surplus is*

$$W_Q = \frac{1}{2}\frac{(E[a] - E[A])^2}{B - b}.$$

In Appendix B (Section 8.2) I show that the maximized expected surplus for a fixed price is

$$W_P = -\frac{b + B}{2b^2} Var(a) + W_Q,$$

where $Var(a)$ *is the variance of the inverse demand intercept,* $a(\cdot)$*.* W_P *exceeds* W_Q *precisely when* $B < -b$*, the slope of marginal cost is less than minus the slope of the inverse demand (the* ***opposite*** *conclusion from Weitzman's): in particular,* ***if cost is not too convex, expected surplus is higher with a fixed price than a fixed quantity****.*

It turns out that the ***expected*** *price at* $q = q^*$ *equals the expected-surplus-maximizing fixed price. The conclusion extends to the case of an* ***arbitrary*** *fixed quantity and fixed price that is set equal to the expected price under the fixed quantity: a fixed price gives higher expected surplus if and only if the slope of marginal cost is less than minus the slope of the inverse demand. The extension is important for applications in which the fixed quantity or price does not maximize expected surplus; for example it could be chosen to maximize profit or the result of an equilibrium which does not maximize expected surplus.* ■

To illustrate why I get the opposite result from Weitzman's, suppose that the inverse demand is steeply sloped, but that marginal cost is nearly flat (and the intercepts for both are random). If firms are price takers, fixing a price will result in a highly variable quantity, though the surplus-maximizing quantity is not very uncertain; so expected surplus is higher with a fixed quantity. If instead firms must meet all demand at a fixed price, then fixing the price near the mean marginal cost implies that the outcome is close to maximizing surplus in every state of the world; so expected surplus is higher with a fixed price.

The example only compares the extremes of a fixed price and a fixed quantity. Yet as the applications in Section 4 illustrate, many I.O. papers include these extremes. And since I.O. applications also often assume concave or linear costs, Theorem 2 and Example 1 broadly suggest that, when demand is uncertain, polices which maximize expected surplus often result in prices which are the least variable among feasible policies.

Although I introduce Theorem 2 to help explain when the assumptions of Theorem 1 are met, it is of independent interest. Scholars often want to know how policies affect consumer welfare, apart from the affect on profit; Theorems 1 and 2 are useful in answering that question. As I mention in Section 4.1, Theorem 2 generalizes Theorem 4 in Deneckere, Marvel, and Peck (1997) to a larger class of models.

4 Applications

I consider several applications, all from Industrial Organization. Except for one they illustrate that the assumptions of Theorem 1 are more likely to be met if the uncertainty is about consumers, rather than producers.

4.1 The Logic of Vertical Restraints

Rey and Tirole (1986) emphasize uncertainty and information in their account of how vertical restraints on retailers affect welfare. A good is produced by a monopolist manufacturer with a constant unit cost, and sold by two identical retailers with a constant unit cost of retailing; the final demand is linear in the retail price. The demand intercept and the unit *retailing* cost are each possibly random. They consider three policies: competition; resale price maintenance (RPM); and exclusive territories. In each the monopolist sets a franchise fee for each retailer and a wholesale price. Under *competition* the retailers play a Bertrand game; under *resale price maintenance*, the monopolist sets a state-independent retail price, a pure fixed-price policy; under *exclusive territories*, each retailer is a monopolist selling to exactly half the consumers. The monopolist imposes the policy terms *before* any uncertainty is resolved; the retailers order the quantity and (if allowed) set the price *after* the uncertainty is resolved. They allow any degree of retailer risk aversion—from risk neutrality to infinite risk aversion—but assume that consumers have quasilinear utility, and so are risk neutral.

Being fixed, the equilibrium price under resale price maintenance is less variable than under the other policies, no matter what the source of uncertainty. Under pure demand uncertainty, Rey and Tirole (1986) show that RPM has the highest aggregate expected surplus, consistent with the conclusion of my Theorem 2. So policy restrictions of my Theorem 1 are met with pure demand uncertainty. But with pure cost uncertainty, both price variability and equilibrium expected surplus are strictly higher for competition than the other two policies.[15] With pure demand uncertainty, their conclusions extend to economies in which consumers have indirect utilities which satisfy increasing differences and Assumption 2 holds; but with pure cost uncertainty they do not.

4.2 *Minimum* Resale Price Maintenance

Deneckere, Marvel, and Peck (1997) propose a theory of ruinous price competition under

[15] See their Proposition 2, the proofs of Propositions 3 and 7, and Tables 1 and 2. Their Proposition 7 imposes both demand and cost uncertainty, but the conclusions I list can be extracted from their proof, or derived independently.

demand uncertainty.[16] The model is a variant of Rey and Tirole's (1986). A good is produced by a monopolist manufacturer with a constant unit cost and sold by a continuum of risk-neutral retailers with zero cost. One difference with Rey and Tirole is timing: retailers must order inventories *before* the demand uncertainty is resolved; any unsold inventories are worthless. They consider two polices. The first is *flexible pricing*: the manufacturer sets a wholesale per-unit price that each retailer pays at the time that inventories are ordered. After the uncertainty is resolved, the retail price adjusts so that demand equals aggregate inventory. The second is ***minimum*** *resale price maintenance* (RPM): the manufacturer sets both a wholesale price and a minimum price, $p_{\min}$, below which the retail price cannot fall. If demand equals supply at a price above $p_{\min}$, then the market-clearing price prevails; otherwise the retail price is $p_{\min}$ and consumers are rationed among retailers to equalize the *ex ante* chance of a sale. Since the good is produced before demand is realized, the flexible price policy is a pure fixed-quantity policy; unlike Rey and Tirole's (1986) RPM, their *minimum* RPM is not a pure fixed-price policy, but merely moves towards that extreme.

They find that aggregate expected surplus can be higher under minimum RPM than under flexible pricing (Theorem 3 on page 632, and Theorems A1 and A2 in the Appendix)—hence 'ruinous' price competition. RPM can even raise expected consumers' surplus (Theorem 4). Intuitively, minimum RPM leads retailers to order more inventory (Theorem 2); the increased output means a lower retail price when demand is high. The trade-off for consumers is between higher prices when demand is low and lower prices when demand is high. Prices are more variable under flexible pricing than under minimum RPM. We reproduce their Figure 1 (see Figure 3), which gives the supply for flexible pricing and minimum RPM: the intersection of the realized demand with supply determines the retail price. If demand is low, the flexible price is lower and if demand is high, it is higher than the price from minimum RPM: the random price from RPM crosses the flexible price once from above. By Theorem 1, if the No Income Effects Assumption holds and each consumer's indirect utility satisfies increasing differences, the conclusion of Deneckere, Marvel, and Peck (1997)'s Theorems 3 and 4 continue to hold: minimum RPM is (*interim*) Pareto Optimal.

Deneckere, Marvel, and Peck (1997)'s Theorem 4 shows that if expected price under minimum RPM is not larger than the expected price under flexible pricing, then expected consumers' surplus is higher under minimum RPM. My Theorem 2 extends their result to a much larger set of economies and clarifies the source of the conclusion: that prices satisfy the single-crossing property and that equilibrium demand and price are comonotonic.

Minimum RPM and competition; connection to Cho and Meyn (2010)

Deneckere, Marvel, and Peck (1997) briefly consider the case in which the manufacturing market is competitive, rather than a monopolist (p. 634). They point out that a competitive equilibrium (flexible prices) maximizes aggregate expected surplus, and so is Pareto Optimal if consumers have quasilinear utility. For concreteness suppose that the unit cost in manufacturing equals a constant k. Since equilibrium expected profit is always 0 under competition, it seems my Theorem 2 should imply that there is a sometimes-binding price floor with higher aggregate expected surplus than a floor of 0. Theorem 2, however, does not apply in this case, since the competitive version of the model violates the assumption that mean price is lowest for the least-variable price: a sometimes-binding price floor implies that the expected price exceeds k. (Firms are rationed when the price floor binds, so if expected marginal *revenue* equals

[16]See also Deneckere, Marvel, and Peck (1996).

