NBER WORKING PAPER SERIES

SAVING LIVES OR SAVING MONEY? UNDERSTANDING THE DUAL NATURE OF PHYSICIAN PREFERENCES

Alice Chen
Darius N. Lakdawalla

Working Paper 21930
http://www.nber.org/papers/w21930

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2016

We gratefully acknowledge comments from Joshua Gottlieb and other participants at the American Economic Association Annual Meeting, the BU-Harvard-MIT Health Seminar, the NBER Health Care Summer Institute, the Midwest Health Economics Seminar, and the USC CESR-Schaeffer Brown Bag. Research reported here was supported by the National Institute on Aging of the National Institutes of Health under award number P01AG033559 and P30AG024968. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or the National Bureau of Economic Research

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w21930.ack

Saving Lives or Saving Money? Understanding the Dual Nature of Physician Preferences
Alice Chen and Darius N. Lakdawalla
NBER Working Paper No. 21930
January 2016
JEL No. I11,I12,I18

ABSTRACT

A longstanding literature highlights the tension between the altruism of physicians and their desire for profit. We develop new implications for how these competing forces determine pricing and utilization. Altruism encourages providers to reduce utilization in response to higher prices, but profit-maximization does the opposite. Rational physicians behave more altruistically when treating poorer patients or those facing higher medical costs, and when foregone profits are lower. These insights help explain observed heterogeneity in pricing dynamics. We show that average price elasticities vary from 0.6 to 1.1 for a given physician, depending on the patient socioeconomic status and out-of-pocket cost burden. This finding has important implications for the design of reimbursement by Medicare Parts A and B, and benefit design within Part D.

Alice Chen
Schaeffer Center for Health Policy and Economics
University of Southern California
635 Downey Way
Los Angeles, CA 90089-3333
alicejc@usc.edu

Darius N. Lakdawalla
University of Southern California
635 Downey Way, VPD 414-K
Schaeffer Center for Health Policy and Economics
Los Angeles, CA 90089-7273
and NBER
dlakdawa@usc.edu

1. Introduction

Economists have long emphasized the peculiarities of healthcare markets, compared to other markets for goods and services. Since at least Kenneth Arrow's pioneering paper on the subject, economists have recognized two features in particular: the altruism of healthcare providers towards their patients and the reliance of patients on their physicians for information and guidance (Arrow 1963). These two insights yield implications for healthcare pricing and utilization that have not yet been explored.

Altruism encourages physicians to faithfully represent their patients' interests. An altruistic physician will tend to economize on the use of scarce inputs and attempt to maximize patient utility subject to her constraints. On the other hand, the informational advantage of physicians creates a classic agency problem that physicians might exploit to pursue their own interests at the expense of their patients' (for example, Dranove and White 1987; Blomqvist, 1991; Emanual and Emanual, 1992; Mooney and Ryan, 1993; Zweifel and Breyer and Kifmann, 1997; and Lu, 1999). These countervailing incentives produce conflicting implications for pricing. Self-interested physicians will respond to higher price by performing more procedures. Altruistic physicians, on the other hand, will protect their patients from higher prices by performing fewer procedures. Thus, a fuller theory of healthcare pricing must present a unified framework for analyzing how altruism and agency problems interact.

[Insert Figure 1 Here]

Figure 1 provides preliminary insight into the importance of this issue. The figure depicts the histogram of price elasticities within Medicare services that experienced an approximately 50 percent increase in annual physician reimbursement rates. Such large and sudden reimbursement changes are unlikely to reflect changes in demand but are instead more likely to reflect administratively determined changes in the prices faced by providers and patients. Price increases coincided with increased quantity in half the services depicted, but with decreased quantity in the other half.[1] This pattern is difficult to explain when relying on either agency problems or altruism alone. Significant policy questions are at stake, since price is often viewed as an important lever for influencing behavior.

[1] The large majority of the procedures depicted in Figure 1 are major or minor procedures. Very few are lab test, imaging, or evaluation and management services.

In this paper, we develop a unified framework that incorporates a spectrum of physician preferences, ranging from pure altruism to pure profit-maximization. We then show how different preferences generate different price responses in the marketplace. In this framework, the physician's degree of altruistic behavior is not simply exogenous. Rather, it varies systematically with patient characteristics. Finally, we show empirically that patient characteristics influence price responses, even for a single physician with fixed preferences.

We rely on well-established models of physician behavior but apply these to problems of pricing and utilization that have not been viewed through this lens before. From a positive perspective, we show that exogenous price increases may raise or lower quantity supplied. When higher prices lower quantity, we say dynamics are primarily "patient-driven," and when the opposite is true, we say they are primarily "physician-driven." Pricing is more likely to be patient-driven (i.e., higher prices are more likely to results in lower quantities) when patients are poorer and when healthcare provision is less profitable. Intuitively, physician altruism is more likely to win out when the value of behaving altruistically is higher and the cost is lower.

From a normative perspective, patient-driven behavior limits the potential for overuse of healthcare resources, while physician-driven behavior exacerbates it. Thus, we expect less overuse when consumers are poorer, patient cost-sharing is higher, input prices are lower, and profitability is lower. Increases in patient wealth, therefore, are expected to increase "waste" in healthcare, as are expansions in the availability of insurance.

Empirically, we test the conjectures of our model by using two exogenous policy shocks to Medicare payments: the 1997 consolidation of geographic payment regions and the 1999 change in estimation of practice expenses. Our results indicate that the size and sign of the own-price elasticity varies with patient characteristics as predicted. Procedures are more likely to follow patient-driven pricing behavior when patient income is lower and patient cost-sharing is higher. The variation in price elasticities matters quantitatively. Within physicians, average price elasticities vary from 0.6 to as high as 1.1, when treating patients in the bottom or top socioeconomic status deciles, respectively.

Our work unifies theories of physician altruism and agency problems into a single framework for healthcare price theory. The literature thus far has offered piecemeal explanations of the observed heterogeneity in response to price changes. Some empirical studies observe that higher reimbursements will lead to increased utilization, and the accompanying theory relies on physicians being profit maximizers (Clemens and Gottlieb, 2003; Gruber et al., 1999; and Jacobson et al., 2006). Other empirical studies show that there is a negative relationship between price and quantity (Rice, 1983; Escarce, 1993; Nguyen and Derrick, 1997; and Yip, 1998). Theories used to explain a negative price-quantity relationship include models of physician induced demand and non-fee-for-service reimbursement schemes. For example, Farley (1986) discusses implications of the target-income model. Ellis and McGuire (1986) demonstrate that a prospective-payment system can lead to under-provision of care if physicians undervalue the benefits of patients relative to hospital profits. Choné and Ma (2011) and Glied and Zivin (2002) discuss how managed care can restrict quantity. Finally, some studies find a low responsiveness between quantity and price, and they infer that there is uncertainty in a physician's objective function (Holohan, 1977; Hurley et al., 1990; and Hurley and Labelle, 1995).

We unify these findings by offering a simple modification to the existing theory. In contrast to several important prior studies, such as Ellis and McGuire (1986), Ellis and McGuire (1990), and Liu and Ma (2013), our model allows physicians to care about patient health *and* patient spending. This latter mechanism generates new insights into predicting when services are likely to be patient-driven versus physician-driven. Our theory highlights how and why certain markets will respond better to demand-side policies and others to supply-side policies in controlling costs.[2] This work relates to Dickstein (2014), who empirically quantifies the contributions of patient and physician incentives to prescription drug utilization.

Our findings have several notable policy implications. First, patient and physician characteristics have systematic and predictable effects on healthcare price elasticities. Second, changes in physician reimbursement rules – for instance, in a public health insurance system – will exhibit systematically different effects across different types of patients and procedures. If policymakers know these effects, they can better target reimbursement reforms within Medicare Parts A and B and benefit-design reforms within Part D, thus considerably strengthening their impacts. Third, relying simply on aggregated estimates of price elasticities may lead to interventions with unintended consequences for certain patients, procedures, or markets. And finally, although price increases are associated with increased utilization on average, overuse will

[2] While policymakers have traditionally focused on controlling Medicare expenditures by altering Medicare payments, demand-side policies, such as changes to patient cost-sharing and supplemental insurance, have been debated recently (e.g., Gruber, 2013; Zuckerman et al., 2010).

be tempered when patient cost-sharing is high or patient income is low: the risk of overtreatment and overmedication fall among poorer patients who face higher out-of-pocket costs.

The rest of the paper is organized as follows. In Section 2, we propose a theoretical framework that unifies physicians' altruism with their role as agents, and we derive the normative implications of our model. In Section 3, we discuss the empirical approach for testing conjectures derived from our model. In Section 4, we present the empirical results, and Section 5 concludes.

2. Theoretical Framework

Physician altruism and agency create unique relationships among pricing, utilization, and other economic forces. We demonstrate these points in a simple theoretical model that traces back to Becker (1957). The model has been used by many health economists to study physician behavior (for example, Ellis and McGuire 1986; Ellis and McGuire 1990; McGuire and Pauly, 1991; and McGuire 2000).

2.1. Simple Illustration

For pedagogical purposes, we begin by illustrating the intuition using a stylized perfectly competitive model without insurance. We assume that physicians earn zero economic profits, and patients bear the full cost of healthcare.

Suppose health is produced using a good or procedure X, according to $F(X)$, where $F_{XX} < 0$. This good is initially health-improving but eventually health-reducing if overused. Imagine first that a fully informed representative patient maximizes the value of health net of the cost of production. This results in the following household production function for health:

$$\max_X vF(X) - p_X X.$$

It is straightforward to show that the derived demand for X is falling in price p_X, as in $\frac{\partial x}{\partial p_x} = \frac{1}{vF_{xx}} < 0$.

Now, however, suppose that the representative patient is not fully informed but instead receives care from a fully informed physician, who bears cost $c(X)$, where $c_{XX} \geq 0$. The physician maximizes a weighted average of patient well-being and physician income, as in:

$$\max_X (1-\alpha)[p_X X - c(X)] + \alpha[vF(X) - p_X X]$$

The parameter α is an index of altruism. Observe in this framework that the physician's objective function can be rewritten as:

$$\max_X \alpha v F(X) - (1-\alpha)c(X) + (1-2\alpha)[p_X X]$$

This has the following first-order condition:

$$\alpha v F_X - (1-\alpha)c_X + (1-2\alpha)p_X = 0$$

Define D as the second derivative for this maximization problem. This allows us to write the comparative static of the problem as:

$$\frac{\partial X}{\partial p_X} = \frac{(2\alpha - 1)}{D}$$

If the problem is strictly concave at the optimum, then $D < 0$. As a result, if $\alpha > \frac{1}{2}$, the own-price elasticity is negative, because the physician is sufficiently altruistic that her decision problem resembles that of the fully informed patient. We call these "patient-driven pricing dynamics." If, on the other hand, $\alpha < \frac{1}{2}$, the opposite dynamics prevail, and the own-price elasticity is positive. We call these "physician-driven pricing dynamics."

2.2. General Model

The derivation above assumed physicians and patients are risk-neutral over consumption, and it abstracted from health insurance. To generalize it, suppose the representative patient derives utility $u(I + vF(X) - \pi - \sigma(X; p_X))$, where I is income, π is an ex ante insurance premium, and σ represents patient out-of-pocket expenditures. For the bulk of the analysis, we assume that patients face at least some cost-sharing in the sense that $\sigma_X > 0$, and in the sense that they bear costs when prices rise (i.e., $\sigma_{X p_X} > 0$). Later, we discuss the polar case of zero cost-sharing. Here and elsewhere, we abstract from effects of physician decisions on the insurance premium. This assumption sacrifices little generality in a public insurance scheme or when studying a relatively focused set of procedures.

