NBER WORKING PAPER SERIES

TOWARDS AN EFFICIENT MECHANISM FOR PRESCRIPTION DRUG PROCUREMENT

Kyna Fong
Michael Schwarz

Working Paper 14718
http://www.nber.org/papers/w14718

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2009

We would like to thank Jeremy Bulow, Robin Lee, Preston McAfee, Michael Ostrovsky, Andy Skrzypacz, Joel Sobel and seminar participants at Chicago GSB, Columbia Department of Economics, NBER Summer Institute and Stanford GSB for helpful comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Towards an Efficient Mechanism for Prescription Drug Procurement
Kyna Fong and Michael Schwarz
NBER Working Paper No. 14718
February 2009
JEL No. D44,I11

ABSTRACT

This paper applies ideas from mechanism design to model procurement of prescription drugs. We present a mechanism for government-funded market-driven drug procurement that achieves very close to full static efficiency -- all members have access to all but at most a single drug -- without distorting incentives for innovation.

Kyna Fong
UC, Berkeley
50 University Hall, MC 7360
Berkeley, CA 94720-7360
kyna.fong@stanford.edu

Michael Schwarz
2397 Shattuck Ave
Berkeley, CA 94704
and NBER
mschwarz@yahoo-inc.com

Towards an Efficient Mechanism for Prescription Drug Procurement*

Kyna Fong† Michael Schwarz‡

February 1, 2009

Abstract

This paper applies ideas from mechanism design to model procurement of prescription drugs. We present a mechanism for government-funded market-driven drug procurement that achieves very close to full static efficiency – all members have access to all but at most a single drug – without distorting incentives for innovation.

Prescription drugs are an essential component of modern healthcare. In the U.S. drug spending has skyrocketed in recent years – while in 1980 it amounted to $12 billion, less than 5% of total health care expenditures, in 2006 that number has increased to $216.7 billion, over 12% of total expenditures. The Center for Medicare and Medicaid Services (CMS) currently projects drug spending to rise to $446.2 billion in 2015, approximately 2.2% of projected U.S. GDP.[1] Those statistics suggest that even minor efficiency improvements in the drug procurement process can lead to significant welfare gains.

Lessons learned from markets for typical healthcare services such as doctor consultations and hospital stays are largely not applicable to markets for prescription drugs. In contrast

*We would like to thank Jeremy Bulow, Robin Lee, Preston McAfee, Michael Ostrovsky, Andy Skrzypacz, Joel Sobel and seminar participants at Chicago GSB, Columbia Department of Economics, NBER Summer Institute and Stanford GSB for helpful comments.

†UC Berkeley and Stanford University, contact: kyna.fong at stanford.edu. Part of this research was conducted during an internship at Yahoo! Research.

‡Yahoo! Research and NBER, contact: mschwarz at yahoo-inc.com

[1]Source: CMS Data and CBO Projections.

to healthcare services, prescription drugs are often supplied by patent-protected monopolists and feature marginal costs that are a very small fraction of price. In this respect, prescription drugs bear a stronger resemblance to information goods such as music and software than to doctor visits and hospital stays.

This paper offers a conceptual framework for efficient design of government-funded prescription drug procurement. We emphasize that this paper is intended primarily as a thought exercise rather than a specific policy proposal. We address a basic question – whether efficiency is attainable in the market for prescription drugs in the presence of monopolists – and then use a highly stylized model to establish a benchmark answer in the affirmative. Our main objective is to use economic theory to develop a framework for thinking about prescription drugs.

For drugs that are zero marginal cost, the socially efficient (static) outcome is for all individuals who derive a positive expected marginal benefit from a drug to have access to that drug: in other words, static efficiency entails universal access.[2] Due to patent protection, however, drugs are often sold by monopolists, and in the benchmark monopoly pricing case, the profit-maximizing price leads to deadweight loss. One contribution of this paper is to describe a novel market-based mechanism for government drug procurement that approaches static efficiency without introducing additional distortions to incentives for innovation (relative to the benchmark monopoly pricing case).

One might argue that appropriately chosen government price controls can accomplish the same goal. The government achieves static efficiency by providing all drugs for free to its citizens. In conjunction, the government can choose appropriate prices such that incentives for innovation are unchanged, thus leading to a Pareto gain in welfare. However, such an arrangement features two major drawbacks: (1) it lacks the power of the market to dynamically correct for inaccurate pricing, imposing strong informational requirements on

[2] We assume away externalities. In some cases, drug consumption imposes positive or negative externalities. For example, the use of vaccines may prevent an epidemic while the use of antibiotics may create drug-resistant strains of bacteria.

the government, and (2) it creates perverse incentives for drug manufacturers to manipulate prices via lobbying.

In the U.S. lawmakers consider the drawbacks of centralized price interventions serious enough to exceed their benefits. The Medicare Modernization Act of 2003, which has established a universal drug benefit for the elderly in the form of Medicare Part D, explicitly prohibits the government from negotiating drug prices. Michael Leavitt, the U.S. Secretary of Health and Human Services, has written that, "government should not be in the business of setting drug prices or controlling access to drugs. That is a first step toward the type of government-run health care that the American people have always rejected." He describes government price-setting as a situation in which "one government official would set more than 4,400 prices for different drugs, making decisions that would be better made by millions of individual consumers."[3] At least in the U.S., a market-driven procurement mechanism that avoids price controls might be a politically palatable way to improve efficiency.

In fact, the U.S. has a long history of using market forces to accompany government subsidization of prescription drug purchases. In the case of Medicaid, which provides health insurance for low-income Americans, in 1990 the federal government instituted a "best-price rule" that ties prices paid by Medicaid to prices paid by private payers. Another example of policy tying U.S. government subsidies in healthcare to market forces is the advantageous tax treatment of employer-provided health insurance. Although in both instances the intention of the government subsidies was to increase access to prescription drugs, studies show that those interventions have instead led to the perverse effect of raising prescription drug prices and in fact reducing access to drug coverage. Morton (1997) is the first to point out that the Medicaid best-price rule has put upward pressure on drug prices, and Duggan & Morton (2006) offers conclusive evidence of that. Using entirely different methods and data, Schwarz (2006) finds that the above policies not only have increased drug prices but, by increasing drug prices, have likely reduced the number of people with access to drugs, perversely the

[3] Source: Opinion editorial, *Washington Post*, January 11, 2007.

opposite effect of what policymakers intended.

In the presence of monopoly-supplied drugs, the main difficulty in designing a market-driven mechanism is that, as long as a monopolist faces a population of consumers with heterogeneous values, he is likely to set his monopoly price above the reservation price of some consumers, leading to potentially large deadweight losses. Even when the purchase of drugs is subsidized in the form of traditional insurance, the same problem persists when monopolists control prices (Newhouse 2004). Intuitively, if the government subsidizes 50% of each drug purchase, monopolists will double prices in response and deadweight losses will remain unchanged.

In this paper, we show that there exists a government-funded, market-driven drug procurement mechanism that gives all consumers access to all but at most a single drug, without introducing additional distortions to innovation. Despite monopoly pricing power, static near-efficiency can be achieved.[4,5]

In our model we assume that drugs have no substitutes and that each drug is sold by a monopolist. Although in practice drugs may have therapeutic substitutes, the setting in which all drugs are supplied by monopolists generates the largest concerns about efficiency and hence is an important benchmark to consider.

Let us describe the intuition behind the market-driven drug procurement mechanism considered herein. Recall that deadweight losses arise because the demand curve is downward-sloping, i.e. different consumers have different willingness to pay. If in fact all consumers were identical, then a monopolist would set price equal to the universal willingness to pay and serve the whole market. The key idea we leverage is to "homogenize" demand.

To illustrate that idea, consider a world in which all individuals have the same probability of disease. The essential source of demand heterogeneity then arises from different tastes

[4]The additional deadweight loss from taxation is zero so long as the funding for the mechanism equals the amount the government spends on subsidizing drug coverage under the status quo.