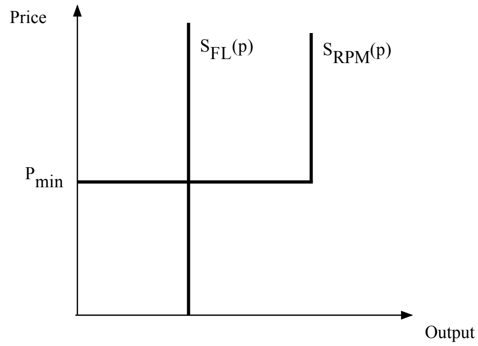

Figure 3: Figure 1 from Deneckere, Marvel, and Peck (1997): The intersection of the realized demand with the output function determines the price; the output function for flexible pricing is $S_{FL}(\cdot)$ and for minimum RPM it is $S_{RPM}(\cdot)$.

k, then expected *price* exceeds k, since some output goes unsold when the floor binds.) Indeed, one can show that if consumers are sufficiently risk averse, rather than risk neutral, then the competitive equilibrium of their model may not be Pareto Optimal: a sometimes-binding price floor can make some consumer types better off, while hurting no type of consumer or firms, the last since expected profit is always 0.

During the crisis in the California electricity market in the late 1990's retail electricity prices fluctuated widely, and conventional wisdom held that the outcome was not optimal (Borenstein (2002)). Motivated in part by this episode Cho and Meyn (2010) prove the striking result that a competitive equilibrium can involve fluctuating prices that almost always diverge from marginal cost, and yet maximize aggregate expected surplus (so is Pareto Optimal if consumers have quasilinear utility). They prove this in the context of a dynamic competitive model with costly capacity adjustment and demand uncertainty. If their model is specialized to a static model with an infinite cost of capacity adjustment, then it collapses to a special case of the flexible-price model just described. As with that model, the competitive equilibrium maximizes aggregate expected consumers' surplus, and so is Pareto Optimal if consumers have quasilinear utility. Also as with that model, it is possible to construct an economy with risk averse consumers which generates the same aggregate demand and equilibrium as the quasilinear economy, but with a competitive equilibrium that is not Pareto Optimal. (I construct such an example in Appendix B, Section 7.3 for the interested reader.) Again the sufficient conditions of Theorem 1 fail, and one reason why is that the sufficient conditions of Theorem 2 fail.

4.3 Regulating a Monopolist with Private Information

Consider a regulator and a firm who produces a single good. The firm has private information about either its cost or its demand. The regulator's objective is a weighted sum of expected consumers' surplus and expected profit, with more weight on consumers's surplus.

4.3.1 Cost Uncertainty

Baron and Myerson (1982) consider a firm whose technology exhibits nondecreasing returns to scale with constant marginal cost. The cost function is private information to the firm, but both the firm and regulator know the demand. The firm reports a cost type to the regulator and the regulator offers a menu of type-contingent prices, quantities and transfers between consumers and the firm to maximize aggregate expected surplus subject to incentive compatibility and participation constraints for the firm. As Baron and Myerson point out, the expected-surplus-maximizing policy may involve setting prices for some cost types which are higher than the highest possible monopoly price in the absence of regulation. Now compare this policy with an unregulated monopoly (so transfer payments are zero). The sufficient condition of Theorem 1 fails here. Indeed, since the worst outcome for consumers can occur under the policy which maximizes aggregate expected surplus, it can happen that all consumers and the firm can prefer no regulation at all to the policy which maximizes aggregate expected surplus.[17]

4.3.2 Demand Uncertainty

Lewis and Sappington (1988) consider the case in which both the firm and regulator know the cost function, but the demand function is private information to the firm. They show that, if the cost function is strictly concave, then the policy which maximizes aggregate surplus (subject to incentive and participation constraints) is to set a price which is *constant* in the demand report of the firm (Proposition 2), a pure fixed price policy. If the no income effects assumption (for the uncertainty case) holds, and each consumer's indirect utility function satisfies increasing differences in the price and income, then Theorem 1 implies that this policy is interim Pareto Optimal.[18]

4.4 Information Acquisition

Suppose firms can acquire information about either demand or cost. Does maximizing aggregate expected surplus result in Pareto Optimal information acquisition? Vives (1999), Section 8.3.3 considers a constant-returns monopolist which is uncertain about the intercept of its linear demand curve. Before it chooses either price or quantity, it can choose the informativeness (precision) of a signal correlated with a demand parameter, with more informative signals costing more.[19] Vives (1999) shows that if the monopolist is a quantity setter (letting price be determined by the realized demand) expected consumer surplus rises with a more informative signal, so aggregate expected surplus rises as well. Since the firm produces more after a good signal, and less after a bad, than it would with less information, prices become less variable with better information. By Theorem 1, choosing a more to a less informative signal

[17] Consider zero fixed-cost case with demand such that an unregulated firm always produces. Suppose regulated firm shuts down for some cost types. In case case the worst outcome for consumers occurs under regulation, since the transfer from consumers is zero (equation (28) on page 920) when the firm shuts down. The same phenomenon can also occur if the firm does not shut down: $\bar{p}(\theta) > p_M(\theta_1) \geq p_M(\theta)$ for some θ, but the firm stays in business.

[18] If the cost function is strictly convex, the optimal regulatory policy is set price equal to marginal cost where demand crosses it. That policy is first-best for every state of the world, and is *interim* Pareto Optimal in a representative consumer economy even without the no-income effects assumption. Absent a representative consumer, the policy might not be interim Pareto Optimal.

[19] Vives (1999) uses the monopoly case to introduce the larger issue of information acquisition and sharing in oligopoly, a question going back to at least to Novshek and Sonnenschein (1982). It's clear that if aggregate expected surplus cannot lead to optimal policies for a monopoly it will not in general do so for an oligopoly.

is interim Pareto Optimal if consumers have indirect utilities satisfying increasing differences. But he shows that if the monopolist is a price setter, both aggregate expected and expected consumers' surplus fall with better information, so a planner interested in expected surplus, but who cannot control the monopolist's pricing decision would prefer zero information to be acquired. Since information acquisition makes prices more variable, this conclusion again fits Theorem 1: every consumer with a increasing-differences indirect utility agrees that no information is the best policy. What if the uncertainty is about cost instead of demand? Schlee (2008), Section 4.1, shows that if uncertainty is about cost, expected consumers' surplus rises with better information for a price (or quantity) setter, but if consumers are sufficiently risk averse, all consumers prefer the *less* informative signal. If information and costly and the firm is just indifferent between acquiring it and not, then the Pareto optimal decision would be not to acquire it, even though aggregate expected surplus is higher with information.