Now suppose physicians derive utility from a weighted average of patient utility and their own utility over consumption, $z(\cdot)$, where u and z are weakly concave utility functions. Physicians may also earn some non-labor income $N \geq 0$. Assume the physician utility function satisfies the assumptions of monotonicity, risk-aversion, and weak prudence, as in $z' > 0$, $z'' < 0$, and $z''' \geq$

0 (Felder & Mayrhofer, 2011). The generalized physician objective function can then be written as:

(1) $$\max_X (1-\alpha)z(N + p_X X - c(X)) + \alpha u(I + vF(X) - \pi - \sigma(X; p_X))$$

The first-order conditions now become:

$$\alpha u' * (vF_X - \sigma_X) + (1-\alpha)z' * (p_X - c_X) = 0$$

The optimality conditions are weighted averages of conditions for physician profit-maximization and patient utility-maximization.

To simplify the analysis, we follow the convention adopted in much of the insurance literature and abstract from the direct income effects associated with patient out-of-pocket payments (Lakdawalla & Sood, 2013). This amounts to holding u' fixed when prices change. The comparative static now becomes:[3]

(2) $$\frac{\partial X}{\partial P_X} = \frac{\left(\alpha u' \sigma_{X p_X} - (1-\alpha)z' - (1-\alpha)z''(p_X - c_X)X\right)}{D}.$$

In the absence of altruism, physicians respond to price increases by providing more care. To see this, observe from the physician's optimality condition that $p_X = c_X$ in the absence of physician altruism. Therefore, without altruism, it follows that $\frac{\partial X}{\partial P_X} = -\frac{(1-\alpha)z'}{D} > 0$.

Notice that the value of health, v, does not alter the price response function, except through its effect on the marginal utility of consumption. This follows from our assumption that patient out-of-pocket prices do not produce any income effects, and the logic is robust to more general utility functions that take consumption and health as two separate arguments. For a given level of output, patient characteristics affect price-responsiveness through the marginal utility of consumption alone. This is helpful empirically, because it specifies a single mechanism through which patient characteristics and socioeconomic status affect price-responsiveness. This sharp result would fail if we allowed for income effects associated with out-of-pocket costs or if we allowed for complex effects of patient characteristics on the second-order condition.

To develop further the implications of the comparative statics, we investigate how changes in exogenous parameters influence the likelihood of patient-driven pricing dynamics. We impose one final assumption, namely that the patient's own private marginal benefit exceeds her marginal out-of-pocket cost, or $vF_X \geq \sigma_X$ at the optimum. The asymmetry of information means this is not

[3] The determinant D is equal to $\alpha u''(vF_X - \sigma_X)^2 + \alpha u'(vF_{XX}) + (1-\alpha)z''(p_X - c_X)^2 + (1-\alpha)z'(-c_{XX})$.

a trivial assumption, but—at least for insured consumers— it would be violated only in extreme cases of overuse. This assumption, coupled with the first-order condition for X, implies that $p_X \leq c_X$.

Pricing dynamics are patient-driven if and only if $\alpha u' \sigma_{Xp_X} > (1-\alpha)z' + (1-\alpha)z''(p_X - c_X)X$. Thus, for a given level of output, pricing dynamics are more likely to be patient-driven if:

1. Physician altruism is higher – i.e., α is higher;
2. Physician non-labor income is higher – i.e., N is higher, which implies that z' and $z''(p_X - c_X)$ are both lower;
3. Patient income is lower – i.e., I is lower and thus u' higher;
4. Patient out-of-pocket spending is higher – i.e., $\pi + \sigma$ is higher and u' higher;
5. The value of health is lower – i.e., v is lower and u' higher;
6. The physician's price-cost margin, $p_X - c_X$, is lower.

Intuitively, pricing is more likely to be patient-driven if: physicians care more about their patients (#1); physicians are richer and thus willing to pay more to purchase patient welfare (#2); patient welfare is more sensitive to out-of-pocket spending growth (#3 and #4); patients derive less value from their health capital (#5); and the opportunity cost to physicians of boosting utilization is lower (#6). Finally, note that in a zero cost-sharing environment, it is true that $\sigma_{Xp_X} = 0$, which implies that pricing is always physician-driven. Intuitively, even altruistic physicians have no incentive to worry about patients' financial situation when there is no cost-sharing. This result also demonstrates that even altruistic physicians may respond to price increases by performing more procedures, if the effects of price increases on patients are minimal. Nonetheless, this case is more pedagogical than practical. The "zero cost-sharing" assumption assumes away any entity (e.g., a third-party payer) that bears the costs of a price increase and influences physician decisions in response.

Variation in any of the above parameters can lead price elasticities to vary widely and even change sign. This result highlights that negative price elasticities could sometimes result from physician altruism alone, without the presence of strong physician income effects and backward-bending labor supply curves. Moreover, in contrast to models of backward-bending labor supply, our theory uniquely predicts that patient characteristics can affect price-responsiveness at a given margin. That is, the same physician on the same day will respond to the same price changes differently for poor patients compared to rich patients.

The theory also has implications for changes in price-responsiveness, $\frac{\partial X}{\partial P_X}$, albeit subject to some caveats. Equation (2) implies that the numerator of this term varies identically in response to the six parameters listed above. For example, $\frac{\partial X}{\partial P_X}$ is likely to be lower for poorer patients with lower income (I) and value of health (v), and for patients with higher cost-sharing $(\pi + \sigma)$. In principle, changes in the second-order condition might undo some or all of these effects, so the impact of patient characteristics on price-responsiveness becomes an empirical matter.

Finally, it is worth noting that our price elasticities do not reflect combinations of traditional demand and supply elasticities. Rather, they reflect the elasticity of supply in a market where producers are acting as agents for patients. More formally, the elasticity implied by Equation (2) can be written as a weighted sum of the elasticities derived from separate maximizations of the profit- and patient-utility components of the physician objective function. If physicians maximized only their own utility of consumption $\mathrm{z}(.)$, the comparative static from the physician's objective function will yield a price elasticity, which we denote ϵ^{C}. Alternatively, if physicians maximized only patient utility $\mathrm{u}(.)$, the comparative static would yield a different price elasticity, which we denote ϵ^{H}. When physician maximize both own- and patient-utility, the elasticity is:

$$\frac{\partial \mathrm{X}}{\partial \mathrm{P_X}} \frac{\mathrm{P_X}}{\mathrm{X}} = \frac{\alpha \mathrm{A}}{\alpha \mathrm{A} + (1-\alpha)\mathrm{B}} \epsilon^{\mathrm{C}} + \frac{(1-\alpha)\mathrm{B}}{\alpha \mathrm{A} + (1-\alpha)\mathrm{B}} \epsilon^{\mathrm{H}},$$

where A = $[\mathrm{u}' \mathrm{v} \mathrm{F_{XX}} + \mathrm{u}''(\mathrm{v}\mathrm{F_X} - \sigma_{\mathrm{X}})^2]\mathrm{X}$ and B = $[\mathrm{z}'(-\mathrm{c_{XX}}) + \mathrm{z}'(\mathrm{p_X} - \mathrm{c_X})^2]\mathrm{X}$.[4] Note that A<0 and B<0. Thus, the implications that we have derived regarding positive versus negative elasticities can be interpreted as identifying when the elasticity of physician consumption dominates the elasticity of patient utility.

2.3. Welfare Implications

Dual physician preferences provide implications for welfare as well. The degree of inefficient input overuse depends on moral hazard and on the over- (or under-) reimbursement of

[4] The elasticity from maximizing physician consumption and patient health are $\epsilon^H = \frac{P_X u' \sigma_{X P_X}}{A}$ and $\epsilon^H = \frac{-p_X[z''(P_X - c_X)X - z']}{B}$, respectively.

physicians. In patient-driven markets, moral hazard is relatively more important to address, while physician reimbursement is more important in physician-driven markets.

To understand these results, observe that Pareto-efficiency requires the standard input efficiency conditions, $vF_X = c_X$. Thus, we can characterize the degree of inefficient overuse by quantifying $c_X - vF_X$. By inspecting the first-order conditions for physician decisionmaking, we can derive:

$$c_X - vF_X = \frac{\alpha u'}{\alpha u' + (1-\alpha)z'} \overbrace{(c_X - \sigma_X)}^{Moral\ hazard} + \frac{(1-\alpha)z'}{\alpha u' + (1-\alpha)z'} \overbrace{(p_X - vF_X)}^{Over-reimbursement}$$

This condition demonstrates that both moral hazard and physician over-reimbursement play a role in input efficiency. The overall degree of input inefficiency is the weighted average of these two sources, with the weights given by the relative importance of patient versus physician consumption. If physicians are perfectly altruistic, the over-reimbursement effect vanishes. On the other hand, if they are perfectly self-interested, the moral hazard effect vanishes. In addition, note that increases in physician consumption levels place more weight on the moral hazard effect, because richer physicians can "afford" to place more value on their patients' consumption than their own.

The relative importance of physician versus patient consumption has implications for which policy levers are most efficient at reducing distortions. If the degree of altruism is high, reimbursement reforms aimed at mitigating moral hazard (measured as the difference, $c_X - \sigma_X$) will be relatively more effective. If low, on the other hand, reforms aimed at physician's over-reimbursement (measured as the difference, $p_X - vF_X$) will be correspondingly more effective. Put heuristically, policymakers should focus on moral hazard in patient-driven markets, but on physician reimbursement in physician-driven markets. More formally, holding all patient and physician incentives constant, reimbursement reforms that compress $(p_X - vF_X)$ will contribute less to efficiency when $\alpha u' > (1-\alpha)z'$, and vice-versa.

Factors that promote physician altruism tend to promote the importance of moral hazard, while factors promoting physician self-interest do the opposite. Thus, increases in patient wealth will tend to increase the importance of aligning physician incentives through payment reform. Our analysis also has implications for reimbursement reforms. Uniform reimbursement changes – either global increases or global decreases in price – may have unintended consequences that depend on the mix of patient-driven versus physician-driven markets or procedures. Targeted reforms that change reimbursement for some markets, but not for others, might be more effective.

3. Empirical Analysis

Our theoretical analysis suggests at least three testable hypotheses:

1. Both the size and even the sign of the price elasticity may vary when physicians balance altruism and self-interest.
2. When patient socioeconomic status is lower, price elasticities may be lower, because the marginal utility of consumption is higher, and physicians behave more altruistically.
3. When patient cost-sharing is higher, price elasticities may also be lower, and physicians will behave more altruistically.

Moreover, the unique feature of our theory is the possibility that the *same* physician may respond differently to price changes, depending solely on the characteristics of the patient being treated. Thus, we wish to assess whether the same physician responds to price changes differently when treating patients of different types. Alternative theories featuring backward-bending supply curves or administratively set prices might produce elasticities of ambiguous sign (Hypothesis #1), but cannot readily explain why the same physician would exhibit different price elasticities simply when treating different types of patients. To be clear, our empirical analysis is not designed to rule out other theories, but rather to "rule in" ours. Multiple theories could be at work in parallel within the marketplace.