[5]The idea that mechanisms excluding one efficient trade can achieve near efficiency has been proposed in a very different setting by McAfee (1992). Specifically, in McAfee (1992) one of many identical items is excluded, whereas in our setting one of many drugs (which are not substitutes) is excluded from coverage.

and wealth. For true monopoly drugs, which have no substitutes, we argue that demand heterogeneity takes a specific form: although different consumers may have different willingness to pay for a given drug, *the ratio between willingness to pay for any two monopoly drugs is constant across all consumers.*[6]

This assumption about the special structure of the demand for drugs is theoretically grounded. In particular, the assumption is motivated by the literature on the value of life, which offers a theoretical and empirical framework for thinking about investments in safety (see Viscusi (1993) for a survey). The main idea is that if a fire detector and an air bag are equally likely to save a life, then any consumer should place equal reservation prices on both devices. In our context, we argue that drugs or insurance for drug coverage can be viewed as particular cases of safety devices that "produce" life or, more accurately, health. Hence, heterogeneity in demand for drugs arises simply from heterogeneity in *willingness to pay for health.* This special demand structure is exactly what allows us to construct a market mechanism that "homogenizes" demand.

The structure of the mechanism at a high level conveniently resembles that of Medicare Part D in the U.S., which is also government-funded and market-driven with private prescription drug plans (PDP's) serving as intermediaries between consumers and drug companies. The government first fixes a per-member budget (subsidy) B to give to each PDP. Drug companies announce prices at which the companies are willing to sell their drugs to PDP's. Drug plans then procure drugs from drug companies and assemble formularies, i.e. the set of drugs that a given PDP will cover, within the budget constraint B. Consumers choose a PDP in which to enroll based on formulary comprehensiveness. The drugs that are not on a formulary continue to be available on the open market at market prices.

Fixing the subsidy serves to homogenize demand: each person now has the same budget to be spent on her behalf, and thus the willingness to pay for access to a given drug becomes

[6] Specifically, we mean the ratio is constant across all consumers with the same probability of disease (i.e. either all healthy consumers, or all consumers who have already developed the disease).

uniform across all people.[7] Thus, the subsidy forces drug companies to "compete" with one another to get on PDP's formularies; as a result, a monopolist can only charge as much as the relative social value of his product warrants, thus making near efficient outcomes possible.

This mechanism imposes much lower informational requirements on the government than administering price controls. In order to specify an optimal per-person budget B, the government needs only information about aggregate profits of all drug companies. Furthermore, in Section 2.1 we argue that the ramifications on drug access of choosing a budget size that is too high or too low are relatively small. In addition, the mechanism generates a natural feedback loop. If in a given year many drugs are left off formularies, then the government knows to increase the budget size in the next year; on the other hand, if all drugs are covered, then the government may consider lowering the budget size. The mechanism also allows budget size to be used as a convenient tool for calibrating incentives for innovation in the pharmaceutical industry.

Kremer has described a patent buyout mechanism that achieves near efficiency in the presence of monopolists (Kremer 1998). Buying out patents hinges on the notion of a two-part tariff.[8] By transferring lump sums of the appropriate size to inventors (would-be monopolists) in exchange for patents, the government attains static efficiency because the invention can then be made available at marginal cost.

The major difficulty with implementing such a system lies in determining the size of the lump sum. Kremer's mechanism extracts the relevant information from a market of private parties, giving them weak incentives to reveal this information truthfully at the cost of introducing a small inefficiency. He points out that the main problem with his mechanism is that it is susceptible to manipulation. The expected joint payoff to a patent holder and any third party whom he controls can be arbitrarily large, making collusion very attractive.

[7] This is a consequence of consumers having the same relative valuations for all drugs. Given the same budget and the same probabilities of disease, consumers will allocate their budget across drugs (or insurance for drugs) in the same way.

[8] Lakdawalla and Sood (2006) also focus on two-part tariffs; in particular, they describe how the two-part design of health insurance contracts combats the inefficiency that results from healthcare providers having market power.

In the context of prescription drugs, manipulation is a serious concern since there are only relatively few pharmaceutical companies and they would be regular repeated participants in patent buyout auctions. The mechanism we describe cannot be manipulated in this fashion.

While Kremer's (1998) patent buyout mechanism relies on the notion of a two-part tariff, the mechanism we describe does not – in fact, it can be implemented using either linear pricing or two-part pricing. Instead, the mechanism relies on a fundamentally different idea, *homogenizing demand*, which can be leveraged in drug procurement due to the unique demand structure for drugs.

There are several features of the prescription drug market that our model ignores completely. We assume that consumers are rational and capable of identifying the highest-value formulary, and we restrict attention to the scenario of pure monopoly drugs, which have no substitutes. We also assume that consumers either have identical probabilities of contracting each disease, or that a mechanism for perfect risk adjustment is available. Furthermore, we assume that the marginal costs of producing drugs and the costs of dispensing drugs are negligible, thus assuming away the importance of existing pharmacy networks. Nonetheless, our highly stylized model may offer a useful framework for thinking about drug procurement.

1 Model

1.1 The Environment

There is a finite number N of diseases $n \in \{1, 2, \ldots, N\}$. Each disease is treated by a unique drug, and a treatment requires exactly one unit of a drug (drugs have no substitutes). Each drug is sold by a patent-protected monopolist.[9] Drugs are produced at zero marginal cost, and drug companies maximize expected profits.

For each disease n, let $\theta_n \in (0, 1]$ be the probability that an individual contracts the

[9] If each drug company owns several drugs, we can rephrase our analysis in terms of each company's portfolio of drugs rather than single drugs, and all the results would remain the same.

disease. Let $v_n \geq 0$ be the ex-post *effectiveness* of drug n. We can think of v_n as the increase in future "health" units (measured, for example, in quality-adjusted life years) that the drug provides to an individual who has developed the disease. Without loss of generality, label drug N as a maximally valuable drug, i.e. $v_N = \max_n v_n$. Finally, define $z_n = \theta_n v_n$ as the *expected benefit* of drug n.

Consumers' utility is strictly increasing in future "health" or life quality and is quasi-linear and strictly increasing in money. Consider some consumer a. Consumer a's willingness to pay for insurance that covers drug n is related to the probability of disease (θ_n) as well as the drug's ability to increase the quality of her future life-years (v_n). Let u^a be consumer a's willingness to pay for insurance coverage of an additional unit of future life quality.[10] We assume that u^a is distributed in the population according to some probability distribution $\rho(\cdot)$ with cumulative distribution function $P(\cdot)$. So, we can write a consumer's willingness to pay (in dollars) for insurance that covers drug n as $u^a \theta_n v_n = u^a z_n$.[11]

The population consists of a unit mass of consumers, although all results still hold for a single consumer or any finite number of consumers.[12] We assume that consumers have access to a market for actuarially fair insurance, which is available "à la carte" on an individual drug basis. Thus, thinking about a consumer's willingness to pay for actuarially fair insurance coverage of a drug is equivalent to thinking about her willingness to pay for that drug.

There are potentially two sources of heterogeneity in willingness to pay for drugs: differences in health status and differences in willingness to pay for health. We focus on the latter by assuming that all consumers have the same probability θ_n of contracting any disease n, or equivalently that a perfect risk adjustment mechanism is available.[13] Assuming perfect risk

[10]Assuming u^a to be constant is a reasonable approximation for small probability events.

[11]We assume that the per-unit drug prices faced by individual consumers on the open market are equal to (or not lower than) the per-unit prices faced by insurance companies that sell actuarially fair insurance. Risk-averse consumers strictly prefer to purchase actuarially fair insurance for drugs. Consequently, if a consumer chooses not to purchase actuarially fair insurance for a particular drug, then she will choose not to purchase that drug in the event she does get sick.

[12]We model consumers as a continuum whose mass is normalized to unity because speaking about the proportion of consumers with access to a drug is more natural than speaking about the probability that a consumer has access to a drug.

[13]Risk adjustment plays an essential role in the design of Medicare Part D and many other insurance

adjustment focuses attention on an already complicated benchmark for prescription drug procurement without additional complications due to adverse selection. Furthermore, when all drugs are supplied by monopolists, it is theoretically possible to construct a perfect risk adjustment mechanism (see Fong & Schwarz (2009)).

Finally, there is a perfectly competitive market of risk-neutral prescription drug plans (PDP's), which engage in Bertrand-style competition. PDP's serve as intermediaries between consumers and drug companies. Each PDP assembles a formulary of drugs to which its members have access. If a drug is not covered by a consumer's PDP, the consumer has the option of purchasing the drug, or actuarially fair supplemental insurance for that drug, on the open market.

Consumers have private information about their willingness to pay for health u^a. However, the distribution $\rho(\cdot)$ is common knowledge. In addition, the expected benefits of each drug $z_n = \theta_n v_n$ are common knowledge among consumers, drug companies, and PDP's.[14] The government, however, does not know the values of each drug.