If the planner *can* control output, it now always prefers more information to less. Persico (1996), Section 6.1, considers information about unknown demand and Athey and Levin (2001), Section 4.1, consider information about unknown cost. They each find that under mild conditions the firm acquires less information than would a planner whose objective is aggregate expected surplus. Let the inverse demand function be $P(q, \omega)$ and cost function be $c(q, \omega)$. A planner would set output so that expected price equals expected marginal cost: $E[P(q, \omega)] = E[c'(q, \omega)]$. If only cost is unknown, it is clear that information can make consumers with increasing-differences indirect utilities worse off: with known demand, price varies only when output varies, and information makes output vary more; indeed if information is costly, it can happen that no one prefers and some consumers are strictly worse off with better information, even though aggregate expected surplus rises. Suppose instead that cost is known but demand is uncertain and increases in the state increase the demand. Price now becomes *less* variable with better information since the planner adjusts output in the same direction as demand after a signal is observed (provided the mild stability condition $p' - c'' < 0$ holds). Summing up, suppose for concreteness that the planner chooses between no information and some information and suppose that the planner chooses output to maximize expected surplus (so the planner always prefers some information to none if it is costless). Once again, when consumers have increasing-differences indirect utilities: *If demand is uncertain, then the planner's information choice is interim Pareto Optimal; if cost is uncertain it might not be.*

4.5 Price Caps with Demand Uncertainty

Earle, Schmedders, and Tatur (2007) consider an oligopoly model with identical firms producing under constant returns and uncertain demand. There is a legal maximum price that firms may charge, the *price cap*. The timing is the same as Deneckere, Marvel, and Peck (1997): Firms choose production before demand is realized; after it is realized price adjusts to the lower of two numbers, the market-clearing price and the price cap. They give sufficient conditions for an increase in the price cap to raise equilibrium output and aggregate expected surplus. Their model seems to fit the framework of Theorems 1 and 2. But with price caps and production in advance of the demand realization, consumers are rationed, putting it outside the scope of my model. Indeed, raising the cap price makes price even higher when demand is high, but even lower when demand is low: raising a price cap makes the random price *more* variable, not less, yet increases expected surplus. Nonetheless it is easy to see that their main conclusion—that raising price caps can raise welfare when demand is uncertain—extends beyond quasilinear utility. If the price cap is set at the unit cost of production, then firms would produce the competitive output for the *lowest* demand realization. If this output is zero, then the price cap causes firms to shut down. If the cap is removed altogether, and firms produce positive output in some states of the world, then no consumer is hurt by removing the price cap, and some of them are strictly better off (as are the firms). If we interpret shut-down as a price so high that demand is zero in each state, raising the price cap in this example results in a *first*-order stochastic dominance shift in the price distribution.

5 *Ex ante* Optimality

As mentioned Rogerson (1980) considers *ex ante* welfare, while I consider *interim* welfare. The usual argument is that which is more suitable depends on what people know when the policy is chosen: if everyone knows their own attributes (but maybe not those of others), then *interim* welfare is more suitable; if no one knows their attributes then *ex ante* is. For example, in the optimal regulation and nonlinear pricing literatures, demand is uncertain because of private preference information, so *interim* optimality seems best there. If demand is uncertain because of business cycles or weather, then *ex ante* optimality seems best.[20]

Rogerson (1980) proves that expected surplus is *ex ante* Pareto Consistent—*ex ante* expected surplus rises whenever all consumers prefer the change—if and only if the indirect utility function of each consumer is additively separable in income any variables that are random (Theorem 1). He proves the "only if" part by considering *all* distributions over the variables which are random. His Theorems 3 and 4 then suggest that expected surplus is more likely to be a good welfare measure if uncertainty concerns producers rather than consumers. In this section I do not aim for full generality; I simply show that by restricting the set of policies and preferences, expected surplus maximizing policies can be *ex ante* Pareto Optimal. I consider two polar cases, a common values model in which consumer demands are comonotonic, and an independent private values model. The common values model might be apt when demand is unknown because of business cycle uncertainty or uncertainty about the quality of a newly introduced product; the private values model might be apt when the uncertainty is not about a good's quality, but how well the good matches each consumer's tastes.

In this section I allow for both income and preference uncertainty but each result includes as important special cases pure preference or pure income uncertainty. Consumer i's *ex ante* preference over policies are represented by

$$\alpha \mapsto \int V(\widetilde{p}(\alpha,\omega), \tau^i(\omega))dF(\omega),$$

and i's *ex ante* expected surplus is $\int_\Omega \left[\int_{\widetilde{p}(\alpha,\omega)}^{\bar{p}} D(p, m_i, \tau^i(\omega))dp + \tau_1^i(\omega)\right] dF(\omega)$, where τ_1^i is the income component of the consumer's type. A policy $\alpha^* \in A$ is *ex ante Pareto Optimal* if there is no $\alpha' \in A$ such that every consumer $i \in \{1, ..., I\}$ *ex ante* weakly prefers α^* to α', and the inequality is strict for at least one i.

In both results which follow I impose all the assumptions of Theorem 1 on *interim* Pareto Optimality. Recall from the discussion of Lemma 1 that a crucial step in the argument for *interim* optimality is to show that the marginal utility of money is higher in states of the world with high prices than it is in states with low prices. From an *interim* perspective, after types are revealed, the property follows from the single crossing property and that V_m is increasing in p for each type. From an *ex ante* perspective, a consumer is uncertain of its type as well as

[20]The two welfare concepts are discussed in Mas-colell, Whinston, and Green (1995), pp. 898-99.

the price. The additional assumptions I impose ensure that the marginal utility of money is still higher in states of the world in which price is high.

In the next Theorem I assume that demand uncertainty is common values: consumer types are comonotonic (implying that they are positively correlated), and demand is increasing in the types (Assumptions (*ii*) and (*iii*)). I also impose the assumption from Theorem 2 that equilibrium price is positively correlated with aggregate demand (Assumption (*i*)). The three assumptions together imply that each consumer's type and the equilibrium price are comonotonic. Finally I assume that the indirect utility function satisfies increasing differences in income and both the price and the type. Together these assumptions ensure that each consumer's marginal utility of money is higher in states of the world with high prices.

Proposition 3 (Ex ante Optimality: Common Values)**.** *Let all the hypotheses of Theorem 1 hold. In addition suppose that*

(i) *aggregate demand* $Q(\cdot,\omega)$ *is strictly decreasing in price on* $[\underline{p},\overline{p})$ *and equilibrium aggregate demand and equilibrium price are comonotonic;*

(ii) *consumer preference types are pairwise comonotonic: for every* i, j *in* $\{1,...,I\}$*, and every* ω,ω' *in* Ω $[\tau^i(\omega)-\tau^i(\omega')]\cdot[\tau^j(\omega)-\tau^j(\omega')]\geq 0$*;*

and for each consumer $i\in\{1,...,I\}$*,*

(iii) *at each price* $p\in P$*, demand* $D(p,t)$ *is increasing in the type* t*; and*

(iv) $V_m(p,t)$ *is increasing in* t *on* $Range(\tau^i(\cdot))$ *for every* $p\in P$*.*

Then the expected surplus maximizing policy α^* *on* $\mathcal{A}$ *is* ***ex ante*** *Pareto Optimal.*

Proof: Appendix A.

Assumption (*iv*)—the marginal utility of money is increasing in the type—deserves comment. I discuss it more fully after Proposition 4, but for now note that Proposition 3 includes the case in which income is *ex ante* uncertain (when the range of τ_1, the income component of the type, is not a singleton). But then *(iv)* implies that V is *convex* in income—which is incompatible with strict risk aversion. Proposition 3 just gives sufficient conditions for the conclusion. Yet it is hard to imagine how *(iv)* could be weakened in an economy in which good 1 is normal, consumer incomes are positively correlated, and demand and price are positively correlated. As mentioned a major example of *ex ante* income uncertainty as a source of demand uncertainty with common values is business cycle risk. One implication of Proposition 3 is negative: *if the main source of demand uncertainty is business cycle fluctuations, and policies change less often than the fluctuations, then aggregate expected surplus has little justification in an economy with risk averse consumers.* This implication agrees with the spirit of Rogerson's Theorem 3, which deals with pure income uncertainty.[21]

Rogerson's Theorem 4 deals with pure preference uncertainty (his case of $\ell = 1$ fits my partial equilibrium framework). He shows that if the marginal utility of money does not depend on the preference type and the demand for good 1 is strictly increasing in the preference type, then the marginal utility of money must depend on the price of good 1—violating the sufficient condition of his Theorem 1. My Proposition 3 includes the case of pure preference uncertainty.