3.1. Data

To explore these hypotheses, we rely on data from 1993 to 2002 from the Medicare Current Beneficiary Survey (MCBS). The MCBS data consists of a nationally representative sample of 12,100 Medicare beneficiaries, and it consists of data from both administrative payment files and patient surveys. The administrative files provide information on the fee-for-service Physician/Supplier Part B claims. Each service provided is identified by a Healthcare Common Procedure Coding System (HCPCS) code, and identifiers allow us to track individual physicians over time. We focus on only physicians providing medical services and procedures (i.e., Level 1

HCPCS codes).[5] The patient surveys provide us with a rich set of covariates that allow us to identify patient demographics and socioeconomic status.

The MCBS provides a few key advantages over using a dataset with solely administrative claims. First, it provides information on a patient's Medicaid status, enrollment into Medigap, and (if applicable) the specific Medigap plan letter.[6] With this information, we can refine the claims-based out-of-pocket (OOP) cost variable to better reflect the actual payments borne by patients, as opposed to the amount that a third-party insurer might cover. Second, the MCBS offers detailed information on an individual's socioeconomic status. We consider a few proxies for the patient's marginal utility of consumption: (1) patient income—which includes information on pre-tax wages, Social- and Supplemental-Security Income, pensions, retirement income, interest from mutual funds and stocks, and other sources, (2) Medicaid enrollment, which identifies the elderly with incomes less than 100% to 120% of the Federal Poverty Level, (3) the patient's highest level of educational attainment, and (4) out-of-pocket medical spending.[7]

We continue to maintain the common assumption that patient cost-sharing changes do not generate income effects. Under this assumption, our theory implies that patient income, socioeconomic status, and other characteristics affect price-responsiveness only through the marginal utility of consumption. This avoids many of the typical endogeneity problems associated with patient socioeconomic status, which often simultaneously influence health-related outcomes through a variety of different channels. As a sensitivity analysis, we conduct our estimation on a subsample of patients where this assumption is most likely to hold.

Because our analysis focuses on how patient characteristics affect physician behavior, we collapse the data to the physician-patient-year level. To obtain an accurate measure of prices physicians face over time, we create a price index. Specifically, for each physician, we identify the universe of procedures that she performs throughout our data period. The physician's price index in a given year is the weighted sum of the Medicare allowed charges for that basket of procedures. The weights are constant over time, and they reflect the share of relative value units (RVUs)—

5 We exclude Level II HCPCS codes, which are used primarily to identify supplies and products, such as durable medical equipment. Additionally, we use the specialty code to exclude suppliers and providers in specialties which do not require an MD or DO degree (e.g., optometry, physical therapists, social workers, nurses, etc.).

6 For Qualified Medicare Beneficiaries and patients with Medigap plans C and F, OOP costs equal zero. For patients with Medigap plans A, B, D, G, M, or N, we set the copay equal to only the deductible, as those plans offer full coverage of the part B copay. Finally, those with Medigap plans K (L) pay 50% (75%) of the copay and the full deductible.

7 In 2000, 100% of the Federal Poverty Level for a one-person household was $8,350.

reflecting differences in the time, skill training, and costs required to perform a procedure—associated with a procedure relative to the total RVUs that make up the physician's basket.[8] For each HCPCS, we identify the corresponding RVUs using data from the Federal Registers in 1993 to 2002. Nineteen HCPCS make up the average basket, and eleven make up the median basket.

To measure OOP costs, we would ideally create a price index that reflects the expected OOP cost a patient faces. However, unlike physicians who perform a relatively constant basket of procedures on an annual basis, patient health can be variable and unpredictable, making it difficult to construct a fixed basket of procedures that remains meaningful for the patient over time. For example, a patient's OOP costs can be fairly low until an unexpected but major health condition arises. An alternative approach is to simply utilize actual OOP payments for Medicare Part B services. However, Medicare OOP costs consist of a deductible, equal to $100 per year in 2003, and 20% copay, so using total or average realized payments conflates the out-of-pocket price with the actual quantity of services consumed. To create a proxy for the a priori expected per-unit price accruing to the patient, we calculate the average OOP per RVU in cells defined by each three digit ICD-9 diagnosis code, comorbidity severity as measured by Quan et al.'s (2005) Charlson Comorbidity index (CCI), physician specialty, and geographic area as measured by the Medicare Payment Localities (MPLs). Finally, using these four variables—diagnosis, CCI, specialty, and location—we define a patient's "expected OOP price" as the average OOP within the matched cell..

[Insert Table 1 Here]

Table 1 shows the summary statistics for the MCBS. All variables measured in dollars are converted to 2010$ using the CPI. The mean income in the MCBS data is $24,080, which is roughly comparable to the Social Security Administration's reported median income of $26,600 among individuals 65 and older. There is considerable distribution in both the individual level measures of income. While Medicare co-insurance is 20%, the average expected OOP of $10.32 is far less than 20% of the total price index of $273.96, because we have considered supplemental insurance such as Medigap and Medicaid. The other patient demographics in the data include gender, race, age, and body mass index (BMI).

[8] The RVU-weights come from summing work, practice, and malpractice RVUs associated with a given HCPCS. These RVUs are the average RVU during the 1993 to 2002 time period.

3.2. The Medicare Payment Structure and Relevant Policy Shocks

For each HCPCS, CMS calculates a payment based on three factors: (1) the RVU, (2) a geographic adjustment factor (GAF), and (3) a conversion factor (CF).[9] While RVUs are procedure specific, GAFs are region specific, so they account for geographic variation in the cost of providing services.[10] The CF is a nationally uniform adjustment factor that converts RVUs into a dollar amount. This factor is updated annually by CMS according to a formula specified by statute, but Congress can and has overridden the statutorily defined formula.[11] To measure price elasticities, we need to identify payment changes within a market that are independent of patient demand, technological change, and supply. If we rely on changes to the overall Medicare payment rate, we will capture variation from RVUs, GAFs, and the CFs. Because GAFs are set across several different markets and the CF is one number set nationally, these two components of Medicare pricing are likely exogenous to dynamics that a given physician faces. However, variation in RVUs may not be exogenous within a market over time. At least once every five years, about 138 physicians from the Specialty Society Relative Value Committee (RUC) and its advisory committee convene to re-evaluate the work component of RVUs, which reflects procedure-specific differences in physician time, skill, and training. If adjustments in work RVUs are systematically correlated with demand, then price elasticity estimates based on work RVU variation may be biased.

While changes in work RVUs may be non-random in theory, the practical case for bias is less clear. The assignment of relative weight is complex and political with battle lines and alliances drawn between specialties (Eaton, 2010). Deliberations are complicated by the fact that the size of the Medicare payment pie is fixed. As such, the final weights have been viewed as somewhat arbitrary. For example, after the first major review of RVUs, the Health Care Financing Administration (HCFA) received "voluminous identical comments from family practitioners

[9] The exact formula for calculating Medicare payments is given by: $Pay = [RVU_W GPCI_W + RVU_{PE} GPCI_{PE} + RVU_{MP} GPCI_{MP}] \times CF$, where W indexes the work component, PE indexes the practice expense component, and MP indexes the malpractice expense component. GPCI represents the geographic practice cost indices, and CF is the conversion factor.

[10] GAF is a weighted sum of the work, practice expense, and malpractice GPCIs. Details can be found in MaCurdy et al. (2012).

[11] The CF in 2013 was $36.61 per RVU. Congress overrode this formula in 1998, 2009, and 2011.

stating that [the HCFA had…] used an arbitrary method for revising the work RVUs" (Department of Health and Human Services, 1996).[12]

Nevertheless, we address potential endogeneity by relying on two policy shocks in Medicare pricing. The first major policy shock occurred in 1997 when the Healthcare Financing Administration (HCFA) consolidated the number of geographic payment regions—known as Medicare Payment Localities (MPLs) from 210 distinct MPLs to only 89 distinct MPLs in 1997. Discussed in Clemens and Gottlieb (2014), this consolidation generated differential price shocks across county groupings within a state. While some states were unaffected by this policy, in about 26 states, the variation in reimbursement rates across counties was either significantly reduced or eliminated because multiple regions were collapsed into one single payment area.

In contrast with the 1997 shock that differentially affected geographies, a second major policy shock in 1999 created differential changes across services. Prior to 1999, the practice expense RVU components (PE-RVUs) were measured using prevailing charges. However, Section 121 of the Social Security Amendments of 1994 and the Balanced Budget Act of 1997 mandated two changes to PE-RVUs to be phased in over a four-year period from 1999 to 2002. First, PE-RVUs were to be determined by relative costs, instead of prevailing charges. Second, PE-RVUs were modified to better account for cost differences of performing a procedure in a "facility"—such as a hospital, skilled nursing facility, or ambulatory surgical center—versus a "non-facility," such as an office or clinic.[13]

[Insert Figure 2 Here]

Using data from Federal Register reports, we depict in Figure 3 the variation in the GAF and PE-RVU components of Medicare reimbursements over time. Plot (a) shows the change in GAF among counties that were affected by the 1997 consolidation versus those that were unaffected. Much of the pre-1997 differentiation across counties was eliminated post-1997. Plot (b) illustrates the change in average facility and non-facility PE-RVUs across HCPCS over time. While the transition from charge- to resource-based estimations was phased in over a four-year period, the differentiation between facility and non-facility RVUs occurred immediately and created a sudden drop in average PE-RVUs. As Appendix Figure A.1 depicts, much of the

[12] Between 1993 to 2002, work RVUs experienced two major reviews which became effective in 1997 and 2002. The change in average work RVU is depicted in Appendix Figure A.1.

[13] Prior to 1999, the non-facility PE-RVU was simply 50% if the facility PE-RVU (Maxwell and Zuckerman, 2007).

observed drop in PE-RVUs in 1999 comes from changes in the non-facility estimates. Changes in the other components of Medicare reimbursements are also discussed in Appendix A.

3.3. Empirical Approach

Estimating Physician Elasticities. We begin with some simple descriptions of the variation in price elasticities across physicians in our data. Using data at the physician-patient-year level, we estimate for each physician i:

(3) $$\log(Q_{ijt}) = \beta^{i} \log(P_{it}) + \Gamma^{i} X_{ijt} + \epsilon_{it}.$$

Q_{ijt} is the total RVUs that patient j received from physician i in year t, P_{it} is the physician-specific price index as described in Section 3.1, and X_{ijt} are patient characteristics, including the Charlson Comorbidity Index, age, BMI, and indicators for white, black, Hispanic, and female. Because we estimate Equation (3) separately for each physician, β^{i} is physician i's price elasticity when seeing the average patient, and it is identified from variation in reimbursement rates that physician i faces over time. Finally, ϵ_{ijt} is an idiosyncratic error term, which we cluster at the 210 pre-1997 MPLs (i.e., the geographic payment regions).

If Medicare price changes are exogenous, β^{i} yields an unbiased estimate of the average price elasticity for each physician. However, as discussed in Section 3.2, there are several reasons why exogeneity might fail. Given the political nature of RVU changes, more popular procedures may draw higher Medicare payment increases. Alternatively, changes in payments may reflect recent or contemporaneous changes in the cost of performing a given procedure.[14] If costs are serially correlated, then changes in overall payment may be correlated with changes in costs. Finally, CMS updates RVUs based on comments submitted by physicians, health care workers, and professional associations and societies, increasing the likelihood of payment changes being correlated with other local supply factors (Federal Register).