1.2 A Simple Example

Before presenting the mechanism in detail and proving general results, we first present a simple example to illustrate the main ideas.

Consider a world in which there are $N = 151$ drugs, 100 of which provide value of one to consumers who have developed the associated disease, i.e. $v_n = 1$, and 51 of which provide value of two to consumers who have developed the associated disease, i.e. $v_n = 2$. We have a unit mass of consumers indexed by a, and in this world, we assume that all are healthy and develop each disease independently with identical probability $\theta_n = 0.01$ for all n. However, consumers do not have the same willingness to pay for drugs due to differences

programs.

[14]The assumption that consumers know the relative values of every drug is not necessary for our results to hold. Instead it is sufficient to assume that consumers are capable of selecting the most valuable formulary from amongst those offered by PDP's. That assumption would be satisfied in a world in which trustworthy entities provide consumers with formulary rankings based on value.

in preferences and in wealth. In particular, we assume that the marginal willingness to pay for coverage of an additional unit of future health, u^a, is uniformly distributed between 0 and 100. Conditional on consumers being healthy then, the willingness to pay for actuarially fair insurance coverage of any drug is uniformly distributed between 0 and v_n.

Notice that each drug company faces a linear, downward-sloping demand curve. In the benchmark case of the open market, each drug company acts as a profit-maximizing monopolist and sets its price at $50v_n$. As a result, actuarially fair insurance costs $0.5v_n$ per drug, and half of the population, i.e. those with willingness to pay above 0.5, purchases insurance and receives drugs if needed. Expected profits to each monopolist are $0.25v_n$, and the total amount spent by consumers is $(0.25 \times 100) + (0.5 \times 51) = 50.5$. In such a market, half the population is excluded from all drugs. Even if the government subsidizes the purchase of drugs by covering half their costs, monopolists will double prices in response and half of the population will continue to be excluded.

Alternatively, consider the following market-based, government-funded drug procurement mechanism. The government commits to a maximum budget of 50 for procuring drugs for the whole population. However, the value of each drug is not known to the government. Hence, it delegates drug procurement and formulary assembly to two private competing prescription drug plans (PDP's), who know the value of drugs, and gives each PDP a budget of 50 scaled by market share.

Through a centralized procurement auction, each drug company is asked to specify the price at which it is willing to sell its drug. PDP's can purchase actuarially fair insurance for drugs at those prices, and each PDP uses those prices to decide which drugs to cover on its formulary subject to the budget constraint. Drug companies can choose not to participate in the auction, and if a drug company opts out, or if its drug is not included on the formulary, the company has the right to sell its drug on the open market at its monopoly profit-maximizing price. In the procurement auction, even though the profit-maximizing market price is $50v_n$, each company is willing to bid a price as low as $25v_n$ because doing so implies an expected profit of $0.25v_n$, the same profit the company obtains from opting out of the procurement auction and selling on the open market. Notice, however, that in equilibrium not all drugs can be fully covered because meeting all drug companies' outside option profits of $0.25v_n$ exceeds the budget, i.e. $(0.25 \times 100) + (0.5 \times 51) > 50$.

In equilibrium, all companies must bid at the exact same price-to-value ratio because otherwise companies bidding at lower ratios can shade upwards and just undercut companies bidding at higher ratios. As a result, the only equilibrium involves all drug companies bidding a price of exactly $25v_n$. Actuarially fair insurance for each drug then costs $0.25v_n$. Due to competition for membership, each PDP assembles its formulary to maximize social value – for example, it is an equilibrium for each PDP to include all of the 100 value-1 drugs as well as 50 of the value-2 drugs randomly chosen from the 51.[15]

Under this alternative mechanism, the entire population can have access to all but a single drug, yielding over 99% efficiency, while total spending remains unchanged. Notice that the approximate efficiency result is not a coincidence. We have chosen the budget to be exactly equal to $\sum_{n=1}^{N-1} r_n$, which is the maximum budget that satisfies Proposition 1 in Section 2.1. In Figure 1, this budget is the highest value for which the mechanism achieves exactly first-best.

Even though both *the total spending on drugs* and *the profits to drug companies remain the same as under the benchmark monopoly regime, the share of the population with access to almost all drugs doubles. Because drug companies' profits are unchanged, incentives for innovation remain unchanged as well.*

1.3 The Mechanism

We now describe the mechanism in detail. While it is publicly-funded, implementation is private – private prescription drug plans (PDP's) receive government subsidies but must compete to attract members.

[15] Note that any other formulary with the same social value is also an equilibrium.

The government first sets a *fixed* per-member subsidy B for PDP's to use in assembling formularies. Each PDP then assembles a formulary of drugs to which members have access at marginal cost (which we assume to be zero). PDP's do not have the option of charging a membership premium. If a member wishes to buy a drug not on her PDP's formulary, she has the option to purchase the drug (or actuarially fair drug insurance) at market prices.[16]

Prices are set for all PDP's via a single procurement auction. In the auction each drug company submits a bid b_n, indicating the minimum price the company is willing to accept for its drug. This creates a vector of bids $b = (b_1, b_2, \ldots, b_N)$. Given the fixed budget constraint B and the bid vector b, each PDP assembles a formulary. Any unused budget becomes a PDP's profit. PDP's compete for membership, and since consumers choose PDP's with the most valuable formulary, each PDP assembles the most valuable formulary possible within the budget in order to attract members.

A formulary consists of a set of drugs that the PDP covers. To avoid integer constraints, we allow PDP's to offer partial coverage. Having a partially covered drug on the formulary means that the drug is covered for a randomly selected fraction, strictly between zero and one, of the PDP's members. In particular, if the remaining budget is ϵ and the next drug the PDP wants to cover has bid $\hat{b} > \epsilon$, then the PDP covers that drug for a fraction $\frac{\epsilon}{\hat{b}}$ of members (in equilibrium at most one drug is partially covered).[17] To break ties, we assume that in assembling formularies PDP's have a lexicographic preference for maximizing the number of fully covered drugs.[18]

The strategic component of the mechanism lies in the procurement auction – the bids submitted directly determine formulary compositions of PDPs. Thus, analysis surrounds equilibrium bidding behavior in the auction.

[16]Note that instead of PDP's, it would be analytically identical to assume consumers assemble their own "à la carte" formularies.

[17]Partial coverage of drugs is not unheard of. Drug plans routinely cover certain drugs on a case-by-case basis, although these decisions aren't meant to be random.

[18]Removing this tie-breaking assumption would have no effect on the expected social value of equilibrium formularies. Including it allows us to focus on the equilibrium formulary that has the maximum number of fully covered drugs.

The auction format we consider is a sealed bid (simultaneous-move) auction. Each drug company n submits a private bid $b_n \in [0, \infty)$ to the auctioneer, or chooses not to participate. Note that drug companies do not learn any information about other bids. One way to view the sealed bid auction format is as a model of unstructured negotiation where monopolist drug companies have all the bargaining power.

2 Analysis

In this section, we lay the necessary groundwork and then analyze equilibrium behavior.

The drug procurement mechanism we have described in the previous section is voluntary. A drug company can choose not to participate and instead to sell its drug on the open market as a monopolist. We characterize a drug company's participation constraint by defining a company's *reserve fee* in the procurement auction: that fee is the minimum "per-person profit," i.e. ratio of expected total profit to population size, that the company requires in the auction. Denote drug company n's reserve fee by r_n. In Lemma 1, we show that this reserve fee is proportional to the expected health benefit z_n of drug n.

We first consider per-person budgets B that are "sufficiently large" to satisfy all drug companies' participation constraints and to achieve full efficiency (in theory), i.e. $B \geq \sum_{n=1}^{N} r_n$. We show that in this case almost all drugs are available to all consumers (at zero cost). Later in Section 2.1, we consider the case when the budget is small and show that the mechanism performs well in that case also.

We claim that under a sealed bid procurement auction, the mechanism achieves near efficiency. The following theorem makes our claim formal.

Theorem 1. *Let $B \geq \sum_{n=1}^{N} r_n$. In any Nash equilibrium of the sealed bid (simultaneous-move) procurement auction, the expected social value of the equilibrium formulary is greater than or equal to a formulary that fully covers $N-1$ drugs, i.e. a formulary that fully covers all drugs but one (which may be the drug with the highest reserve fee).*

Before proceeding to the proof, we introduce some preliminary notation and results. Let bid b_n denote the price for which drug company n indicates it is willing to sell one unit of its drug. Notice that b_n implies that in expectation, the per-person price of covering drug n is $\theta_n b_n$. To make a non-negative profit, a PDP must choose a formulary where the sum of per-person prices of all drugs on the formulary does not exceed the budget B.