[21] Rogerson's Theorem 3 shows that $V_{pm} = 0$ and $V_{mm} = 0$—his sufficient and necessary conditions for *ex ante* expected consumer's surplus to represent *ex ante* preferences over *all* policies which affect the price of good 1 and income—implies that $D_m = 0$ (so if $V_{pm} = 0 = V_{mm}$ for all consumers, then income changes can't be a source of demand changes).

It shows that aggregate expected surplus can be justified as an *ex ante* guide to policies in the case of pure preference uncertainty if the Assumptions of Theorems 1 and 2 hold, demand uncertainty is common values, and V_m is increasing in the preference type.

I now consider the polar opposite case, independent private values. I must now also restrict how the shape of the demand varies with the type, and I assume that each consumer is "informationally small" in the sense that each consumer's beliefs about the equilibrium price distribution does not depend on that consumer's type realization (an assumption that is inconsistent with a finite number of consumers, just as the price-taking assumption is).

Proposition 4 (Ex Ante Optimality: Independent Private Values)**.** *Let all the hypotheses of Theorem 1 hold, suppose that α^* maximizes aggregate expected surplus on $\mathcal{A}$, and suppose that*

(i) Consumer types are independently distributed; and each consumer i's belief about the equilibrium price distribution does not depend on i's type realization, $\tau^i(\omega)$;

(ii) $D(p,\cdot)$ is increasing in t on T for every $p \in P$ and $V(p,\cdot,\theta)$ is concave in m on $\mathbb{R}_{++}$ for every $(p,\theta) \in P \times \Theta$.

*(iii) At least **one** of the following holds for every $i \leq I_C$ and $\alpha \in \mathcal{A}$.*

(iii-a) $V_m(p,\cdot)$ is increasing in t on Range$(\tau^i(\cdot))$; the demand D satisfies increasing differences in p and t on $P \times$ Range$(\tau^i(\cdot))$; and $\int_\Omega \widetilde{p}(\alpha,\omega)dF(\omega) \geq \int_\Omega \widetilde{p}(\alpha^*,\omega)dF(\omega)$;

(iii-b) $V_m(p,\cdot)$ is increasing in $-t$ on Range$(\tau^i(\cdot))$; the demand D satisfies increasing differences in p and $-t$ on $P \times$ Range$(\tau^i(\cdot))$; and $\int_\Omega \widetilde{p}(\alpha,\omega)dF(\omega) \leq \int_\Omega \widetilde{p}(\alpha^*,\omega)dF(\omega)$; or

(iii-c) Demand is additively separable in p and t and $\int_\Omega \widetilde{p}(\alpha,\omega)dF(\omega) = \int_\Omega \widetilde{p}(\alpha^*,\omega)dF(\omega)$.

*Then α^* is **ex ante** Pareto Optimal.*

Proof: Appendix A.

Except for the important special case of *(iii-c)*, both Propositions 3 and 4 assume monotonicity of V_m in the preference type θ (increasing or decreasing). To pin down each possibility from preferences, suppose that the indirect utility function is C^2 and that the demand is C^1. Write the direct utility function as $u(x,y,\theta)$, where x is consumption of good 1 and y is spending on all other goods, and assume that it is strongly quasiconcave. Routine calculation reveals that (subscripts denote partial derivative)

$$V_{\theta m} = D_m D_\theta H + u_{y\theta} \tag{11}$$

where $H = -u_{xx} + 2u_{xy}p - u_{yy}p^2 > 0$ and p is the price of good 1. So if demand is increasing in both income and the preference type and $u_{y\theta} \geq 0$, then V_m is increasing in θ. These conditions are satisfied, for example, if u is additively separable and strongly concave, and θ enters only the good-1 component function—$u(x,y,\theta) = f(x,\theta) + g(y)$. On the other hand, if demand does not depend on income (quasilinear preferences) but the consumer is risk averse over money lotteries, then V_m is decreasing in θ (since $u_{y\theta} \leq 0$ in that case).

The case of high consumer risk aversion is of special interest. In particular suppose that the indirect utility is of the form $V(p,m,\theta) = \phi(\widetilde{V}(p,m,\theta))$, where ϕ is strictly increasing and strictly concave with $\phi' > 0$. Differentiate with respect to m and θ and rearrange to find that

$$V_{m\theta} = \phi' \left[\frac{\phi''}{\phi'} \widetilde{V}_m \widetilde{V}_\theta - \widetilde{V}_{m\theta} \right].$$

As with equation (8), the equation immediately implies that if Arrow-Pratt risk aversion is uniformly high enough and $V_\theta < 0$, then V_m is *increasing* in θ; if risk aversion is high enough and $V_\theta > 0$, then V_m is *decreasing* in θ. Since demand is increasing in the preference type, the conclusion can be rephrased: if the consumer is risk averse enough, then V_m is increasing in θ if demand and utility move in *opposite* directions in response to changes in θ; and V_m is decreasing in θ if demand and utility move in *same* direction in response to changes in θ.

Examples abound for each possibility. Suppose the uncertain state of the world affects the weather. Demand for heating and cooling is highest when the weather is extreme; and since most prefer moderate to extreme weather, a consumer's utility and demand for climate control and move in opposite directions. But demand for many outdoor recreation activities is highest when weather is moderate. If the uncertainty is about the quality of a *substitute* for the good, then utility and the demand for the good move in opposition directions—for example car gas mileage and the demand for gasoline; and possibly health care quality and the demand for hospital days (Feldstein (1977)). If the uncertainty is about the quality of the good itself—with 'higher' quality defined as something everyone prefers—then utility is by construction increasing in quality, but the demand for the good could be either increasing or decreasing in quality: if the durability of a firm's product increases, then its demand might increase if it has competitors, but might decrease if the firm is a (protected) monopolist.

Propositions 3 and 4 contain a long list of assumptions, so it might not be obvious how to apply them. Here I only consider pure preference uncertainty. Suppose first that the demand uncertainty is common values, as in Proposition 3. If one adds the assumption that V_m is increasing in the preference type θ, then the results in Lewis and Sappington (1988), Deneckere, Marvel, and Peck (1997), and Rey and Tirole (1986) that were shown to be *interim* Pareto Optimal in Section 4 are also *ex ante* optimal for common-values demand uncertainty.

For private values, notice that condition *(iii-c)* includes the familiar case of linear demand with intercept uncertainty: $D(p, m, \theta) = g(\theta, m) - p$. Rey and Tirole (1986) impose this form in their main theorem on vertical restraints and welfare (Proposition 7). If retailers are risk neutral or infinitely risk averse, then the expected retail price is the same across the three policies for pure demand uncertainty and condition *(iii-c)* holds, so by Proposition 4, RPM is *ex ante* Pareto Optimal for private-values demand uncertainty. If retailers have intermediate risk aversion, then the expected retail price for exclusive territories can be higher than for RPM and competition, so that conclusion extends if we add the assumption that V_m is increasing in θ. Deneckere, Marvel, and Peck's (1997) Theorem A.2 imposes linear demand. They prove that the mean price is the same for flexible pricing and minimum RPM, so again condition *(iii-c)* holds. Deneckere, Marvel, and Peck (1997) also consider nonlinear demand. Their maintained assumptions on the curvature on the inverse demand are in general consistent with D satisfying increasing differences in either p and θ or p and *minus* θ but in some important special cases, the second must hold. They also point out that expected price under minimum RPM can be either higher or lower than under flexible prices, but for important classes of demands and costs, the mean price for flexible prices is no less than the mean price under minimum RPM. Thus assumption (*iii-b*) of Proposition 4 fits best for nonlinear demand when the demand uncertainty is iid in their model.