In light of the potential threats to exogeneity, we use the 1997 geographic-specific shock and the 1999 PE-RVU procedure-specific shock for identification. These two shocks generated exogenous variation in Medicare reimbursements that is arguably unrelated to the local demand for and supply of services. We use them as instruments for observed Medicare payments, and our

[14] Although CMS uses the decennial census to determine certain indices, such as employee wage indices, it also uses the most recent retrospective data to determine other indices, such as office rental expenses.

first stage identifies the predictability of PE-RVU and GAF changes on overall Medicare payment changes while controlling for the covariates specified in Equation (3).

To isolate changes in GAF that are due to the 1997 policy change, our GAF instrument is calculated as the change in GAF from 1996 to 1997 that a given physician experiences. Specifically, in years prior to 1997, the instrument is equal to zero, and from 1997 onward, the instrument equals the one-time 1996 to 1997 change in GAF. Our second instrument isolates changes due to the 1999 PE-policy change. The PE-RVU instrument equals zero prior to 1999. From 1999 to 2002 when the new cost-based methodology was phased in, the instrument equals the annual change in each physician's PE-RVU index. Like the physician price index, we construct the PE-RVU index by taking a RVU-weighted sum of PE-RVUs across the basket of procedures that each physician performs. The weights, which are constant over time, are the same as those used in constructing P_{it}. Because the 1999 policy shock differentially affected payment changes for services performed in a facility versus non-facility setting, there may also be a change in PE-RVU induced through physician preferences, altruism, or other confounding factors. As such, we hold constant the share of a given service performed in the facility setting to the pre-policy years. Specifically, for physician i performing HCPCS h in year t, we calculate:

$$P\widehat{ERV}U_{iht} = \begin{cases} s_{iht} * PERVU_t^f \; + (1 - s_{iht}) * PERVU_{it}^{nf} & if\ t < 1999 \\ \overline{s_{ih}} * PERVU_t^f \; + (1 - \overline{s_{ih}}) * PERVU_t^{nf} & if\ t \geq 1999 \end{cases}$$

where the f and nf superscripts denote facility and non-facility components, respectively, and s_{iht} is the share of services performed in a facility setting for a given physician-HCPCS-year. In the post-1999 policy years, we use the average share $\overline{s_{ih}}$ of services from 1996 to 1998 that a given physician performed in a facility setting. Next, we calculate a PE-RVU index equal to a weighted sum across each physician's basket of procedures H^i: $P\widehat{ERV}U_{it} = \sum_{h \in H^i} \frac{RVU_{iht}}{\sum RVU_{iht}} * P\widehat{ERV}U_{iht}$. Finally, we calculate the annual change in $P\widehat{ERV}U_{it}$ for the policy transition years.

Appendix Figure A.2 plots the price variation we capture from each of these instruments. Plot (a) shows the change in the geographic GAF consolidation from 1996 to 1997, whereas plot (b) shows the annual change in the PE-RVU physician baskets from 1998 to 2002.

Identifying the Impact of Patient Income and Cost-Sharing. To identify when patient- or physician-driven behavior is more likely to occur, we identify how patient income and cost-sharing affect the magnitude of physician elasticities. For physician i treating patient j in year t, we estimate the following model:

(4) $\log(\mathrm{RVU}_{ijt}) = \beta \log(\mathrm{P}_{it}) + \eta \mathrm{Z}_{ijt} + \alpha[\log(\mathrm{P}_{ijt}) \times \mathrm{Z}_{ijt}] + \Gamma \mathrm{X}_{ijt} + \xi_i + \phi_t + \epsilon_{ijt}$

Here, Z_{ijt} represents either log income or expected OOP price for patient j. The coefficient on the interacted term (α) identifies the responsiveness of price elasticities to patient income or cost-sharing. If $\alpha < 0$, then physicians are more likely to exhibit patient-driven behavior (i.e., elasticities will be more negative), whereas if $\alpha > 0$, physicians are more likely to exhibit physician-driven behavior (i.e., elasticities will be more positive). We control for the patient demographics stated earlier (X_{ijt}), and we include physician (ξ_j) and year (ϕ_t) fixed effects so that the remaining variation is within a physician over time and across patients. Using this variation allows us to identify how the price-responsiveness of an individual physician varies with differing patient characteristics. Robust standard errors are again clustered by pre-1997 MPLs. To ensure that each physician is considered equally, we weight the sample by the inverse number of patients associated with each physician.

4. Results

4.1. Prediction 1: Heterogeneity in Elasticities

[Insert Figure 3 Here]

First, we show that the size and sign of price elasticities vary across physicians by estimating Equation (3) using the MCBS data. Ordering physicians by their price elasticities, we plot the price elasticities estimated via OLS in Figure 3a and 2SLS with both instruments in Figure 3b. To counteract the problem of multiple comparisons, we apply a Bonferroni correction and show only estimates that are statistically significant at the (0.05/3,521) level. Both subplots clearly indicate that physicians can appear either patient-driven with negative price elasticities or physician-driven with positive price elasticities. [15]

[Insert Table 2 Here]

[15] While the tail ends of these elasticities may appear large, it is important to keep in mind that we are estimating the price responsiveness relative to RVUs, as opposed to procedures. In our data, the average procedure is equivalent to 2 RVUs; the first and 99th percentiles of RVUs per procedure range from 0.24 to 20.3, respectively.

To check if the observed heterogeneity in price elasticities is driven by weak instruments, we examine the distribution of first stage F-statistics for the estimates shown in Figure 3b. As Panel A in Table 2 indicates, the 50th percentile of the F-statistic distribution is 8.93, and the 75th percentile has an average F-statistic is 25.56. Using the "rule of thumb" that F-statistics should exceed 10, this suggests that first-stage F-statistics are sufficiently large for a meaningful number of physicians in the sample (Olea and Pfluger, 2013; Staiger and Stock, 1997; Stock and Yogo, 2005). Moreover, in Appendix Figure B.1, we limit the price elasticities to those with first stage F-statistics greater than 10 and significant at the 5% level with a Bonferroni correction. A similar pattern of positive and negative elasticities emerges.

To test whether we should use one instrument or two, we perform a Sargan-Hansen test. The first row of Panel B in Table I indicates that about 29 percent of the physician elasticity estimates have a J-Statistic with p-value<0.05. In other words, for most of the estimates, the p-value is large, and the overidentifying restrictions are not rejected. Therefore, we rely on both instruments in subsequent IV estimates presented in this paper. Finally, to evaluate whether IV methods are required to estimate the model, we calculate the C-statistic under the null that exogeneity is supported by the data and both OLS and IV are consistent estimators. This statistic is calculated by examining the difference of two Sargan-Hansen statistics: one where payments are treated as endogenous (i.e., 2SLS) and another where payments are treated as exogenous (i.e, OLS).[16] Unlike the Durbin-Wu-Hausman test, this statistic is robust to violations of homoskedasticity (Sargan, 1958; Hansen, 1982). For most estimates, the p-values are small, suggesting that the price index is endogenous and that data rejects the use of OLS in favor of IV.

Our results highlight the heterogeneity in pricing dynamics, a fact that can be easily masked when estimating average price elasticities. We demonstrate this point by estimating the price responsiveness for the average physician. In Appendix Table B.1, we provide our estimates of an average price elasticity using a model akin to Equation (3). The overall impact of price changes on quantity is positive. Column (2), which uses only the GAF instrument, yields an aggregate elasticity of 1.390, which is comparable to a conceptually similar exercise performed by Clemens and Gottlieb (2014). Column (3), which uses only the PE-RVU instrument, also yields a positive elasticity of 1.022. In the remaining sections, we focus on variables from our model that can

[16] Under homeoskedasticity, this test is numerically equivalent to a Hausman test (Hayashi, 2000).

explain differences in the within-physician price response: patient SES and patient out-of-pocket costs.

4.2. Predictions 2 and 3: Patient SES and Out-Of-Pocket Costs

[Insert Table 3 Here]

We estimate Equation (4) to test the conjecture that lower patient SES and higher OOP costs increase the likelihood of patient-driven pricing for a given physician. The results are shown in Table 3. With the pooled data, the first stage F-statistic is large, suggesting that our instruments are sufficiently strong, and the Hansen J-statistics continue to suggest that the over-identifying assumptions are valid. The endogeneity test consistently rejects the OLS model, so we focus on the 2SLS results. Column (2) indicates that the interacted effect of log price and individual income is positive, meaning that price elasticities increase with patient income. All else equal, the same physician will have a larger quantity response when seeing a patient of higher income (i.e., the price responses become more physician-driven). Similarly, when patient income is low, as measured by Column (3)'s indicator for being dually enrolled in Medicaid, physician price responses tend to be more negative. When patient education is high (shown in Column 4), the coefficient is positive, but not statistically significant. Together, the evidence in columns (2)-(4) suggest that physicians are more likely to exhibit physician-driven behavior when their patients have higher consumption.

One may be concerned that patients with higher consumption demand and receive more intensive treatments from physicians. While we have not allowed inputs to vary with SES in our theoretical model, Equation (2) suggests that a positive correlation between input intensity (X) and consumption will bias our interaction of price and SES toward zero. Specifically, if physicians offer more intensive treatments to the rich, richer patients will be associated with not only lower levels of patient marginal utility from consumption (which increase price-responsiveness), but also greater relative risk-aversion by the physician (which reduces price-responsiveness). In the presence of this force, our estimated interaction between price elasticities and patient characteristics should be viewed as conservative.

Turning to the OOP results, Column (4) of Table 3 indicates that the interaction coefficient between expected patient OOP price and price is negative, providing support for the implication that physicians become more patient-driven when patients face higher OOP costs. The OOP effect on price elasticity (-0.12) is statistically significant and almost twice the size of the income effect (0.07). Finally, in Column (5), we include a price-interaction with all patient covariates of interest. While all estimates are not precisely estimated, the positive relationship between income and price responsiveness and the negative relationship between OOP and price responsiveness remain evident. For measures of SES, the interacted income and education coefficients are positive, though not significant, and the interacted Medicaid eligibility indicator is weakly negative. The effects of OOP on price responsiveness is robust, even when accounting for the impact of patient SES on physician responses.

Following Aiken and West (1991), we can identify the marginal effects of income and OOP evaluated at two standard deviations above or below the mean. This exercise traces out within-physician price elasticities across most of the patient income and OOP distributions. Using notation from Equation (4), this equates to calculating $\beta + \alpha[Z \pm 2\sigma_z]$, where Z is the mean log income or log OOP and $2\sigma_z$ is two standard deviations away from the mean. Standard errors are also calculated according to Aiken and West (1991). The bottom row of Table 3 indicates that patient income changes the average within-physician response from about 0.94 to 1.19. In other words, the average physician seeing poor patients with income two standard deviations below the mean behaves with a price elasticity of 0.94, but that same physician behaves with a price elasticity of 1.18 when seeing rich patients with income two standard deviations above the mean. Columns (3) and (4) show a similar result: the probability of Medicaid enrollment creates a within-physician elasticity range of 0.86 to 1.17, and patient education creates a within-physician elasticity range of 0.95 to 1.12. On the flip side, physicians have a higher price elasticity (1.04) when seeing patients with low expected OOP and a lower price elasticity (0.70) when seeing patients with high expected OOP.

[Insert Table 4 Here]

4.3. Robustness Tests

While the estimates in Tables 3 are identified by variation within physicians over time, it is possible that patient characteristics do not vary much within physicians. To ensure that our estimates are driven by those physicians with significant heterogeneity in patient characteristics, we examine the 2SLS results when limiting the sample to physicians with wide distributions in patient income and OOP costs. We divide the income and OOP variables into terciles and define physicians as having a wide distribution if they have patients in at least two different terciles. Physicians with narrow distributions are those with patient characteristics in only one tercile. Shown in Table 4, the results indicate that our estimates for income and OOP costs are being identified by physicians with a wide distribution of patient characteristics, as opposed to those with more uniform patients.