It is convenient to consider *normalized bids*, defined for drug company n as $e_n = \frac{\theta_n b_n}{r_n}$. A normalized bid is equal to the expected per-person price of drug n, scaled by the reserve fee r_n, which is proportional to drug n's expected health benefit. Focusing on normalized bids is more convenient analytically, as doing so places drug companies' bids on the same "scale" of expected per-person price per unit of expected benefit. Normalized bids are particularly useful because, given that PDP's assemble formularies to maximize total expected benefit, a PDP will never include a drug with some normalized bid without also including all drugs with lower normalized bids. Notice that to satisfy a company's participation constraint, the normalized bid must be at least 1, i.e. $\theta_n b_n \geq r_n$.

We now explicitly characterize a drug company's reserve fee, i.e. r_n. Recall that if a drug company does not participate in the auction, it sells its drug on the open market as a monopolist instead. Because our model considers only drugs without substitutes, the profit the drug company receives from not participating in the government mechanism is exactly equal to the monopoly profits the company would receive in the absence of any government intervention. We can show that a company's total monopoly profit is in fact linear in its drug's average expected benefit, i.e. $z_n = \theta_n v_n$, and hence each company's reserve fee r_n is linear in z_n.[19] Formally, we have the following lemma:

Lemma 1. *There exists a constant $K > 0$ such that $r_n = K z_n$ for all $n \in \{0, 1, \ldots, N\}$.*

Proof. This proof follows Schwarz (2006). Denote the profit-maximizing monopoly price

[19]Since all PDP's are identical and consider the same bids from drug companies, in equilibrium all formularies will be the same. If a drug is covered on any formulary, it will be covered on all formularies and hence be available to the entire population. Thus, to match its outside option, a drug company need only submit a bid that is equal to its outside option profits (i.e. open market monopoly profits) divided by the entire population.

by ψ^*. Recall that a market for actuarially fair insurance exists and that a consumer is willing to pay $u^a \theta_n v_n$ for insurance coverage of drug n. When the price of drug n is set at ψ_n, the demand for insurance coverage for drug n is equal to the mass of consumers for whom $u^a \theta_n v_n$ exceeds $\theta_n \psi_n$, i.e. $u^a v_n \geq \psi_n$. We can write (expected) conditional demand as $D(\psi_n) = \theta_n \int_{\psi_n/v_n}^{\infty} dP(u^a)$. That represents the fraction of consumers who are willing to buy insurance coverage for drug n, when the price of drug n is ψ_n. So ψ^* solves the following profit maximization problem:

$$\max_{\psi_n} \psi_n \theta_n \int_{\psi_n/v_n}^{\infty} dP(u^a)$$

Notice that demand depends only on the ratio between price ψ_n and effectiveness v_n. Therefore each drug company maximizes its expected profits by serving the same proportion of the population, i.e. choosing the same ratio of price ψ_n to effectiveness v_n. As a result, a drug company's total monopoly profit is linear in its drug's average expected benefit z_n. Therefore, for any n, the company's reserve fee r_n is also linear in z_n. $\square$

The proof of Theorem 1 relies on three additional lemmas.

Lemma 2. *If drug company n submits a normalized bid of $\underline{e} = \frac{B}{\sum_{i=1}^{N} r_i}$, then in any equilibrium drug n will be fully covered. As a result, submitting a normalized bid of $e_n < \underline{e}$ is a strictly dominated strategy for any company n.*

Proof. If all drug companies submit a normalized bid of $\underline{e}$, then full coverage of all drugs will just fit within the budget. Hence, a normalized bid of $\underline{e}$ guarantees drug n is fully covered on the formulary, and bidding any lower leads only to a strictly lower payoff. (Note that bidding $\underline{e}$ satisfies drug company n's participation constraint since $B \geq \sum_{i=1}^{N} r_i$.) $\square$

Lemma 3. *There exists a constant $\zeta > 0$ such that for any equilibrium bid for any drug, the probability of that drug being at least partially covered is greater than or equal to ζ.*

Proof. Recall from Lemma 2 that $\underline{e}r_n$ is a lower bound on the expected payoff to drug company n. We now determine an upper bound on that payoff.

If drug n is partially or fully covered, a strict upper bound on its maximum payoff is $B + r_n$, since B is the maximum payoff drug n receives from its covered portion and r_n is a strict upper bound on the maximum payoff drug n receives from its uncovered portion. On the other hand, if drug n is not covered, the maximum payoff is r_n. Thus, if ζ is the probability of being at least partially covered, a strict upper bound on drug company n's payoff is $\zeta(B + r_n) + (1 - \zeta)r_n = \zeta B + r_n$.

Hence ζ must be large enough that, for all n, $\zeta B + r_n > \underline{e}r_n$, i.e. $\zeta > (\underline{e} - 1)\frac{r_n}{B} \geq 0$. $\square$

Before stating the next lemma, we need to introduce two definitions. First, for a given vector of bids, define d as the sum of reserve fees of all fully uncovered drugs *plus* the reserve fees of all partially covered drugs multiplied by the uncovered share of those drugs. In other words, d measures the total profits earned by drug companies on the open market (not covered by the formulary). One consequence then is that the sum of all drug companies' profits combined is equal to $B + d$.

Secondly, since d is a random variable in any mixed strategy equilibrium, define $\bar{d}$ as the expected value of d.

Lemma 4. *If $\bar{d}$ is greater than the monopoly profit of the most valuable drug, i.e. $\bar{d} > r_N$, then in* any *equilibrium, for* any *equilibrium realization of bids, there is at least one drug that has zero coverage and there are at least two drugs that are not fully covered (but may be partially covered).*

Proof. First, we remark that with probability one there is at most one partially covered drug in equilibrium. Consider any two drugs with a positive probability of not being fully covered. The probability of both drugs having the same normalized bid is an event of measure zero because otherwise any of the bidders could strictly increase its payoff by lowering its bid by an arbitrarily small amount in this event. Thus, without loss of generality we assume that there is at most one partially covered drug in any equilibrium formulary realization.

Suppose there exists an equilibrium realization of bids such that all drugs are covered at least partially and therefore at least $N-1$ drugs are fully covered. Since any drug that is fully covered under a particular realization of bids receives a profit that is greater than or equal to the expected profit that the drug receives in equilibrium, it follows that $B+r_N$ must be greater than or equal to the total expected profit of the drug companies, which is equal to $B+\bar{d}$. Hence, it follows that $\bar{d} \leq r_N$, which contradicts the assumption of the lemma.

$\square$

Proof. **(Theorem 1)** We prove Theorem 1 by showing that $\bar{d}$ must be smaller than the reserve fee of the most valuable drug. We begin by assuming the contrary and then arrive at a contradiction.

Suppose that $\bar{d}$ is greater than the reserve fee (i.e. monopoly profit) of the most valuable drug. Lemma 4 then tells us that under any equilibrium realization of bids at least one drug has zero coverage and at least two drugs are not fully covered. But as long as there is always one fully uncovered drug for any equilibrium vector of bids, we show below that there must exist an equilibrium action that leads to a drug not being covered with arbitrarily high probability, hence contradicting Lemma 3.

Consider the range of normalized bids that are consistent with an equilibrium (a mixed strategy involves a drug company randomizing over a range of normalized bids). Denote the supremum of that range for drug company n as e_n^{max}. If the set is unbounded, let $e_n^{max} = \infty$. Also, let $e^{max} = \max_n e_n^{max}$ so that e^{max} is the highest normalized bid encountered in equilibrium.

Now consider some drug company i for whom $e_i^{max} = e^{max}$. As company i's bid approaches e^{max}, the probability that drug i will have zero coverage approaches one. However, this violates Lemma 3. Hence, we have a contradiction and therefore conclude that $\bar{d}$ must be less than the reserve fee of the most valuable drug. $\square$

2.1 Small Budgets

We have shown that sufficiently large budgets yield near efficient outcomes. We now show that when the budget is small, the mechanism performs very well.