Finally consider Lewis and Sappington (1988) and private values. Although linear demand is their leading example of a functional form which fits their assumptions, Proposition 4 requires that the mean price from the policy which maximizes aggregate expected surplus be higher or lower than the mean price from every other policy. That part of the sufficient condition fails in Lewis and Sappington (1988): the best hope for extending their results beyond quasilinear utility is for *interim* optimality, or common-values uncertainty for *ex ante* optimality.

6 What is to be done?

I extend important welfare conclusions under uncertainty and incomplete markets substantially beyond the knife edge case of quasilinear utility often assumed in applied microeconomics. For *interim* optimality I argue the conditions are plausible if uncertainty is about consumers rather than producers; for *ex ante* optimality they are less so, but still applicable to some welfare conclusions. I do not intend the paper to be an apology for misspecified welfare functions. Rather I encourage applied researchers to examine the robustness of their welfare conclusions, especially to consumer risk aversion; and encourage theorists to prove theorems about welfare from limited information, in the spirit of Willig (1976), Radner (1993), and Vives (1987). To illustrate, my model is a partial equilibrium one. I imposed this restriction mainly to replicate the positive predictions of the partial equilibrium applications I consider in Section 4. Still, a natural and important extension is to determine when an optimal partial-equilibrium policy remains optimal in general equilibrium. The existing literature on 'piecemeal' second-best policy—for example, Jewitt (1981)—is not directly applicable because it considers second-best economies in which some relative prices are fixed exogenously; but in many applied microeconomics models, markets under uncertainty are incomplete. An important open question is: if markets are incomplete—in particular only spot markets exist, as here—then when is piecemeal second-best policy Pareto Optimal?

7 Appendix A: Proofs

The proofs appear in the same order as the results in the text.

7.1 Proof of Proposition 1

Let $\mathbf{y}(\alpha) \in \mathbb{R}^L$ be the equilibrium production plan, $\mathbf{y}'(\alpha)$ its derivative, $\mathbf{e}^g$ the aggregate endowment of goods, and $\widetilde{\mathbf{p}}(\alpha) \in \mathbb{R}^L_{++}$ the equilibrium price vector at policy α. (I follow the usual sign convention for production plans that positive coordinates represent net outputs and negative ones net inputs.) Under my partial equilibrium assumption α only affects the first coordinate of $\widetilde{\mathbf{p}}(\alpha)$. With the equilibrium production plan, equilibrium prices, and consumer utility all differentiable, and the equilibrium allocation interior, the economy fits the framework of Radner (1993). His main result is if $\widetilde{\mathbf{p}}(\widehat{\alpha}) \cdot \mathbf{y}'(\widehat{\alpha}) > 0$—the value of the change in the equilibrium production plan from a small increase in α from $\widehat{\alpha}$ is positive—then a small-enough increase in α is a potential Pareto improvement.[22] I will show that $S'(\alpha) = \widetilde{\mathbf{p}}(\alpha) \cdot \mathbf{y}'(\alpha)$, from which the result follows. Equilibrium aggregate income is $\sum_i m_i(\alpha) = \widetilde{\mathbf{p}}(\alpha) \cdot (\mathbf{e}^g + \mathbf{y}(\alpha))$, the value of the aggregate endowment plus profit. Differentiate with respect to α to find that

$$\sum_i m_i'(\alpha) = \widetilde{\mathbf{p}}'(\alpha) \cdot (\mathbf{e}^g + \mathbf{y}(\alpha)) + \widetilde{\mathbf{p}}(\alpha) \cdot \mathbf{y}'(\alpha), \tag{12}$$

[22] Intuitively, the allocation of the total supply $\mathbf{e}^g + \mathbf{y}(\widehat{\alpha})$ to consumers is a competitive equilibrium allocation with respect to the price vector $\widetilde{\mathbf{p}}(\widehat{\alpha})$. Since the allocation is interior, the first order condition implies that the price vector equals the vector of marginal utilities divided by the marginal utility of money for each consumer. If each consumer's share of the new supply is proportional to the *reciprocal* of that consumer's marginal utility of money, then each consumer's utility rises with a small increase in α starting from $\widehat{\alpha}$.

and differentiate equation (1) with respect to α and sum over consumers to find (recall that $\widetilde{p}(\alpha)$ is the equilibrium price of just good 1)

$$S'(\alpha) = -\widetilde{p}'(\alpha) \sum_i D(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i) + \sum_i \widetilde{m}_i'(\alpha). \tag{13}$$

Insert (12) into (13), and use market clearing and the assumption that the policy only affects the price of good 1 to conclude that $S'(\alpha) = \widetilde{\boldsymbol{p}}(\alpha) \cdot \mathbf{y}'(\alpha)$. If $S'(\widehat{\alpha}) > 0$, and the conclusion now follows from in Radner (1993). ■

7.2 Proof of Lemma 1

Let the hypotheses of Lemma 1 hold. The conclusions hold for any type of consumer $i > I_C$ by assumption, so fix a consumer $i \leq I_C$ of type $t = (m, \theta)$. Let $\Omega_+ = \{\omega \in \Omega \,|\, \widetilde{p}(\alpha, \omega) > \widetilde{p}(\alpha^*, \omega)\}$ and let $\Omega_- = \{\omega \in \Omega \,|\, \widetilde{p}(\alpha, \omega) < \widetilde{p}(\alpha^*, \omega)\}$. If Ω_- has 0 measure, namely $\int_{\Omega_-} dF = 0$, then a type t consumer i weakly prefers α^* to α, since the policy does not affect that consumer's income; and if that type's expected surplus is strictly higher at α^*, then the preference is strict. If Ω_+ has 0 measure then, since expected surplus at α^* is at least equal to expected surplus at α for a type t consumer i—and that consumer's surplus in each state is decreasing in the price of good 1—the price at α must equal the price at α^* F-almost everywhere, so a type-t consumer i is indifferent between α and α^*.

For the rest of the proof suppose that both Ω_+ and Ω_- have positive measure and let r be a crossing price for the two policies (that is, r satisfies (15)). Since $V_m(\cdot, t)$ is increasing in p,

$$V_m(p^+, t) \geq V_m(r, t) \geq V_m(p^-, t) \tag{14}$$

for every $p^+ \in [\widetilde{p}(\alpha^*, \omega_+)\,,\ \widetilde{p}(\alpha, \omega_+)]$ and $p^- \in [\widetilde{p}(\alpha, \omega_-)\,,\ \widetilde{p}(\alpha^*, \omega_-)]$ where $\omega_+ \in \Omega_+$ and $\omega_- \in \Omega_-$. It follows that

$$\begin{aligned}
\int & \Big(V^i(\widetilde{p}(\alpha^*, \omega), t) - V^i(\widetilde{p}(\alpha, \omega), t)\Big) dF(\omega \,|\, i, t) \\
&= \int_\Omega \left[\int_{\widetilde{p}(\alpha,\omega)}^{\widetilde{p}(\alpha^*,\omega)} V_p(p, t) dp\right] dF(\omega \,|\, i, t) \\
&= \int_\Omega \left[\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p, t) V_m(p, t)\ dp\right] dF(\omega \,|\, i, t) \\
&= \int_{\Omega_+} \left[\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p, t) V_m(p, t)\ dp\right] dF(\omega \,|\, i, t) + \int_{\Omega_-} \left[\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p, t) V_m(p, t)\ dp\right] dF(\omega \,|\, i, t) \\
&\geq V_m(r, t) \int_\Omega \left[\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p, t)\ dp\right] dF(\omega \,|\, i, t) \\
&= \ V_m(r, t)\big(E[S(\alpha^*, \omega, t) \,|\, i, t] - E[S(\alpha, \omega, t) \,|\, i, t]\big).
\end{aligned} \tag{15}$$