Another concern may be that there are insufficient patients per physician in our data to accurately identify variation in physician responses across patient types. In Appendix Table B.2, we sequentially limit the sample to physicians with "practice sizes" (defined by the number of patients per physician observed in the data) that are greater than the 25th, 50th, and 75th percentiles of the practice-size distribution. By doing so, the average number of patients per physician increases from 28 to 42 to 79, respectively. While these cuts in sample size limit the power of our data, the positive interacted income coefficient and negative interacted Medicaid indicator generally persist. The interacted OOP coefficient remains negative among physicians with practice sizes greater than 25th percentile. Taken together, Table 3 and Appendix Table B.2 provide confidence that our estimates are being identified from variation in patient-specific differences within a physician.

[Insert Table 5 Here]

Throughout this paper, we have also adopted the commonly applied assumption that patient cost-sharing changes do not generate income effects, allowing us to mitigate endogeneity concerns associated with patient SES influencing health through channels other than the marginal utility of consumption. However, as Nyman (2003) notes, this assumption is only valid when big fluctuations in cost-sharing are not present. When cost-sharing is high, such as when paying for a liver transplant, there can be large income effects. In Table 5, we assess whether our SES results persist in the range where our model assumptions are most realistic. We repeat our SES analysis, excluding instances where patients faced the top 10 (column 2) and 20 (column 3) percent of expected OOP spending. The results indicate that focusing on patients with lower cost-sharing magnifies the positive and significant price responsiveness to patient SES. The log income effect increases by about 27 to 34 percent, and the education effect increases by about 54 to 75 percent and becomes significant.

4.4. Policy Implications

[Insert Figure 4]

Our analysis highlights the heterogeneity of elasticities across both physicians and patients. In Figures 4, we calculate elasticities across the patient distribution for the average physician. This equates to calculating $(\hat{\beta} + \hat{\alpha}Z_{ijt})$ at different points of the patient characteristic distribution. In Figure 4, we demonstrate the variation in quantity responses for the average physician seeing patients in the 10th, 50th, and 90th percentile distribution of log income, education, and log OOP costs. We also consider Medicaid enrollment, which in the 10th percentile is zero (i.e., no Medicaid enrollment) and in the 90th percentile is one (i.e., dually enrolled). We test whether elasticities at the middle and top of the respective distributions are statistically different from those at the bottom. The hypothesis testing is performed using 1000 draws from a block bootstrap approach that resamples at the physician level.

Figure 4 shows that the average physician's price elasticity will range from approximately 0.88 to 1.14 as a result of variation in patient SES and from 0.76 to 0.98 as a result of variation in OOP spending. The bootstrap errors indicate that the within-physician elasticity estimates at the 10th and 50th percentiles of patient SES and OOP are statistically different for all variables except for education.[17] A similar conclusion holds for the elasticity estimates at the 50th and 90th percentiles. In the last column, we compute the variation in elasticities when considering all four characteristics simultaneously. The average physician seeing a patient with Medicaid, 10th percentiles log income and education, and 90th percentile OOP spending will temper their quantity response significantly: the price elasticity is only 0.6. However, that same physician seeing a patient without Medicaid, 90th percentiles log income and education, and 10th percentile OOP spending responds with a price elasticity of 1.06. These estimates are statistically different at the 1% level. These suggest that a uniform price change that ignores physician and patient heterogeneity will result in overuse of healthcare services, particularly among less altruistic physicians treating richer patients with low out-of-pocket costs. A more efficient design of price reform will target price increases toward patient-driven physicians treating low-SES, high-OOP patients, and it will lower prices for physician-driven physicians seeing high-SES, low-OOP patients.

The impact of patient characteristics on physician price responsiveness offer important implications for pricing and benefit-design policies more broadly within Medicare Parts A, B, and D. Although price reforms in Medicare Part B have historically been independent of patient income or patient out-of-pocket costs, Medicare Part A payments have differentially increased for hospitals serving a disproportionate share (DSH) of low-income Medicare, as well as Medicaid, patients. In Medicare Part D, cost-sharing amounts vary widely across plan choices, and over time, the median cost sharing for among branded drugs has increased (Hoadley et al., 2014). Our analyses imply that as a patient's cost-sharing burden increases, physicians will be more likely to prescribe cheaper alternatives, and overprescription will fall. Similarly, all else equal, richer patients are more likely to receive prescriptions for costlier drug options. In cases where drugs can be covered under either Part B or D, the added payment incentive from Part B prescriptions magnifies the role of patient out-of-pocket costs in assuring patient-driven behavior (Marrufo et al., 2011).

5. Conclusion

In this paper, we present a model of decision-making by physicians with imperfect altruism towards their patients. The interaction between altruism and self-interest explains how and why price increases can have dramatically different effects on quantity within different markets and procedures. Economic theory provides insights into the likelihood of positive or negative price elasticities, which we identify as physician-driven or patient-driven pricing behavior. Specifically, patient-driven behavior is more common when patient income is low and patient health care spending is high. We provide empirical evidence in support of these conjectures and illustrate these two patient characteristics alone can generate within-physician price responses that vary from 0.6 to 1.1 for the average physician. The theory suggests two remaining implications that could be tested in future work: patient-driven behavior is more common when physician altruism is high and when the physician's price-cost margin is low.

While our measures of price elasticities rely on plausibly exogenous variation, future research could incorporate not only exogenous changes in reimbursement, but also randomization in patient cost-sharing across physicians. In the absence of such an experiment, we cannot rule out selection bias from the sorting of physicians and patients. That said, as we discuss earlier, if richer patients seek out physicians who provide more services, our estimated variation in the price elasticity will be conservative. In addition, we have taken as given that physicians choose treatments on behalf of their patients. In some cases, patients may have more influence over the outcome; such shared decision–making and its implications for price elasticity represents another fertile area for future research.

The health economics literature has long recognized the tension between physician altruism and physician profit-maximization. In other healthcare contexts, economists have developed elegant and tractable models accounting for this tension. We exploit these tools to generate novel testable predictions about pricing and utilization behavior in healthcare markets. Our analysis demonstrates that the unique preferences and objectives of physicians create pricing dynamics in healthcare that depart from those in other product markets.

These implications seem consistent with the data and provide useful guidance for policymakers and researchers. First, physicians are systematically more “altruistic” – in the sense of pursuing patient interests – when treating more vulnerable and disadvantaged patients. Second, heterogeneity in the effect of reimbursement changes is to be expected, and can be exploited to increase the effectiveness of reimbursement reforms. Reimbursement reductions might be useful tools for containing costs when physicians are largely profit-maximizing, but they may be counterproductive when they are more altruistic. Being able to differentiate when a service or market is physician- versus patient-driven allows policy makers to more effectively target supply- and demand-side incentives. More generally, economic theory provides policymakers with guidance on the source and nature of variation in price elasticities. Suitably directed empirical analysis can help inform more targeted approaches to reforming reimbursement policy, particularly when the goal is to restrain or boost the quantity of healthcare utilized.

REFERENCES

Aiken, Leona S. and Stephen G. West. 1991. "Multiple Regression Testing and Interpreting Interactions," Sage Publications: Thousand Oaks, CA.

Andrews, Donald W.K., Marcelo Moreira, and James Stock. 2006. “Optimal Two-Sided Instrumental Variables, Regression,” *Econometrica*, 74: 116-132.

Arrow, Kenneth J. 1963. “Uncertainty and the Welfare Economics of Medical Care,” *American Economic Review*, 53: 941-973.

Becker, Gary. 1975. The Economics of Discrimination, 2nd ed. Chicago, IL: The University of

Chicago Press.

Bhattacharya, Jay and Darius Lakdawalla. 2006. "Does Medicare Benefit the Poor?" *Journal of Public Economics*, 90(1-2): 277-292.

Blomqvist, Ake. 1991. “The Doctor as Double Agent: Information Asymmetry, Health Insurance and Medical Care,” *Journal of Health Economics*, 10: 411- 432.

Chetty, Raj, John Friedman, and Jonah Rockoff. 2014. "Measuring the Impacts of Teachers 1: Evaluating Bias in Teacher Value-Added Estimates," *American Economic Review* 104: 2593-1521.

Clemens, Jeffrey and Joshua D. Gottlieb. 2014. "Do Physicians’ Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review*, 104 (2014), 1320-49.

———, "Bargaining in the Shadow of a Giant: Medicare's Influence on Private Payment Systems," *Journal of Political Economy*, Forthcoming, 2016

Choné, Philippe and Ching-to A. Ma. 2011. "Optimal Health Care Contract Under Physician Agency," *Annals of Economics and Statistics*, 101/102: 229-256.

Dickstein, M. 2014. "Physician vs. Patient Incentives in Prescription Drug Choice," Draft.

Dranove, David. 1988. "Demand Inducement and the Physician-Patient Relationship," *Economic Inquiry*, 26:281-296.

Dranove, David and William D. White. 1987. “Agency and the Organization of Health Care Delivery,” *Inquiry*, 24: 405-415.

Eaton, Joe. 2008. "Little-known AMA Groups has Outsized Influence on Medicare Payments," The Center for Public Integrity, accessed December 22, 2015, <http://www.aafp.org/fpm/2008/0200/p36.html>.

Ellis, R. and McGuire, T. 1985. "Provider Behavior Under Prospective Reimbursement: Cost Sharing and Supply," *Journal of Health Economics*, 5: 129-151.

Ellis, Randall P. and Thomas G. McGuire. 1986. "Optimal Payment Systems for Health Services," *Journal of Health Economics*, 9:375-396.

Emanual, Ezekial J. and Linda L. Emanual. 1992. "Four Models of the Physician-Patient Relationship," *Journal of the American Medical Association*, 267: 2221-2226.

Department of Health and Human Services. 1996. "Medicare Program; Revisions to Payment Policies and Five-Year Review of and Adjustments to the Relative Value Units Under the Physician Fee Schedule for Calendar Year 1997," Federal Register, 61: Report 59490.

Farley, Pamela J. 1986. "Theories of the Price and Quantity of Physician Services," *Journal of Health Economics*, 5: 315-333.

Felder, Stefan and Thomas Mayrhofer. 2011. Medical Decision Making: a Health Economic Primer, Heidelberg: Springer.

Gaynor, Martin. 1994. “Issues in the Industrial Organization of the Market for Physician Services,” *The Journal of Economics & Management Strategy*, 211-255.

Glied, Sherry and Joshua Zivin. 2002. "How do Doctors Behave When Some (But Not All) of Their Patients are in Managed Care?" *Journal of Health Economics*, 21: 337-53.

Gruber, Jonathan. 2013. "Proposal 3: Restructuring Cost Sharing and Supplemental Insurance for Medicare," Brookings: The Hamilton Project, 1-6.

Hadley, Jack and Robert Lee. 1978. “Toward a Physician Payment Policy: Evidence From the Economic Stabilization Program,” *Policy Sciences*, 10: 105-120.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50: 1029-1054.

Hayashi, Fumio. 2000. Econometrics. Princeton, NJ: Princeton University Press.