We begin by defining a "first-best" benchmark for comparison. Recall that in our model we consider a world in which the government does not know the values of each drug. Let us now consider a counterfactual world in which the government knows exactly the values of each drug (and hence the monopoly profits each drug would generate on the open market). Moreover, the government has the power to make lump-sum take-it-or-leave-it offers to drug companies in exchange for any of its citizens being able to consume a company's drug when needed. In this benchmark world, the government assembles the most valuable formulary subject to two constraints: (1) the cost per person does not exceed the budget B, and (2) each drug company receives exactly its monopoly profits (thus ensuring that participation is voluntary and there are no distortions to innovation). We will refer to the value of that formulary as *first-best.*

In our model, because the government is not informed about the values of each drug, first-best may not be achievable. For small budgets, however, Proposition 1 shows that in fact the first-best formulary value is achievable.

Proposition 1. *Let $B \leq \sum_{n=1}^{N-1} r_n$. In the unique equilibrium of the sealed bid procurement auction, the expected social value of the equilibrium formulary is exactly equal to* first-best*, i.e. the maximum formulary value achieved in the benchmark world in which the government knows exactly the values of each drug and has the power to make lump-sum take-it-or-leave-it offers to drug companies.*

For small budgets, the mechanism we have described leads to a formulary in which all on-formulary drugs are covered at exactly a normalized bid of one. By paying for on-formulary drugs an amount that is equal to the profits that would have been generated under monopoly pricing, the mechanism provides access to on-formulary drugs for all of the population, and

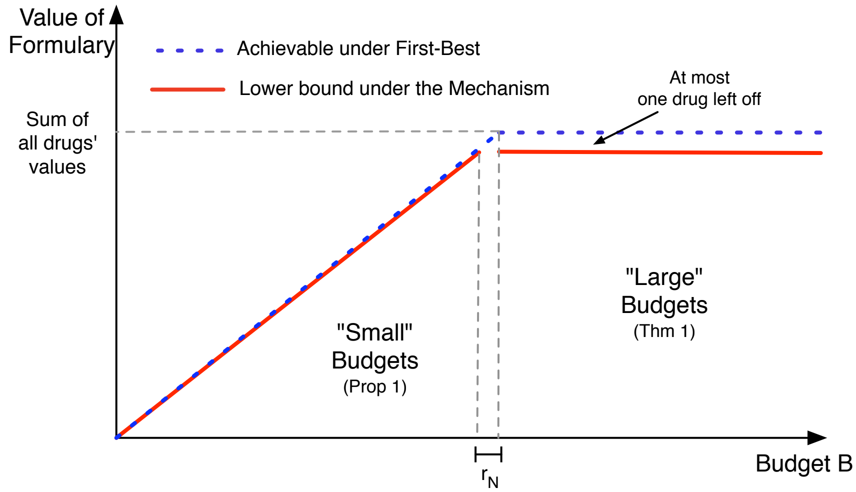

Figure 1: Illustration of formulary value achieved by mechanism, as a function of budget B

not just the fraction that can afford those drugs at monopoly prices.

The proof of Proposition 1 is straightforward. If the budget is not large enough to meet all drug companies' reserve fees, then competition to get on the formulary will drive each company to bid down to a level that generates expected profits equal to its reserve fee of monopoly profits. No drug company can get on the formulary with a normalized bid higher than one because it will be undercut by some other drug company. Meanwhile, no company will submit a normalized bid less than one because its participation constraint will be violated.

Figure 1 presents an illustration, under both "small" budgets (Proposition 1) and "large" budgets (Theorem 1), of the formulary value achieved by the mechanism relative to first-best. Notice that when the budget is small, those two values are identical – the formulary value attained by the mechanism is equal to the first-best. When the budget is large, the formulary value attained by the mechanism falls below the first-best, but only by an amount less than or equal to the value of a single drug.

2.2 Incentives for Innovation

Determining what level of profits correspond to "optimal" incentives for innovation is outside the scope of this paper, and we remain agnostic about what profit levels are optimal – that is the topic of a separate, well-developed literature.[20]

We can show, however, that the mechanism performs well at preserving relative incentives for innovation across drugs. When the budget is small, i.e. $B < \sum_{n=1}^{N-1} r_n$, then from Section 2.1 above we know that incentives for innovation are exactly unchanged – all drug companies receive exactly their monopoly profits.

Now consider the case of large budgets $B \geq \sum_{n=1}^{N} r_n$. We can show that total profits are distributed across all drug companies in close proportion to expected benefit. Fix any equilibrium and let e_n^{avg} denote the expected normalized profit to drug company n. We know from Lemma 2 that $e_n^{avg} > \underline{e}$, where $\underline{e} = \frac{B}{\sum_{n=1}^{N} r_n}$ is a lower bound on the normalized profit any drug company can receive. To understand the distribution of profits across drug companies, we are interested in how much drug companies' expected profits deviate from that lower bound and compute the following ratio:

$$\frac{\sqrt{\sum_{n=1}^{N} (e_n^{avg} r_n - \underline{e} r_n)^2}}{B + \bar{d}},$$

where the numerator, i.e. the square root of the sum of squared residuals between drug companies' expected profits and the lower bound, is normalized by total profits across all drug companies.

As a corollary of Theorem 1, we show that this ratio goes to zero as the value of the largest drug becomes small relative to the sum of the values of all drugs. To do so, we consider the worst-case scenario. The numerator takes its maximum value when all drug companies have expected normalized profits of exactly $e_n = \underline{e}$ except for one company, which receives all remaining profits. Thus, in the worst-case the total deviation of expected profits divided

[20]For example, see Boldrin and Levine (2006), Garber,ones and Romer (2006), Hopenhayn, Llobet and Mitchell (2006), and Scotchmer (1999).

by total profits, i.e. the ratio above, is exactly $\frac{r_N}{B+\bar{d}}$, which is less than or equal to $\frac{r_N}{\sum_{n=1}^{N} r_n}$. Thus, as the value of the largest drug becomes small relative to the sum of the values of all drugs, the total deviation of expected profits among drugs from the lower bound goes to zero.

Hence, the mechanism described in the above section continues to do well at preserving relative incentives for innovation even when the budget is large. In fact, Proposition 2 in the next section shows that under an ascending clock auction format, relative incentives for innovation across drugs are perfectly preserved even when budgets are large. Notice that the budget size B serves as a convenient tool for calibrating the strength of incentives for innovation in prescription drug markets.

2.3 Extension

The previous section assumes that PDP's are risk neutral players. The mechanism described in the above section could easily be modified so that we can get the same efficiency results even if PDP's were risk averse and unable to diversify the risk. In particular, the modification would involve changing the bid format to *per-person prices*.

Thus far the bid format we have been considering is a conventional linear price: the bid specifies the price that drug companies charge to PDP's per unit of the drug. An alternative bid format is a per-person price: the bid specifies the price that drug companies charge to PDP's per consumer covered by the PDP regardless of whether that consumer will need the drug. Collecting a per-person price for a consumer obligates the drug company to provide the drug for free in the event the disease is contracted. Note that the results of the previous section continue to hold if bids are submitted as per-person prices rather than linear prices.

Per-person pricing improves risk-sharing. Linear prices create uncertainty about how much PDP's will have to pay to drug companies and hence create artificial risk for both PDP's and drug companies, which produce drugs at zero marginal cost. Under per-person pricing, that uncertainty disappears.[21]

We also point out that, in the mechanism described in the previous section, generically all equilibria are in mixed strategies (see Appendix A.1). The fact that no pure strategy equilibrium exists may be seen as an undesirable property. However, mixed strategy equilibria arise due to the simultaneous-move nature of the mechanism; in mechanisms where players move sequentially, pure strategy equilibria can be attained.

In this section, we consider auctions in which drug companies bid in terms of per-person prices rather than linear prices, so PDP's need not be risk neutral. To avoid ties, we assume that for all $i \neq j$, $v_i \neq v_j$. We consider two sequential-move auctions: the ascending clock auction and the descending clock auction. We prove that each auction features a unique equilibrium in pure strategies that achieves near efficiency.

Fix some $p^{max} \in [B, \infty)$. The rules of the auctions are as follows:[22]

- *Ascending Clock Auction.* The clock begins at per-person price $p = 0$ and increases continuously until $p = p^{max}$. Each drug company n chooses a level at which to leave the auction, and that becomes its per-person bid p_n. There is no re-entry permitted. Upon observing a dropout at p, a bidder can drop out only at a level strictly higher than p. Ties are broken in favor of later bidders.