The second equality uses Roy's Identity, $D = -V_p/V_m$, and the inequality follows from (14). If the difference in expected surplus (12) is nonnegative, so is the difference in expected utility; if the difference in expected surplus is positive, then so is the difference in expected utility. ■

7.3 Proof of Theorem 2

Suppose that the hypotheses of the theorem hold and let $\beta \in \mathcal{A}$ be the policy which is least variable and has the lowest mean: for any $\alpha \in \mathcal{A}$, $\widetilde{p}(\beta, \cdot)$ crosses $\widetilde{p}(\alpha, \cdot)$ at most once from above—equation (6) holds with $\alpha^* = \beta$—and $\int_\Omega \widetilde{p}(\alpha, \omega) d\omega \geq \int_\Omega \widetilde{p}(\beta, \omega) d\omega$. Let $Q(\widetilde{p}(\alpha, \omega), \omega)$ be the equilibrium aggregate demand for good 1 in state ω under policy α. Fix $\alpha' \in \mathcal{A}$, let r be a crossing price between polices β and α', and let $\underline{Q} = \inf_{\omega \in \Omega}\{Q(\widetilde{p}(\alpha', \omega), \omega) \,|\, \widetilde{p}(\alpha', \omega) \geq r\}$. The difference in expected surplus between policy β and some $\alpha' \in A$ is

$$
\begin{aligned}
&\int_\Omega \left[\int_{\widetilde{p}(\beta,\omega)}^{\widetilde{p}(\alpha',\omega)} Q(p, \omega) dp \right] dF(\omega) \\
&\geq \int_\Omega \big(\widetilde{p}(\alpha', \omega) - \widetilde{p}(\beta, \omega)\big)\, Q(\widetilde{p}(\alpha', \omega), \omega) dF(\omega) \\
&\geq \underline{Q} \int_\Omega \big(\widetilde{p}(\alpha', \omega) - \widetilde{p}(\beta, \omega)\big)\, dF(\omega) \geq 0. \qquad (16)
\end{aligned}
$$

The first inequality holds since $Q(p, \omega) \geq Q(\widetilde{p}(\alpha', \omega), \omega)$ for states with $\widetilde{p}(\alpha', \omega) \geq r$ and prices $p \geq \widetilde{p}(\beta, \omega)$; and $Q(p, \omega) \leq Q(\widetilde{p}(\alpha', \omega), \omega)$ for states with $\widetilde{p}(\alpha', \omega) \leq r$ and and prices $p \leq \widetilde{p}(\alpha', \omega)$. The second inequality follows from the single-crossing property and comonotonicity of equilibrium demand and price at α': $Q(\widetilde{p}(\alpha', \omega), \omega) > \underline{Q}$ only if $\widetilde{p}(\alpha', \omega) \geq r$ (hence $\widetilde{p}(\alpha', \omega) \geq \widetilde{p}(\beta, \omega)$). The last inequality follows from the assumption that β has the lowest mean price. ■

7.4 Proof of Proposition 3

First, I show that Assumptions *(i)*-*(iii)* of Proposition 3 imply that for each consumer $i \in \{1, ..., I\}$ and every policy $\alpha \in \mathcal{A}$, the equilibrium price and i's preference type are comonotonic: for all ω, ω' in Ω,

$$[\tau^i(\omega') - \tau^i(\omega)][\widetilde{p}(\alpha, \omega') - \widetilde{p}(\alpha, \omega)] \geq 0. \qquad (17)$$

Fix a policy $\alpha \in \mathcal{A}$, a consumer $j \in \{1, ..., I\}$ and a pair of states ω and ω' in Ω. Suppose that $\widetilde{p}(\alpha, \omega') - \widetilde{p}(\alpha, \omega) > 0$. (If the prices are equal, (17) holds.) Since by *(i)*,

$$\left[Q\left(\widetilde{p}(\alpha, \omega'),\ \omega'\right) - Q\left(\widetilde{p}(\alpha, \omega),\ \omega\right)\right] \times \left[\widetilde{p}(\alpha,\ \omega') - \widetilde{p}(\alpha,\ \omega)\right] \geq 0$$

it follows that $Q\left(\widetilde{p}(\alpha, \omega'),\ \omega'\right) \geq Q\left(\widetilde{p}(\alpha, \omega),\ \omega\right)$. Since aggregate demand is strictly decreasing in price at each state and the price is higher in state ω', it must be that $\tau^j(\omega') > \tau^j(\omega)$ for at least one consumer $j \in \{1, ..., I\}$ (aggregate demand must be higher in state ω' than in state ω at price $\widetilde{p}(\alpha, \omega)$, and each consumer's demand is increasing in θ). But then by *(iii)*, $\tau^i(\omega') \geq \tau^i(\omega)$ for every $i \in \{1, ..., I\}$ and (17) holds.

Now I use (17), Assumption *(iv)*, and the Assumptions imposed in Theorem 1 to prove Proposition 3. Suppose that α^* maximizes expected surplus on $\mathcal{A}$, and consider $\alpha \in \mathcal{A}$. As in the proof of Lemma 1, let $\Omega_+ = \{\omega \in \Omega \,|\, \widetilde{p}(\alpha, \omega) > \widetilde{p}(\alpha^*, \omega)\}$ and Ω_- be the set of states for which the inequality is reversed. If Ω_- has 0 measure, then every consumer $i \leq I_C$ at least weakly prefers α^* to α and the conclusion holds. If Ω_+ has 0 measure, then expected surplus of every consumer $i \leq I_C$ under α^* is less than or equal to its expected surplus under α. Since α^* maximizes aggregate expected surplus, at least one consumer $i > I_C$ weakly prefers α^* to α and the conclusion holds.

For the rest of the proof suppose that Ω_+ and Ω_- have positive measure and let r be a crossing price (that is, r satisfies (15).) Now consider a consumer $i \leq I_C$. Let

$$K = \inf\{V_m(\widetilde{p}(\alpha^*, \omega), m_i, \tau^i(\omega)) \,|\, \omega \in \Omega_+\}.$$

By (17) and the assumption that V_m is increasing in (p, θ), it follows that $V_m(\widetilde{p}(\alpha,\omega), m_i, \tau^i(\omega)) \leq K$ for every $\omega \in \Omega_-$. Use this fact and Roy's Identity to find that

$$\int \Big(V(\widetilde{p}(\alpha^*,\omega), m_i, \tau^i(\omega)) - V(\widetilde{p}(\alpha,\omega), m_i, \tau^i(\omega)) \Big) dF(\omega)$$
$$= \int_\Omega \Big(\int_{\widetilde{p}(\alpha,\omega)}^{\widetilde{p}(\alpha^*\omega)} V_p(p, m_i, \tau^i(\omega)) \Big) dF(\omega)$$
$$= \int_\Omega \Big(\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} V_m(p, m_i, \tau^i(\omega))(D(p, m_i, \tau^i(\omega))) \Big) dF(\omega)$$
$$\geq K \int_\Omega \left[\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p, m_i, \theta)\ dp \right] dF(\omega). \tag{18}$$

The inequality gives a version of Lemma 1 for *ex ante* welfare: if for consumer $i \leq I_C$, α^* has *ex ante* expected surplus no lower than at α, then i weakly prefers α^* to α. Since this statement is also true for any consumer $i > I_C$, it follows that any policy which maximizes *ex ante* aggregate expected surplus is Pareto Optimal. ■

7.5 Proof of Proposition 4

Suppose that *(i)* and *(ii)* hold and consider inequality (15). By *(i)* $F(\omega)$ can replace $F(\omega \,|\, i, t)$ in (15). Multiply by $\lambda_i(t)$—the probability that consumer i is type t—and sum over types to find that

$$\sum_t \int \Big(V^i(\widetilde{p}(\alpha^*,\omega), t) - V^i(\widetilde{p}(\alpha,\omega), t) \Big) dF(\omega) \lambda_i(t)$$
$$\geq \sum_t \left[V_m(r,t) \times \int_\Omega \Big(\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p,t)\ dp \Big) dF(\omega) \right] \lambda_i(t) \tag{19}$$