Hoadley, Jack, Laura Summers, Elizabeth Hargave, et al. 2014. Medicare Part D in Its Ninth Year: THe 2014 Marketplace and Key Trends, 2005-2014. *Kaiser Family Foundation.*

Lakdawalla, Darius and Neeraj Sood. 2013. "Health Insurance as a Two-Part Pricing Contract," *Journal of Public Economics*, 102: 1–12.

Liu, Ting and Ching-to Albert Ma. 2013. "Health Insurance, Treatment Plan, and Delegation to Altruistic Physician," *Journal of Economic Behavior & Organization*, 85: 79-96.

Lu, Mingshan. 1999. “Separating the ‘True Effect’ from ‘Gaming’ in Incentive-based Contracts in Health Care,” *Journal of Economics and Management Strategy*, 8: 383-432.

MaCurdy, Thomas, Jason Shafrin, Thomas DeLeire, Jed DeVaro, Mallory Bounds, David Pham, and Arthur Chia. 2012. "Geographic Adjustment of Medicare Payments to Physicians: Evaluation of IOM Recommendations," Burlingame, CA: Acumen, LLC.

Marrufo, Grecia M, Emil Rusev, Kristy Piccinini, Elizabeth Coombs, Ken Ueda, and Erica Schecter. 2011. "Estimating the Effects of Consolidating Drugs under Part D or Part B." Burlingame, CA: Acumen, LLC.

Maxwell, Stephanie and Stephen Zuckerman. 2007. "Impact of Resource-Based Practice Expenses on Medicare Physician Volume," *Health Care Financing Review*, 29: 65-79.

McClellan, Mark and Jonathan Skinner. 2006. "The Incidence of Medicare," *Journal of Public Economics*, 90(1-2): 257-276.

McGuire, Thomas. 2000. "Physician Agency" in Handbook of Health Economics, Volume I, Culyer, Anthony J Culyer and Joseph P. Newhouse, eds. Amsterdam, The Netherlands: Elsevier Science.

Mitchell, Janet B. and Susan G. Haber, "State Pyament Limitations on Medicare Cost-Sharing: Impacts on Dually Eligible Beneficiaries and Their Profiders," Health Economics Research, Inc., (2003), CMS Contract No. 500-95-0058.

Mooney, Gavin and Mandy Ryan. 1993. “Agency in Health Care: Getting Beyond First Principles,” *Journal of Health Economics*, 12: 125-135.

Nyman, John. 2003. "The Theory of Demand for Health Insurance." Stanford University Press: Stanford, California.

Olea, José Luis Montiel and Carolin Pfluger. 2013. "A Robust Test for Weak Instruments," *Journal of Business and Economic Statistics*, 31: 358-369.

Pauly, Mark. 1980. Doctors and Their Workshops: Economic Models of Physician Behavior,” Chicago, IL: University of Chicago Press.

Quan, Hude, Vijaya Sundararajan, Patricia Halfon, Andrew Fong, Bernard Burnand, Jean-Christophe Luthi, L. Duncan Saunders, Cynthia A. Beck, Thomas E. Feasby and William A. Ghali. 2005. "Coding Algorithms for Defining Comorbidities in ICD-9-CM and ICD-10 Administrative Data," *Medical Care*, 43(11): 1130-1139.

Sargan, J.D. 1985. "The Estimation of Economic Relationships Using Instrumental Variables," *Econometrica*, 26: 393-415.

Staiger, Douglas and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65: 557-586.

Stock, James H. and Motohiro Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression," in Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg, Donald W.K. Andrews and James H. Stock, ed. Cambridge: Cambridge University Press.

U.S. Government Accountability Office, "Medicare Physician Fees Geographic

Adjustment Indices are Valid in Design, But Data and Methods Need Refinement," Pub No. GAO-05-119, 2005.

Zweifel, Peter, Friedrich Breyer, and Mathias Kifmann. 1997. Health Economics, 2nd ed. New York, NY: Oxford University Press.

Zuckerman, Stephen, Baoping Shang, and Timothy Waidman. 2010. "Reforming Beneficiary Cost Sharing to Improve Medicare Performance," *Inquiry*, 47: 215-225.

FIGURE 1: HISTOGRAM OF ELASTICITIES MEASURED USING LARGE PRICE INCREASES

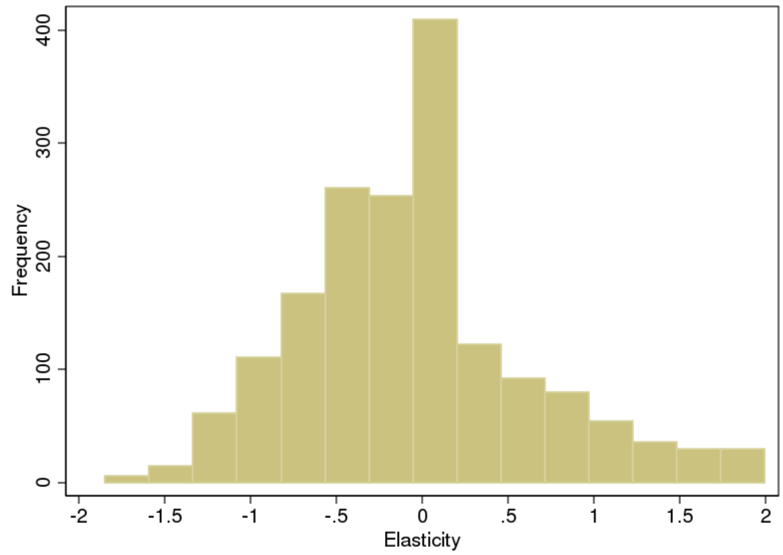

Notes: Data from MCBS, 1993-2002. This figure shows the elasticities (calculated simply as the annual percent change in quantity divided by the annual percent change in price) for HCPCS with annual physician payment increases above 50 percent. It is evident that quantity increases for about half of the HCPCS, while quantity falls for the other half. The long right tail has been truncated.

TABLE 1: SUMMARY STATISTICS

	MCBS Data	
	Mean	SD
Price Index ($)	273.97	893.81
Total RVUs	8.04	22.95
ΔGAF	0.0011	0.014
ΔPE-RVU	-0.11	0.54
CCI	2.51	1.96
1(Male)	0.42	0.49
1(White)	0.86	0.35
1(Black)	0.082	0.28
1(Hispanic)	0.047	0.21
Age	74.57	12.97
BMI	25.02	5.65
Income	24,080	45,571
Education	11.58	3.58
1(Medicaid)	0.22	0.42
Expected OOP	10.32	20.47
No. Observations	257,669	

Notes: Data from the Federal Register and MCBS at the physician-patient-year level.

FIGURE 2: SHOCKS IN MEDICARE PAYMENT COMPONENTS

(a) Geographic Consolidation of Payment Regions (GAF)

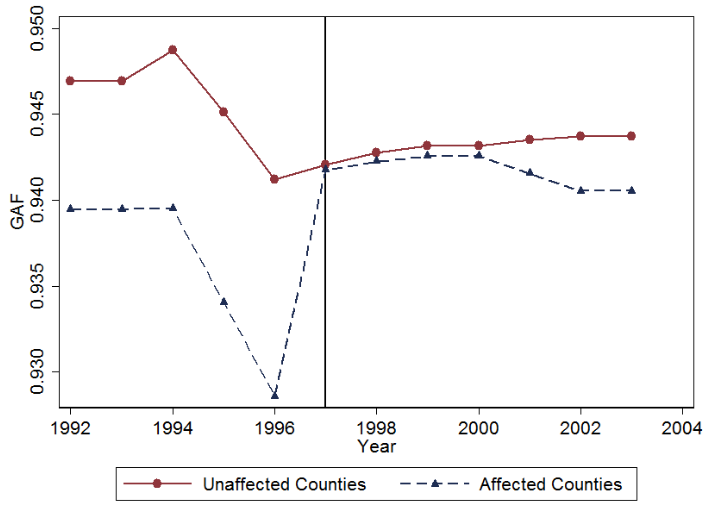

(b) Change in Reimbursement Calculation Method (PE-RVU)

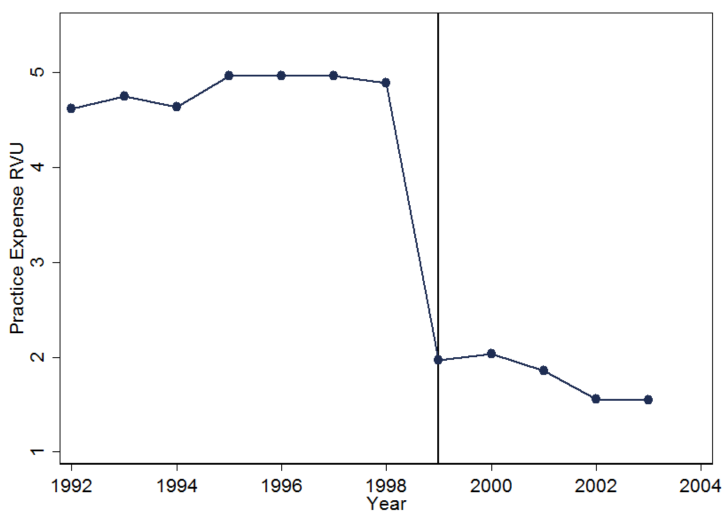

Notes: Data from the Federal Register 1992-2003. The sample is limited to HCPS observed in all years. Plot (a) shows the average GAF across counties that were or were not affected by the 1997 consolidation of payment regions from 210 to 89 payment regions. Plot (b) depicts the change in average PE-RVUs across HCPCS. In 1999, HCFA phased in a cost-based methodology of calculating PE-RVUs and more accurately priced non-facility services.

FIGURE 3: HETEROGENEITY IN PHYSICIAN ELASTICITIES

(a): OLS-Estimated Price Elasticities

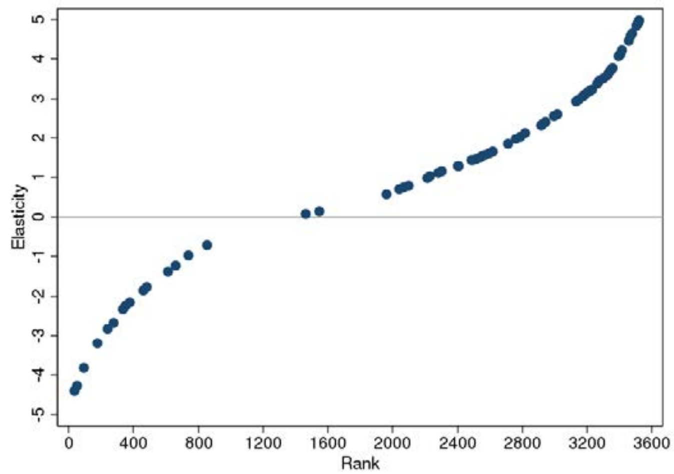

(b): 2SLS-Estimated Price Elasticities

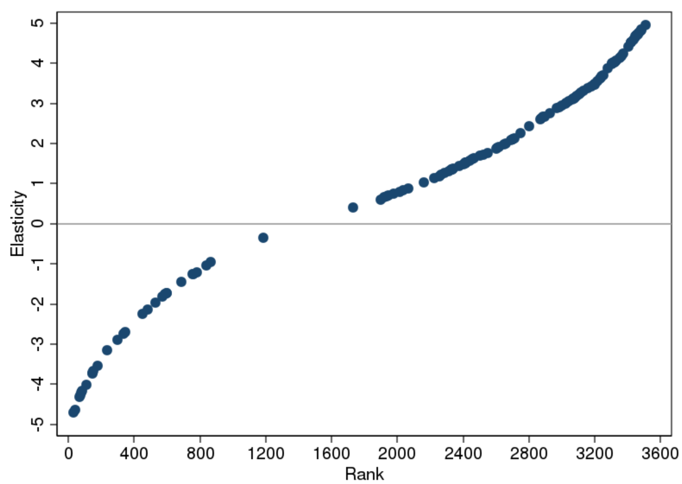

Notes: Data from the Federal Register and MCBS. Each dot represents the price elasticity for a given physician, estimated using Equation (3). Elasticities are ordered and plotted. Outliers with elasticities outside of ± 5 are not shown. A Bonferroni correction has been applied; for both plots, only HCPCS with statistically significant price elasticities with p-value<(0.05/3,521) are shown.