- *Descending Clock Auction.* The clock begins at per-person price $p = p^{max}$ and decreases continuously until $p = 0$. Each drug company n chooses a level at which to leave the auction, and that becomes its per-person bid p_n. There is no re-entry permitted. Upon observing a dropout at p, a bidder can drop out only at a level strictly lower than p. Ties are broken in favor of later bidders.

The following propositions show that, like the sealed bid auction, both clock auctions

[21] Admittedly, in reality the risks generated by linear drug pricing are easily diversifiable. Nonetheless, it is theoretically interesting that such risks can be avoided using per-person pricing. For the auctions considered in this section, per-person pricing makes the analytical treatment simpler and more transparent.

[22] For completeness, we include non-participation as a viable action in both auctions (although it is not important for equilibrium analysis). So the action space for bidders in each auction consists of either submitting a per-person bid $p \in [0, p^{max}]$ in the auction or not participating at all.

lead to near efficient outcomes. Unlike the sealed bid auction, however, those outcomes are achieved in unique, pure-strategy equilibria. Proofs are relegated to the Appendix.

Proposition 2. (Ascending Clock Auction) *Let $B \geq \sum_{n=1}^{N} r_n$. The unique subgame perfect equilibrium of the ascending clock auction is a pure strategy equilibrium in which all drugs are fully covered on the formulary except for the most valuable drug, which is partially covered. Moreover, the total profit of each drug company n is in constant proportion to its drug's expected health benefit z_n.*

In the ascending clock auction, the most valuable drug, or the drug with the highest expected benefit z_n, is only partially covered on the formulary. The intuition is that in the ascending clock auction, the last bidder to bid, i.e. the company with the most valuable drug, can bid arbitrarily high (up to the maximum bid p^{max}) so that it collects not only the excess budget but also some monopoly profits from the open market since its drug is only partially covered. Other bidders cannot act to lower the payoffs of this highest-value drug company, because it has the last-mover advantage – it can always threaten to undercut previous bidders and get its drug fully covered on the formulary.

Proposition 3. (Descending Clock Auction) *Let $B \geq \sum_{n=1}^{N} r_n$. The unique subgame perfect equilibrium of the descending clock auction is a pure strategy equilibrium in which all drugs are fully covered on the formulary except for one drug, which is partially covered. The partially covered drug is* not *the most valuable drug.*

The intuition is similar to the ascending clock auction. The main differences are that it is (a) not the highest-value drug that is left partially covered and (b) not necessarily the last company to bid that is left partially covered. Instead, the identity of the drug company that remains partially covered depends on the distribution of drugs' values. What drives these differences is that in the descending clock auction, a given bidder's maximum bid is bounded from above because bidders with lower-value drugs bid later. So no company can submit a higher bid than what a company with a higher value drug has already submitted.

Although both clock-auctions achieve near-efficiency, covering all drugs but at most one, we show in Appendix A.3 that the descending clock auction generally leads to a formulary with higher social value than the ascending clock auction.

3 Concluding Remarks

There are several intuitive approaches for increasing access to prescription drugs using government-funded market-driven mechanisms. One approach is for the government to use market prices as benchmarks for the prices it pays, as in the case of Medicaid. Another approach is for the government to directly subsidize health insurance by making its purchase tax-deductible, as in the case of employer-provided health insurance. Research has shown, however, that both those approaches may have the unintended effect of raising drug prices and thus reducing rather than increasing the number of people with access to prescription drugs (Morton (1997), Schwarz (2006), Duggan & Morton (2006)).

Yet a third approach, which we study in this paper, is to make monopolists "compete" for scarce portions of a fixed subsidy, as in the case of Medicare Part D.[23,24] We show that at least in theory this third approach may lead very close to universal access: we describe a market-driven mechanism with fixed government subsidies in which all consumers have access to all but at most one drug.

The structure of the mechanism we describe is similar in many ways to the current form of Medicare Part D. Medicare's drug benefit is a government-funded, market-driven federal entitlement program in which private prescription drug plans receive fixed per-member subsidies from the government but must compete through assembling formularies and setting premiums to attract members.

There are several major points of distinction. First, our efficiency results hinge on creating

[23]For details on the design of Medicare Part D, see McAdams & Schwarz (2007).

[24]Notice that this logic has interesting antitrust implications. Because non-substitue drugs compete with one another based on cost effectiveness, a merger between two drug companies can put upward pressure on prices paid by Medicare even if no two drugs, one from each company's portfolio, are substitutes.

competition between monopolists for coverage by PDP's. Hence, it is crucial that every drug face the threat of being left off the formulary. The design of Medicare Part D, however, removes that threat for various drugs. By law the government requires coverage of "all or substantially all" drugs within various "protected" classes of drugs and has created a pathway for identifying and protecting additional medications (see the Medicare Improvements for Patients and Providers Act, 2008).

Moreover, the mechanism we describe depends critically upon consumers choosing PDP's based on formulary comprehensiveness. PDP's then assemble formularies to maximize expected health benefit. In contrast, the current implementation of Part D makes it difficult for beneficiaries to compare comprehensiveness. To assist seniors in choosing the right drug plans, the government-run online Medicare Prescription Drug Plan Finder ranks plans based only on information about drugs that beneficiaries are currently taking (or know they are likely to take). The government provides no information about formulary comprehensiveness.[25]

The mechanism we describe highlights the importance of formulary comprehensiveness in determining efficiency, suggesting a novel way to measure a drug's market penetration. The current measure predominantly used by the pharmaceutical industry is the percentage of people diagnosed with the relevant condition who take that drug. Our analysis suggests another meaningful measure of market penetration, namely formulary penetration: the percentage of formularies that include that drug, which thus reflects the percentage of people who have access to that drug.

The analysis in this paper applies principles of economic theory to prescription drug markets using a highly stylized model. As we have emphasized, this paper is intended primarily as a thought exercise rather than a blueprint for market design. Our main objective is to offer a conceptual framework for efficient design of government-funded drug procurement.

[25] Even obtaining the list of drugs covered by a particular PDP is difficult.

A Appendix

A.1 No Generic Pure Strategy Equilibrium under Sealed Bid Auction

Lemma 5. *A pure strategy equilibrium in the sealed-bid auction does not generically exist.*

Proof. Suppose there exists a pure strategy equilibrium in which each drug company submits a bid of b_n. One potential outcome is that all drugs are fully covered on the formulary. In this case, any drug company can profitably deviate by changing its bid to $b^{max}(\geq B)$ so that its drug becomes partially covered and collects not only a portion of the budget B but also a portion of monopoly profits on the open market. So there is no pure strategy equilibrium in this case.

The other potential outcome is that some drugs are fully covered on the formulary, and some drugs are not. That implies that some drug companies are submitting normalized bids strictly greater than 1 and getting on the formulary. Hence, any drug company whose drug is not fully covered on the formulary can profitably deviate by changing its bid to just undercut the normalized bid of an on-formulary drug and thus become fully covered on the formulary at a normalized bid level greater than 1.

Hence, no generic pure strategy equilibrium exists. □

A.2 Proofs of Proposition 2 and 3

In the following proofs, we label bidders in the order that they bid. So the first bidder to bid is bidder 1, the second bidder to bid is bidder 2, and so on so that the last bidder to bid is bidder N. In the ascending clock auction, bidder N will turn out to be the highest value bidder (i.e. the bidder with the most valuable drug); meanwhile in the descending clock auction, bidder N will turn out to be the lowest value bidder (i.e. the bidder with the least valuable drug).

When we refer to bidders in terms of their drugs' values, we use parentheses. For example, the reserve value for the most valuable drug is $r_{(1)}$ and the reserve value for the least valuable drug is $r_{(N)}$.

Recall that we have set up our model so that there are no ties. First, $r_i \neq r_j$ for all $i \neq j$ so that no two bidders have both the same bid and the same normalized bid. Secondly, in the rules of the sequential-move auction formats, we have indicated that ties are broken in favor of later bidders – so if two drug companies submit bids at the same normalized level, the later bidder has priority for inclusion on the formulary.

Lemma 6. *In a clock auction (ascending or descending) with $B \geq \sum_{n=1}^{N} r_{(n)}$, there does not exist an equilibrium in which all bidders bid $e = 1$.*

Proof. Suppose all bidders are bidding $e = 1$. Consider the last bidder to bid. This bidder has incentive to deviate and bid as high as possible, which allows it to collect both the excess budget *and* a positive fraction of monopoly profits. Notice that if a bidder bids at a normalized level higher than $e = 1$ and receives any of the budget, then that bidder receives a total payoff strictly higher than monopoly profits (i.e. the bidder's reserve fee).