Now impose Assumption *(iii-a)*, so V_m is increasing in t. The other term in brackets is the change in expected surplus for a type-t consumer. If that term is increasing in t then the conclusion follows since the second line of (19) would be at least as large as

$$\sum_t \Big[V_m(r,t)\lambda_i(t) \Big] \times \sum_t \int_\Omega \Big(\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} D(p,t)\ dp \Big) dF(\omega) \lambda_i(t).$$

which equals

$$\sum_t \Big[V_m(r,t)\lambda_i(t) \Big] \times \sum_t \big(E[S(\alpha^*,\omega,t)] - E[S(\alpha,\omega,t)] \big) \lambda_i(t).$$

To confirm that the change in expected surplus is increasing in t under *(ii)* and *(iii-a)*, let t'' and t' be any points in the range of $\tau^i(\cdot)$ with $t'' > t'$ and consider the difference in expected surplus

$$\Delta = \int_\Omega \Big(\int_{\widetilde{p}(\alpha^*,\omega)}^{\widetilde{p}(\alpha,\omega)} [D(p,t'') - D(p,t')]\ dp \Big) dF(\omega). \tag{20}$$

It suffices to show that $\Delta \geq 0$. If r be a crossing price for α^* and α (that is, r satisfies (15), then for any $p^- \leq r \leq p^+$ we have

$$D(p^-,t'') - D(p^-,t') \leq D(r,t'') - D(r,t') \leq D(p^+,t'') - D(p^+,t'),$$

which implies that

$$\Delta \geq [D(r,t'') - D(r,t')] \times \int_\Omega \widetilde{p}(\alpha,\omega) - \widetilde{p}(\alpha^*,\omega) dF(\omega).$$

By *(ii)* the bracketed term is nonnegative and by *(iii-a)* expected price difference is nonnegative, so $\Delta \geq 0$ and the change in expected surplus is increasing in the preference type. This completes the proof for the case of *(iii-a)*. The proof in the case of *(iii-b)* is similar. For *(iii-c)*, note that if demand is additively separable in p and (m,θ), then Δ in (20) is 0 if the mean prices are the same for policies α and α^*. ■

8 Appendix B: Supplementary Derivations

In this appendix I gather together derivations and proofs omitted in the text.

8.1 Strong Quasiconcavity implies optimality of surplus-maximizing policies in complete markets

In Section 2.2, I asserted that, if each S_i is strongly quasiconcave, then any surplus-maximizing policy is Pareto Optimal. I now prove that assertion. Suppose each S_i is strongly quasiconcave. Then by equation (3) each consumer i's equilibrium indirect utility $W_i(\alpha) = V(\widetilde{p}(\alpha), \widetilde{m}_i(\alpha), \theta_i)$ is also strongly quasiconcave in α on $\mathcal{A}$. Suppose that $\alpha_0 \in \mathcal{A} = [0,1]$ is not Pareto Optimal: there is an $\alpha_1 \in \mathcal{A}$ with $W_i(\alpha_1) \geq W_i(\alpha_0)$ for all i, with the inequality strict for at least one consumer j. I will show that α_0 cannot maximize aggregate surplus on $\mathcal{A}$. There are two cases. First suppose that $\alpha_1 > \alpha_0$. Then by strong quasiconcavity, $W_i'(\alpha_0) \geq 0$ for every consumer i with a strict inequality for consumer j, and so $S'(\alpha_0) > 0$ by (3). Since $\alpha_0 < 1$, it cannot maximize surplus on $\mathcal{A}$. If $\alpha_1 < \alpha_0$, then $W_i'(\alpha_0) \leq 0$ for every i with a strict inequality for consumer j, and so $S'(\alpha_0) < 0$; since $\alpha_0 > 0$, it cannot maximize S on $\mathcal{A}$. ■

8.2 Prices vs Quantities (Example 1)

Here I calculate the maximized expected surplus under a fixed-price policy. If price is fixed at p, then output in state ω is $\frac{a(\omega)-p}{-b}$. Insert this expression into (10) to find that expected surplus at price p is

$$E\Big[(a(\omega) - A(\omega)) \left(\frac{a(\omega)-p}{-b} \right) + \frac{1}{2}(b-B) \left(\frac{a(\omega)-p}{-b} \right)^2 \Big]. \tag{21}$$

Differentiate with respect to p and set equal to 0 to find

$$E[a(\omega)] - E[A(\omega)] - (E[a(\omega)] - p)) \frac{b-B}{b} = 0,$$

so the expected surplus maximizing price is

$$\overline{p} = E[a(\omega)] + \frac{b(E[a(\omega)] - E[A(\omega)])}{B-b}.$$

Insert this expression into (21) and rearrange to find that

$$W_P = -\frac{b+B}{2b^2} Var(a) + W_Q, \tag{22}$$

as asserted in the text.

To demonstrate the assertion of the last paragraph of Example 1, note that for an arbitrary fixed quantity $\overline{q}$, the expected price is $E[a(\omega)] + b\overline{q}$. Substitute this expression for p in (21) and simplify to find that the the difference between expected surplus with a fixed price of $E[a(\omega)] + b\overline{q}$ and a fixed quantity of $\overline{q}$ is $-(b + B)/2b^2$, exactly as in (22).

8.3 Calculations for Cho and Meyn (2010)

I now construct an example which generates the same aggregate demand and supply as the static version of Cho and Meyn (2010) which has a competitive equilibrium that is not Pareto Optimal. Following Cho and Meyn (2010), I suppose that aggregate demand in state $\omega \in \Omega = [0, 1]$ is $v\omega$ if $p \leq \omega$ and is 0, otherwise, where $v > 0$. Before the demand is realized a production capacity $q \geq 0$ is installed. As before the distribution of the state ω is given by a c.d.f. F. There is a continuum of identical firms. Before the demand uncertainty is realized each firm installs a production capacity, which cannot later be changed. The cost of capacity q for any firm is $c(q) = kq$, where $0 < k < v$. After the capacity is installed, the demand state ω is realized. In state ω with capacity q the equilibrium price, $\widetilde{P}(q, \omega)$, is given by the intersection of the aggregate demand with the inelastic supply: $\widetilde{P}(q, \omega) = 0$ if $\omega < q$, $\widetilde{P}(q, \omega) = v$ if $\omega > q$ (and anything in the interval $[0, v]$ if $\omega = q$). The static version of the model with no capacity adjustments is just a special case of the competitive version of Deneckere, Marvel, and Peck (1997) with flexible prices, and corresponds to a pure fixed quantity policy. As with that model, Theorem 2 does not apply, so it remains an open question whether the conditions of Theorem 1 are satisfied.

In a competitive equilibrium the expected price equals k: $v(1 - F(q)) = k$. Expected aggregate surplus for a given q is

$$v \int_0^q \omega dF(\omega) + qv(1 - F(q)) - kq \tag{23}$$

which is maximized when $v(1-F(q)) = k$, so the competitive equilibrium choice of q is Pareto Optimal whenever each consumer has quasilinear utility. The competitive equilibrium quantity equals the fixed quantity that social planner who maximizes expected surplus.