TABLE 2: SUMMARY OF IV RELATED STATISTICS

	(1)	(2)	(3)
	Panel A: First Stage F-Statistics		
	25^{th} percentile	Median	75^{th} percentile
First Stage F-Stat	2.56	8.93	25.56
	Panel B: Endogeneity and Over-Identification		
	Fraction with p-value<0.10	Fraction with p-value<0.05	Fraction with p-value<0.01
Hansen J-Statistic	0.403	0.289	0.126
Endogeneity Test	0.937	0.863	0.296

Notes: Data from the Federal Register and MCBS. This table shows the IV related summary statistics used to estimate the statistically significant elasticities at the 5 percent level shown in Figure 3(b). Panel A shows the distribution of first stage F-statistics when using both PE-RVU and GAF as instruments. Panel B shows distribution of p-values for the Hansen J statistic for the over-identifying restrictions being valid and the C-statistic test of endogeneity for log price being exogenous.

TABLE 3: EFFECT OF PATIENT CHARACTERISTICS ON PRICE ELASTICITIES

	(1)	(2)	(3)	(4)	(5)	(6)
Log(Price)	1.026***	0.394	1.056***	0.892***	1.119***	0.824**
	(0.0956)	(0.340)	(0.0963)	(0.140)	(0.100)	(0.371)
Log(Price) x						
Log(Income)		0.0695**				0.0311
		(0.0339)				(0.0370)
1(Medicaid)			-0.181***			-0.146*
			(0.0777)			(0.0838)
Education				0.0123		0.00351
				(0.00845)		(0.00854)
Log(OOP)					-0.123***	-0.129***
					(0.0415)	(0.0417)
No. of Obs.	257,669	257,669	257,669	257,669	257,669	257,669
First F-Stat	55.8	202.2	63.0	295.1	31.56	31.56
Hansen J-Stat	0.510	0.702	0.152	0.312	0.124	0.124
Endogeneity	1.34e-7	1.75e-6	2.27e-7	1.22e-6	8.56e-6	8.56e-6
Implied Price Elasticity at Interacted Variable's:						
Mean – 2sd	---	0.941***	1.167***	0.947***	1.046***	---
Mean	---	1.064***	1.012***	1.035***	0.871***	---
Mean + 2sd	---	1.187***	0.858***	1.123***	0.695***	---

Notes: Data from the Federal Register and MCBS. All regressions control for patient characteristics (CCI, male, white, black, Hispanic, age, married, widowed, BMI, log income, Medicaid indicator, years of education, and log expected OOP), physician fixed effects (53,351 indicators) and year fixed effects. Standard errors, shown in parentheses, are clustered by the pre-1997 MPL. *10 percent level, **5 percent level, ***1 percent level.

TABLE 4: PHYSICIANS WITH WIDE AND NARROW DISTRIBUTIONS OF PATIENTS

	All (1)	Wide Distribution (2)	Narrow Distribution (3)
	A. Income		
Log(Price) x	0.0695**	0.0606*	0.484
Log(Income)	(0.0339)	(0.0329)	(0.979)
No. Observations	257,669	216,100	41,569
No. Physicians	53,351	36,653	16,698
	B. Education		
Log(Price) x	0.0123	0.0105	-0.0176
Education	(0.00845)	(0.00747)	(0.112)
(yrs.)			
	257,669	205,094	52,575
No. Observations	53,351	32,744	20,607
No. Physicians			
	C. Expected OOP		
Log(Price) x	-0.123***	-0.156***	0.0455
Log(OOP)	(0.0415)	(0.0345)	(0.465)
No. Observations	257,669	218,266	39,403
No. Physicians	53,351	38,049	15,302

Notes: Data from the Federal Register and MCBS. See notes to Table 3. Physicians with wide distributions are those with patients in at least two terciles of the income (Panel A), education (Panel B), or expected-OOP (Panel C) distributions. Physicians with narrow distributions are those seeing patients with characteristics within only one tercile of the patient-characteristic distribution. Standard errors, shown in parentheses, are clustered by the pre-1997 MPL. *10 percent level, **5 percent level, ***1 percent level.

TABLE 5: EDUCATION EFFECTS, EXCLUDING HIGH OOP INCIDENCES

	All (1)	All but top 10% OOP (2)	All but top 20% OOP (3)
	A. Income		
Log(Price) x	0.0695**	0.0931**	0.0882**
Log(Income)	(0.0339)	(0.0387)	(0.0399)
No. Observations	257,669	226,927	196,176
No. Physicians	53,351	48,719	43,795
	B. Education		
Log(Price) x	0.0123	0.0189**	0.0215**
Education	(0.00845)	(0.00962)	(0.00972)
(yrs.)			
	257,669	226,927	196,176
No. Observations	53,351	48,719	43,795
No. Physicians			

Notes: Data from the Federal Register and MCBS. See notes to Table 3. In column (2), we drop all patients in the tenth decile of the expected OOP distribution. In column (3), we drop all patients in the top ninth and tenth deciles of the expected OOP distribution. Standard errors, shown in parentheses, are clustered by the pre-1997 MPL. *10 percent level, **5 percent level, ***1 percent level.

FIGURE 4: ELASTICITIES ACROSS PATIENT CHARACTERISTICS

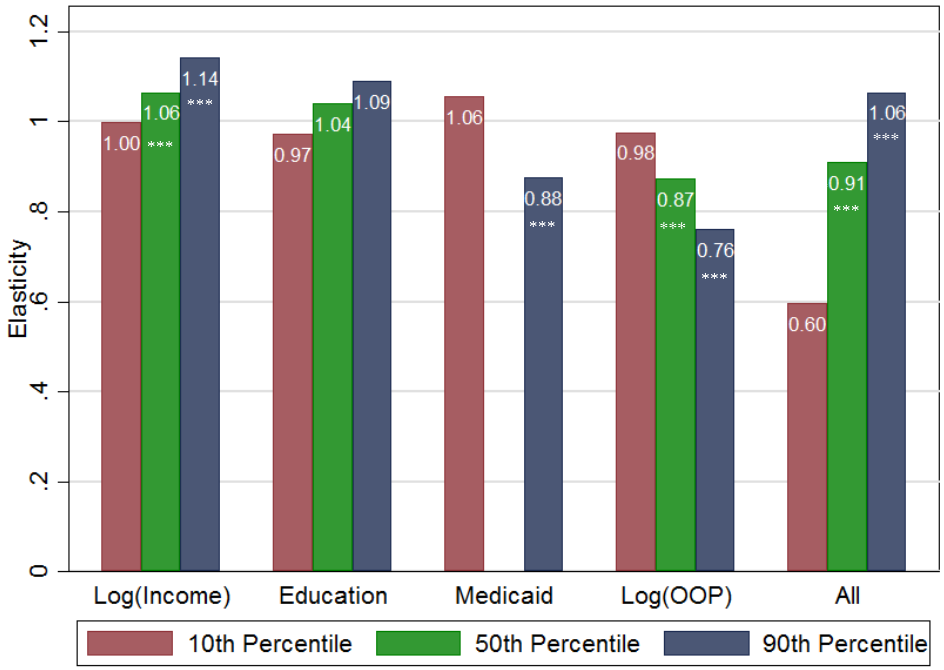

Notes: Data from the Federal Register and MCBS. This figure shows the within-physician quantity response to a 10% price change. The 10^{th}, 50^{th}, and 90^{th} percentile corresponds to patients with log income of 8.7, 9.7, 10.8; education of 6.5, 12, and 19 years; no/yes enrollment into Medicaid; and log expected OOP costs of 1.2, 2.0, 2.9. The first column of the "All" category examines patients with Medicaid in the 10^{th} percentiles of log income and education and 90^{th} percentile of log OOP costs. The middle column corresponds to patients without Medicaid in the 50^{th} percentile distribution of log income, education, and log OOP costs. The last column considers patients again without Medicaid in the 90^{th} percentiles of log income and education and 10^{th} percentile of log OOP costs. *** indicates that elasticity in the bar and the right of the bar (i.e., 50^{th} and 10^{th} percentile, or 90^{th} and 50^{th} percentile) are statistically different at the 1% level.

FOR ONLINE PUBLICATION ONLY

Appendix A

In this section, we further illustrate the distribution of price variation from our GAF and PE-RVU instruments, shown in Figures A1 and A2. We also discuss the remaining policy changes during 1992 to 2003 that affected Medicare payments. Other than the GAF and PE-RVU components detailed in Section 3.2, variation in Medicare payments come from changes in the work RVU, malpractice RVU, and CF. On average, work RVUs, PE-RVUs, and malpractice RVUs account for 52 percent, 44 percent, and 4 percent of total payments, respectively (US Government Accountability Office, 2005). Because the malpractice component accounts for such a small share of payments, we do not focus on it.

From 1993 to 2002, work RVUs experienced two major reviews which became effective in 1997 and 2002. Plot (a) of Figure A.3 shows the average work RVU over time for HCPCS. After the RUC committee met to re-assess work RVUs, we see clear jumps in the RVU. However, with competing political pressures and physician incentives, it is unlikely that RUC committee changes are exogenous to local demand and supply factors.

The CF also experienced a major change during our study period. Prior to 1998, there were three different CFs: one for surgery, primary care, and non-surgical services. The CF for surgical procedures led to surgeons earning a 17 percent bonus payment relative to all other procedures. This generated political discontent and led to a budget-neutral merger of CFs in 1998 (Clemens and Gottlieb, 2013). Plot (b) shows the CFs over time. After 1998, the CF for surgical procedures fell by about 11 percent, whereas the CF for non-surgical procedures increased by about 6 percent. We do not use this policy shock as another instrument for two reasons. First, CFs are constant across all geographic regions and all procedures, so their explanatory power for payment changes within physicians is weak. Second, the shock in CF payments occurs mainly for surgical procedures, while changes in CF for non-surgical and primary care procedures are much less pronounced.