To be more specific, in the ascending clock auction the last bidder can deviate to b^{max} and take excess budget $B - \sum_{n=1}^{N-1} r_n = \epsilon$ $(\geq r_N)$ in addition to a positive fraction $(1 - \frac{\epsilon}{b^{max}})$ of monopoly profits. Similarly, in the descending clock auction the last bidder can deviate to the second last bidder's bid and take excess budget $B - \sum_{n=1}^{N-1} r_n$ $(\geq r_N)$, in addition to a positive fraction of monopoly profits. Thus all bidders bidding $e = 1$ is not an equilibrium. $\square$

Lemma 7. *In a clock auction (ascending or descending) with $B \geq \sum_{n=1}^{N} r_{(n)}$, there does not exist an equilibrium in which some bidder receives none of the budget.*

Proof. Since $B \geq \sum_{n=1}^{N} r_{(n)}$, there always exist a small enough $\epsilon > 0$ such that if drug company i bids ϵ above its reserve fee r_i, then in any equilibrium drug i will at least be partially covered on the formulary. In the worst case, all other drug companies bid at normalized levels slightly lower than that of drug company i and excess budget $B - \sum_{n \neq i} r_n -$

$\epsilon' > 0$ remains for drug company i to receive at least partial coverage. Since any drug company i always has the option of bidding arbitrarily close to its reserve fee in order to gain partial coverage and receive a payoff strictly higher than its reserve fee, there does not exist an equilibrium in which some drug company receives none of the budget.

□

Lemma 8. *In any equilibrium of a clock auction (ascending or descending) with* $B \geq \sum_{n=1}^{N} r_{(n)}$, *exactly one bidder gets less than full coverage.*

Proof. First we show that at most one bidder gets less than full coverage. Suppose there is an equilibrium in which at least two bidders are less than fully covered. Lemma 7 tells us that these bidders must be receiving some of the budget and hence they must be partially covered. However, at most one bidder can be partially covered since there are no ties.

Now we show that at least one bidder gets less than full coverage. Suppose there is an equilibrium in which all bidders are fully covered. To be an equilibrium, the entire budget must be depleted. The last bidder can achieve a higher payoff by deviating and increasing its bid, thereby becoming partially covered, collecting the remaining portion of the budget and receiving extra payoff from the open market. □

Lemma 9. *The unique equilibrium of the ascending clock auction involves all bidders bidding at the same normalized level* e^* *except for the bidder with the highest value drug, who bids* b^{max}. *That last bidder is partially covered on the formulary while all the other bidders are fully covered. All bidders receive normalized profits of* e^*.

Proof. First, we characterize the normalized bid $e \geq 1$ such that if the first $N-1$ bidders all bid at normalized level e, then the optimal bid for bidder N is to bid b^{max} and be partially covered on the formulary. The only other bid worth considering is for bidder N to bid e and be fully covered. Bidding b^{max} is preferred so long as

$$er_N \leq (B - e\sum_{n \neq N} r_n) + \left[1 - \frac{B - e\sum_{n \neq N} r_n}{b^{max}}\right] r_N$$

$$\implies e \leq \frac{B + r_N(1 - \frac{B}{b^{max}})}{\sum_n r_n - \sum_{n \neq N} r_n \frac{r_n}{b^{max}}} = e^*. \tag{A.1}$$

Recall that $r_N = r_{(1)}$. Notice that e^* is the (maximum) value of e that exactly satisfies the above inequality. If the first $N-1$ bidders all bid at e^* or lower, then the optimal bid for bidder N is indeed to bid b^{max}, which leads to expected normalized profits of e^*.

Now consider the first $N-1$ bidders. A key observation is the following: the maximal payoff a bidder $j \in \{1, 2, \ldots, N-1\}$ can hope to receive from bidding $e > e^*$ and being partially covered is strictly bounded from above by $(B - e^* \sum_{n \neq j} r_n) + \left[1 - \frac{B - e^* \sum_{n \neq j} r_n}{b^{max}}\right] r_j$. That payoff is never higher than the payoff from being fully covered on the formulary at e^*:

$$e^* r_j \geq (B - e^* \sum_{n \neq j} r_n) + \left[1 - \frac{B - e^* \sum_{n \neq j} r_n}{b^{max}}\right] r_j$$

$$\implies e^* \geq \frac{B + r_j(1 - \frac{B}{b^{max}})}{\sum_n r_n - \sum_{n \neq j} r_n \frac{r_j}{b^{max}}}.$$

Plugging in (A.1) for e^*, we get

$$\implies \frac{B + r_N(1 - \frac{B}{b^{max}})}{\sum_n r_n - \sum_{n \neq N} r_n \frac{r_N}{b^{max}}} \geq \frac{B + r_j(1 - \frac{B}{b^{max}})}{\sum_n r_n - \sum_{n \neq j} r_n \frac{r_j}{b^{max}}}$$

which is always true since the LHS numerator is larger than the RHS numerator, while the LHS denominator is smaller than the RHS denominator. Equality holds if and only if $j = N$. Therefore, if any bidder $j \in \{1, 2, \ldots, N-1\}$ has the opportunity to be fully covered at a normalized bid of $\bar{e} \geq e^*$, then bidding $\bar{e}$ is always preferred to bidding at some $e > \bar{e}$ and being partially covered.

We now use backwards induction to prove that it is a subgame perfect equilibrium for each bidder $j \in \{1, 2, \ldots, N-1\}$ to bid exactly e^*. Consider bidder $N-1$, the second last bidder to bid. Suppose that all previous bidders have bid e^*. If bidder $N-1$ bids e^*, then bidder N will bid b^{max} and hence bidder $N-1$ will be fully covered on the formulary at e^*. On the other hand, if bidder $N-1$ bids anything higher than e^*, then bidder N will match that bid and bidder $N-1$ will be partially covered. Hence, bidder $N-1$ prefers to bid e^* and be fully covered. Suppose instead that some previous bidder has bid $e > e^*$. Then both bidder $N-1$ and bidder N will choose to match e and be fully covered on the formulary (leaving the previous bidder that has bid $e > e^*$ to be partially covered).

Now, consider bidder $N-2$. An identical argument holds. Suppose that all previous bidders have bid e^*. If bidder $N-2$ bids e^*, then we know from above that bidder $N-2$ will be fully covered on the formulary. On the other hand, if bidder $N-2$ bids $e > e^*$, then bidder $N-2$ will be partially covered. Hence, bidder $N-2$ prefers to bid e^* and be fully covered. Suppose instead that some previous bidder has bid $e > e^*$. Then bidder $N-2$ prefers to match e and be fully covered.

Continuing the backwards induction for all bidders $j = N-3, N-4, \ldots, 1$, we find that bidder 1 prefers to bid e^* and be fully covered rather than bid some $e > e^*$ and be partially covered. Hence, we conclude that it is a unique subgame perfect equilibrium for all bidders $1, 2, \ldots, N-1$ to bid e^* and for bidder N to bid b^{max}.

□

Lemma 10. *The unique equilibrium of the descending clock auction involves all bidders but one being fully covered and one bidder being partially covered.*

Proof. From Lemma 8 we know that in any pure strategy equilibrium, exactly one bidder gets less than full coverage. Our first step is to show that this bidder cannot be completely uncovered. One way a bidder can receive no coverage is by bidding $b^{max}(\geq B)$. In the descending clock auction, no bidder bids b^{max} in equilibrium because doing so requires that bidder to move first and effectively eliminate itself from the auction completely. The only other way in which some bidder j can get no coverage is if it bids at a finite level and all the other drugs bid such that the budget is exactly depleted and bidder j has the highest normalized bid. But in this case, since $B \geq \sum_{n=1}^{N} r_{(n)}$, bidder j has a profitable deviation to a lower normalized bid near 1 that guarantees a spot on the formulary. Hence in any pure strategy equilibrium, exactly one bidder gets partial coverage.

Our second step is to show that in any equilibrium, a later bidder never bids at a normalized level below that of an earlier bidder. This is obvious if the earlier bidder is fully covered on the formulary in equilibrium since then the later bidder can match that bid and therefore be fully covered on the formulary as well. If instead the earlier bidder is not fully covered on the formulary in equilibrium (and hence is the only bidder not fully covered), then it is optimal for the later bidder to match the earlier bidder and hence obtain full coverage on the formulary.