Now let $\mathbb{I} = [0, 1]$ be the set of consumers. Each consumer i has a unit demand for the good and is one of two types: type 1 values the good at v dollars, type 2 values the good at 0 dollars. Let the mapping from consumer labels to types be the indicator function $\chi_{[0,\omega]}$ for $\omega \in [0, 1]$: if consumer i learns that he is of type 1, then he knows that all consumers with labels less than i are type 1: consumer i however is uncertain about the types above i; if $i < q$ then a type-1 consumer i is uncertain about the equilibrium price. If $i > q$, then both potential types of consumer i know that the equilibrium price is v. Consumer i's vN-M utility in state ω is $u(x, y; i, \omega) = \phi(xv\theta_i(\omega) + y)$, where $x \in \{0, 1\}$ is consumption of good 1, y is consumption of good $L = 2$, and $\theta_i(\omega) = 1$ if $i < \omega$ and 0 otherwise. The function ϕ is concave and strictly increasing. Normalize it so that $\phi(m) = 0$. When the capacity is $q > i$ and consumer i's type is 1, then its equilibrium *interim* expected utility is

$$\phi(v + m)\frac{F(q) - F(i)}{1 - F(i)}$$

Now consider the following policy of minimum RPM: a price floor $p_F \geq 0$ is set so that if the market-clearing price is below p_F, then the actual price is held at p_F, and each firm only sells a fraction of its capacity (the same for all firms). The equilibrium output $Q(p_F)$ is

$$p_F \int_0^{Q(p_F)} \omega dF(\omega) + v(1 - F(Q(p_F))) = k.$$

It follows immediately that the equilibrium output is increasing in p_F. Equilibrium *interim* expected utility for a type 1 consumer i in this economy is

$$\phi(v - p_F + m)\frac{F(Q(p_F)) - F(i)}{1 - F(i)}.$$

The tradeoff in raising p_F from 0 is that each consumer pays more the good, but is more likely to get some gains to trade. Differentiate this last expression with respect to p_F to find that

$$-\phi'(v - p_F + m)\frac{F(Q(p_F)) - F(i)}{1 - F(i)} + \phi(v - p_F + m)\frac{f(Q(p_F))}{1 - F(i)}Q'(p_F).$$

Evaluate at $p_F = 0$ to find that

$$-\phi'(v + m)\frac{F(Q(0)) - F(i)}{1 - F(i)} + \phi(v + m)\frac{f(Q(0))}{1 - F(i)}Q'(0).$$

Now normalize ϕ so that $\phi(v + m) = 1$ and $\phi(m) = 0$. As risk aversion increases uniformly without bound $\phi'(v + m)$ tends to zero, so that consumer i preferences are strictly increasing in p_F in a neighborhood of 0: all consumer types weakly prefer a small price floor to none, and a positive mass of type one consumers strictly prefer it.

References

Athey, S., and J. Levin (2001): "The Value of Information in Monotone Decision Problems," working paper,.

Baron, D., and R. Myerson (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–30.

Blackorby, C. (1999): "Partial-Equilibrium Welfare Analysis," *Journal of Public Economic Theory*, 1, 359–74.

Borenstein, B. (2002): "The Trouble With Electricity Markets: Understanding Californias Restructuring Disaster," *Journal of Economic Perspectives*, 16, 191–211.

Chipman, J., and J. C. Moore (1976): "The Scope of Consumer's Surplus Arguments," in *Evolution, Welfare and Time in Economics*, ed. by A. Tang, F. Westfield, and J. Worley, pp. 69–124. D.C. Heath and Company.

——— (1980): "Compensating Variation, Consumer's Surplus, and Welfare," *American Economic Review*, 70, 933–49.

Cho, I.-K., and S. Meyn (2010): "Efficiency and Marginal Cost Pricing in Dynamic Competitive Markets with Friction," *Theoretical Economics*, 5, 215–239.

Cohen, A., and L. Einav (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97, 745–788.

Debreu, G. (1954): "A Classical Tax-Subsidy Problem," *Econometrica*, 22, 14–22.

——— (1959): *The Theory of Value*. Yale University Press, New Haven.

Deneckere, R., H. Marvel, and J. Peck (1996): "Demand Uncertainty, Inventories, and Resale Price Maintenance," *The Quarterly Journal of Economics*, 111, 885–913.

——— (1997): "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition," *American Economic Review*, 50, 619–641.

Duhamel, M. (2006): "The Optimality of Arbitrary Partial-Equilibrium Welfare Analysis," *Journal of Public Economic Theory*, 8, 321–343.

Earle, R., K. Schmedders, and T. Tatur (2007): "On Price Caps Under Uncertainty," *Review of Economic Studies*, 74, 93–111.

Facchini, G., P. J. Hammond, and H. Nakata (2001): "Spurious Deadweight Gains," *Economics Letters*, 72, 33–37.

Feldstein, M. (1977): "Quality Change and the Demand for Hospital Care," *Econometrica*, 45, 1681–1702.

Hayashi, T. (2008): "A note on small income effects," *Journal of Economic Theory*, 139, 360–397.

Jewitt, I. (1981): "Preference Structure and Piecemeal Second Best Policy," *Journal of Public Economics*, 16, 215–231.

Just, R., D. Hueth, and A. Schmitz (2004): *The Welfare Economics of Public Policy*. Edward Elgar, Northampton, MA.

Kihlstrom, R., and J.-J. Laffont (1979): "A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion," *The Journal of Political Economy*, 87, 719–748.

Lewis, T., and D. Sappington (1988): "Regulating a Monopolist with Unknown Demand," *American Economic Review*, 78, 986–98.

Mas-colell, A. (1982): "The Cournotian Foundations of Walrasian Equilibrium Theory: An Exposition of Recent Theory," in *Advances in Economic Theory*, ed. by W. Hildendbrand, pp. 183–224. Cambridge University Press.

Mas-colell, A., M. Whinston, and J. Green (1995): *Microeconomic Theory*. Oxford University Press, New York.

Meyer, D., and J. Meyer (2005): "Relative Risk Aversion: What Do We Know?," *Journal of Risk and Uncertainty*, 33, 243–262.

Novshek, W., and H. Sonnenschein (1982): "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release," *Bell Journal of Economics*, 13, 214–18.

Persico, N. (1996): "Information Acquisition in Affiliated Decision Problems," Northwestern University Discussion Paper N. 1149.

Radner, R. (1993): "A Note on the Theory of Cost-Benefit Analysis in the Small," in *Capital, Investment and Development*, ed. by K. Basu, M. Majumdar, and T. Mitra, pp. 129–141. Blackwell.

Rey, P., and J. Tirole (1986): "The Logic of Vertical Restraints," *American Economic Review*, 76, 921–39.

Rogerson, W. (1980): "Aggregate Expected Consumer Surplus as a Welfare Index with an Application to Price Stabilization," *Econometrica*, 48, 423–36.

Samuelson (1942): "Constancy of the Marginal Utility of Income," in *Studies in Mathematical Economics and Econometrics, in Memory of Henry Schultz*, ed. by O. Lange, pp. 75–91. University of Chicago Press.

Schlee, E. (2008): "Expected Consumer's Surplus as an Approximate Welfare Measure," *Economic Theory*, 34, 127–55.

Silberberg, E. (1972): "Duality and the Many Consumer's Surpluses," *American Economic Review*, 62, 942–952.

Slesnick, D. (1998): "Empirical Approaches to the Measurement of Welfare," *Journal of Economic Literature*, 36, 2108–2165.

Tirole, J. (1988): *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Turnovsky, S., H. Shalit, and A. Schmitz (1980): "Price Instability and Consumer Welfare," *Econometrica*, 48, 135–52.

Vives, X. (1987): "Small income effects: A Marshallian theory of consumer surplus and downward sloping demand," *Review of Economic Studies*, 54, 87–103.

——— (1999): *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, Cambridge, MA.

Waugh, F. V. (1944): "Does the Consumer Benefit from Price Instability?," *The Quarterly Journal of Economics*, 58, 602–614.

Weitzman, M. (1974): "Prices vs. Quantities," *The Review of Economic Studies*, 41, 477–491.

Willig, R. (1976): "Consumer's Surplus Without Apology," *American Economic Review*, 66, 589–597.