FIGURE A1: PRACTICE EXPENSE RVU, BY FACILITY

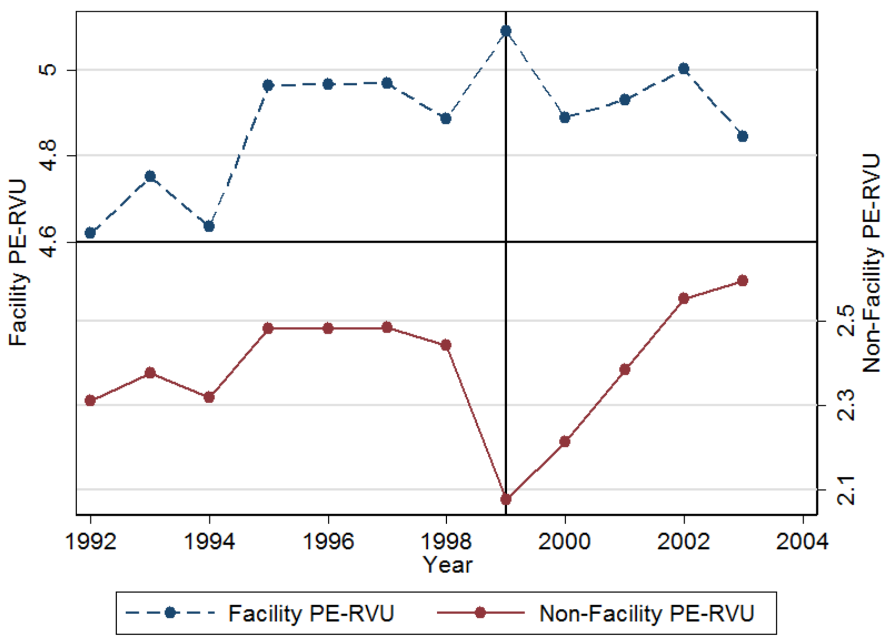

Notes: Data from the Federal Register 1992-2003. The top line shows changes in the facility PE-RVU. The bottom line shows changes in the non-facility PE-RVU. Sample restricted to HCPCS observed in all years.

FIGURE A2: DISTRIBUTION OF INSTRUMENT-INDUCED PRICE SHOCKS

(a) Price Shock from GAF-Consolidation

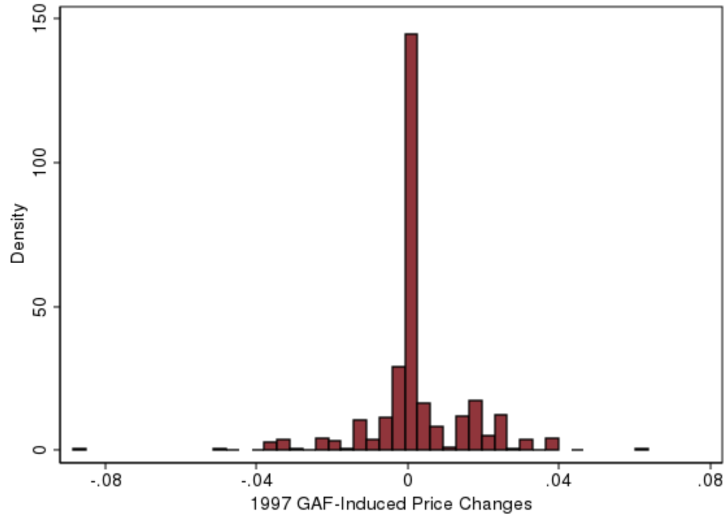

(b) Price Shock from PE-RVU Methodology Change

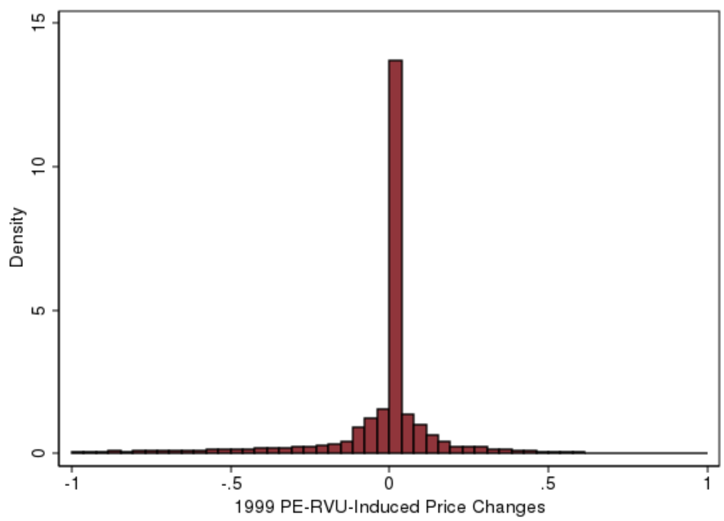

Notes: Data from the Federal Register and MCBS at the physician-patient-year level. Plot (a) shows the distribution of the GAF-induced price change from 1996 to 1997. Plot (b) shows the distribution of PE-RVU induced price changes for physician baskets of goods from 1998 to 2002.

FIGURE A2: REMAINING VARIATION IN MEDICARE PAYMENTS

(a) Average Work RVU Over Time

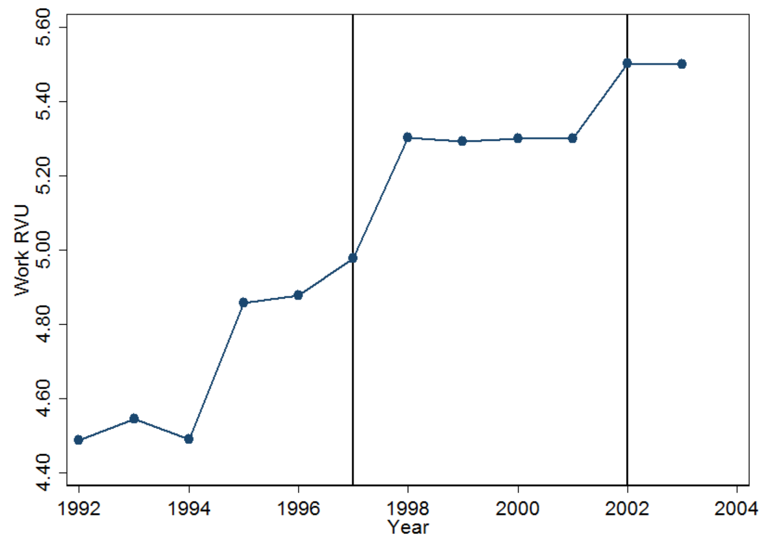

(b) Conversion Factor

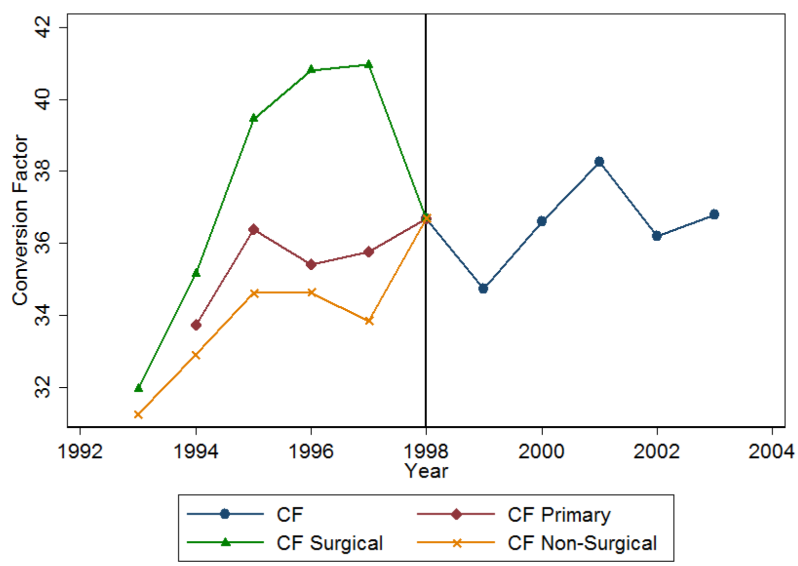

Notes: Data from Federal Register 1992-2003. Plot (a) show the change in work-RVUs. Evident from the graph are the two major reviews by the RUC committee in 1997 and 2002. The sample is restricted to HCPCS observed in all years. Plot (b) shows the change from three CFs (primary care, surgical, and non-surgical) to a single budget-neutral CF in 1998.

FOR ONLINE PUBLICATION ONLY

Appendix B

APPENDIX TABLE B.1: ELASTICITES USING DIFFERENT METHODS

	OLS (1)	IV=GAF (2)	IV=PE-RVU (3)	2SLS (4)
Log(Price)	0.132***	1.390***	1.022***	1.026***
	(0.0152)	(0.627)	(0.0963)	(0.0956)
First Stage F-Stat	---	19	1142	578.7
No. of Observations	257,669	257,669	257,669	257,669

Notes: Data from Federal Register and MCBS at the physician-patient-year level. The dependent variable is log(total RVU). The main independent variable is log(price per RVU). Covariates included are patient age, CCI, gender, and race dummies, and we control for year and physician fixed effects. Robust standard errors clustered by the 210 pre-1997 Medicare Practice Localities are shown in parentheses * 10 percent level, ** 5 percent level, *** 1 percent level.

APPENDIX FIGURE B.1: 2SLS-ESTIMATED PRICE ELASTICITIES WITH F-STATISTIC HIGHER THAN 10

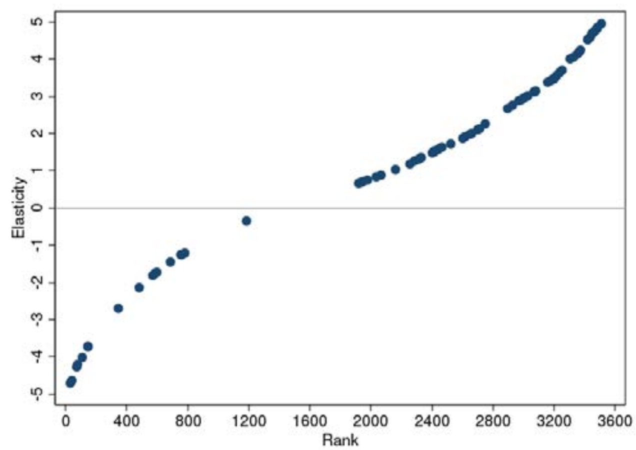

Notes: Data from the Federal Register and MCBS. This plot replicates Figure 3b, including only physician-specific regressions where the first stage F-statistic is greater than 10 and the price elasticity estimate is significant at the (0.05/3,521) level.

APPENDIX TABLE B.2: EFFECT OF PATIENT CHARACTERISTICS ACROSS PHYSICIANS WITH VARYING MEDICARE PRACTICE SIZES

	All Physicians	# Patients >25th Pctile	# Patients >50th Pctile	# Patients >75th Pctile
	(1)	(2)	(3)	(4)
Log(Price)	0.824**	0.891***	0.281	0.493
	(0.371)	(0.358)	(0.454)	(0.525)
Log(Price) x				
Log(Income)	0.0311	0.0289	0.0964**	0.191***
	(0.0370)	(0.0361)	(0.0446)	(0.0563)
1(Medicaid)	-0.146*	-0.144*	-0.122	-0.228
	(0.0838)	(0.0769)	(0.0775)	(0.175)
Education	0.00351	0.00423	-0.00817	-0.00317
	(0.00854)	(0.00778)	(0.0114)	(0.0126)
Log(OOP)	-0.129***	-0.109**	0.0129	-0.0372
	(0.0417)	(0.0446)	(0.0484)	(0.108)
No. of Obs.	257,669	217,387	139,408	65,973
No. Physicians	55,616	33,621	10,975	1,986
Avg No. Pat. Per Phy.	24	28	42	79

Notes: Data from the Federal Register and MCBS. All regressions control for patient characteristics (CCI, male, white, black, Hispanic, age, married, widowed, BMI, log income, Medicaid indicator, years of education, and log expected OOP), physician fixed effects (53,351 indicators) and year fixed effects. Standard errors, shown in parentheses, are clustered by the pre-1997 MPL. *10 percent level, **5 percent level, ***1 percent level.