Our third step is to prove that any equilibrium takes the following form: for some $j \in \{2, \ldots, N\}$, (a) bidders 1 through $j-1$ bid at some normalized level $e_1 \geq 1$ and are fully covered, (b) bidder j bids at some higher normalized level $e_2 > e_1$ and is partially covered, and (c) all remaining bidders $j+1$ through N bid at normalized level e_2 and are fully covered. To prove that, we show that in any equilibrium, for any two bidders k and $k+1$, if both bidders are fully covered on the formulary, then both bidders must be bidding at the same normalized level. Suppose not and that bidder k is bidding at a lower normalized level than bidder $k+1$. Then we claim that bidder k has a profitable upward deviation, at least by ϵ, that keeps bidder k still fully covered on the formulary. The only concern when bidder k deviates upward is if bidder $k+1$ now revises its bid downward to undercut bidder k. Even if bidder $k+1$ does so, however, we know there is some other bidder j that, in the original equilibrium, is partially covered and bidding higher than what bidder $k+1$ was originally bidding. That bidder j will still prefer to be partially covered because (a) either $j < k$ so bidder j has already bid and the conclusion is trivial, or (b) $j > k+1$ and bidder j faces a larger excess budget than before.

Our final step is to show that there exists an equilibrium and it is unique. That is, there exists an equilibrium of the form described above with a unique identity of bidder j that is partially covered and unique normalized bids e_1 and e_2.

Suppose bidder 1 has chosen some normalized bid e_1. We assume e_1 is appropriate in the sense that it leads to subgame behavior that resembles the form of the equilibrium described above. In particular, e_1 is not too high, so that bidder 1 will be on the formulary[26]; and e_1 is not too low, so that any bidder $j > 1$ bidding optimally at a normalized level higher than e_1 will always be matched by all remaining bidders that follow bidder j.[27] Instead, e_1 satisfies the constraint that at least one later bidder will have incentive to bid at a normalized level higher than e_1, and that later bidder will be the one that is partially covered. By continuity, we know such an e_1 must exist.

The identity j of the partially covered bidder is uniquely determined in the subgame as the first bidder willing to bid at some $e_2 > e_1$. By our constraints on e_1, this then leads to all remaining bidders bidding e_2 and being fully covered. We know that such a bidder j must exist because in the extreme, if bidders 1 through $N-1$ bid at the same level (and are all on getting on the formulary), then bidder N will prefer to bid some $e_2 > e_1$ and take the remaining budget in addition to partial coverage.

Conditional on e_1, bidder j chooses its normalized bid $e_2 > e_1$ as follows. Let $B_j = B - e_1 \sum_{n=1}^{j-1} r_n - e_2 \sum_{n=j+1}^{N} r_n$. Then bidder j chooses e_2 to solve

$$\max_{e_2 \geq e_1} B_j + \left(1 - \frac{B_j}{e_2 r_j}\right) r_j$$

subject to the following three constraints:

1. All remaining bidders prefer to bid e_2 (so that bidder j ends up being partially covered).

2. Bidder j cannot bid higher than the bid of the bidder just before, so for $r_{j-1} > r_j$,

$$e_2 \leq e_1 \frac{r_{j-1}}{r_j}$$

3. Bidding at $e_2 > e_1$ must lead to a higher payoff for bidder j than bidding at e_1 and

[26] The minimum constraint is that there exists some j such that $e_1 < \frac{B}{\sum_{n \neq j} r_n}$.

[27] If e_1 is too low, then it can happen that some bidder $j > 1$ bids higher than e_1 and some even later bidder $k > j$ bids even higher than bidder j. Clearly, this is suboptimal for bidder 1.

getting fully covered on the formulary:

$$e_1 r_j \leq B_j + (1 - \frac{B_j}{e_2 r_j}) r_j.$$

Since e_2 can be arbitrarily close to e_1 and B is sufficiently large, we know that given e_1, a solution e_2 exists and, by monotonicity, is unique.

Finally, we show that bidder 1's choice of e_1 is unique. This follows immediately from bidder 1's payoff $e_1 r_1$ being strictly increasing in e_1 because bidder 1 simply chooses the highest e_1 possible, subject to the constraints detailed above. That unique e_1 makes bidder j just indifferent between bidding e_1 (full coverage) and bidding e_2 (partial coverage).

□

A.3 Comparison between Ascending and Descending Clock Auctions

Lemma 11. *As b^{max} increases, the social value of the formulary obtained in the unique equilibrium of the ascending clock auction converges towards $\sum_{n=1}^{N} z_n - \max_n z_n$. Meanwhile, as b^{max} increases, the social value of the formulary obtained in the unique equilibrium of the descending clock auction remains unchanged and is strictly greater than $\sum_{n=1}^{N} z_n - \max_n z_n$.*

Therefore, for any set of drug values $\{z_n\}_n$, there exists a b^* such that if $b^{max} \geq b^*$, then the formulary obtained in the unique equilibrium of the descending clock auction has a social value strictly greater than that of the ascending clock auction.

Proof. In the unique equilibrium of the ascending clock auction, we know that all drugs are fully covered on the formulary except the highest value drug, which is the last to bid. From the proof of Lemma 9, we know that the highest value drug always bids b^{max} and has a fraction x covered on the formulary, where the numerator of x equals the remaining budget (which does not increase nearly as fast as b^{max}) and the denominator equals b^{max}. Hence, as

b^{max} increases, the fraction x converges to 0 and the social value of the formulary converges towards $\sum_{n=1}^{N} z_n - \max_n z_n$.

In the unique equilibrium of the descending clock auction, we know that all drugs are fully covered on the formulary except for one drug, which is partially covered and which is *not* the highest value drug. From the proof of Lemma 10, we know that the equilibrium does not depend on b^{max} and hence as b^{max} increases, the social value of the formulary does not change. □

References

Boldrin, Michele & David Levine. 2006. "Growth and Intellectual Property." NBER Working Paper Series.

Cornelli, Francesca & Mark Schankerman. 1999. "Patent Renewals and R&D Incentives." *Rand Journal of Economics* 30:197–213.

Duggan, Mark & Fiona Scott Morton. 2006. "The Distortionary Effects of Government Procurement: Evidence for Medicaid Prescription Drug Purchasing." *Quarterly Journal of Economics* 121:1–30.

Fong, Kyna & Michael Schwarz. 2009. "Per-Person Pricing and Alternative Risk Adjustment Schemes in Prescription Drug Procurement.".

Garber, Alan, Charles Jones & Paul Romer. 2006. "Insurance and Incentives for Medical Innovation." *Forum for Health Economics and Policy* 9(2):4.

Hopenhayn, Hugo, Gerard Llobet & Matthew Mitchell. 2006. "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts." *Journal of Political Economy* 114(6):1041–1068.

Kremer, Michael. 1998. "Patent Buyouts: A Mechanism for Encouraging Innovation." *The Quarterly Journal of Economics* 113(4):1137–1167.

Lakdawalla, Darius & Neeraj Sood. 2006. "Health Insurance as a Two-Part Pricing Contract." NBER Working Paper 12681.

McAdams, David & Michael Schwarz. 2007. "Perverse Incentives in the Medicare Prescription Drug Benefit." *Inquiry* . Forthcoming.

McAfee, Preston. 1992. "A Dominant Strategy Double Auction." *Journal of Economic Theory* 56(2):434–450.

McClellan, Mark, I. Spatz & S. Carney. 2000. "Designing A Medicare Prescription Drug Benefit: Issues, Obstacles, and Opportunities." *Health Affairs* 19(2):26–41.

Morton, Fiona Scott. 1997. "The Strategic Response by Pharmaceutical Firms to the Medicaid Most-Favored-Customer Rules." *The RAND Journal of Economics* 28(2):269–290.

Newhouse, Joseph. 2004. "How Should Medicare Pay For Drugs?" *Health Affairs* 21(1):89–102.

Schwarz, Michael. 2006. "Do Tax Subsidies and Medicaid Increase Access to Essential Medications?".

Scotchmer, Suzanne. 1999. "On the Optimality of the Patent Renewal System." *Rand Journal of Economics* 30:181–196.

Viscusi, W. Kip. 1993. "The Value of Risks to Life and Health." *Journal of Economic Literature* 31(4):1912–1946.

Wright, Brian D. 1983. "The Economics of Invention Incentives: Patents, Prizes, and Research Contracts." *American Economic Review* 73:691–707.