Published in final edited form as:
J Polit Econ. 2007 ; 115: 103–140.

PRODUCTIVITY SPILLOVERS IN HEALTHCARE: EVIDENCE FROM THE TREATMENT OF HEART ATTACKS*

Amitabh Chandra and

Harvard University and the National Bureau of Economic Research

Douglas O. Staiger

Dartmouth College and the National Bureau of Economic Research

Abstract

A large literature in medicine documents variation across areas in the use of surgical treatments that is unrelated to outcomes. Observers of this phenomena have invoked “flat of the curve medicine” to explain these facts, and have advocated for reductions in spending in high-use areas. In contrast, we develop a simple Roy model of patient treatment choice with productivity spillovers that can generate the empirical facts. Our model predicts that high-use areas will have higher returns to surgery, better outcomes among patients most appropriate for surgery, and worse outcomes among patients least appropriate for surgery, while displaying no relationship between treatment intensity and overall outcomes. Using data on treatments for heart attacks, we find strong empirical support for these and other predictions of our model, and reject alternative explanations such as waste or supplier induced demand, for geographic variation in medical care

I. Introduction

Since Sir Allison Glover (1938) first documented significant variation in tonsillectomy rates across areas in the United Kingdom, an enormous body of literature in economics and medicine has documented variations in the use of intensive treatments (that is surgical or technologically intensive treatments) across comparable locales. Surprisingly, the use of more intensive procedures is not associated with improved satisfaction, outcomes, or survival but is associated with significantly higher costs. Traditional explanations such as sampling variation, differences in income and insurance, patient-preferences, and underlying health status do not explain these variations.¹ These facts stand in sharp contrast to the results of Randomized Clinical Trials (RCT) that consistently find gains from the surgical management of acute conditions that are routinely interpreted as evidence in support of more intensive management of patients.²

These apparently conflicting findings can potentially be explained by a model of diminishing returns. Whereas RCTs are performed on a group of patients most likely to benefit from the intervention, the lack of a cross-sectional relationship between intensity and outcomes is

*We are grateful to our editors, three astute reviewers, and seminar participants at the University of Chicago, Cornell, Dartmouth, Harvard, the LSE, Oregon, Penn, Princeton, UCLA, UCSD, RAND, Syracuse, Virginia and Yale for helpful comments. Jeffrey Geppert, Dan Gottlieb and Lee Lucas provided expert assistance with the data. We received support from NIA P01 AG19783-02. The opinions in this paper are those of the authors and should not be attributed to the National Institute of Aging or the National Bureau of Economic Research.

¹Similar observations have been made by comparing treatments and outcomes in the US with those in Canada and Europe. Phelps (2000) provides an economists perspective of the immense literature on geographic variation in physician practice style. The medical literature on this topic is succinctly summarized in Chassin et. al (1987), Fisher et al (1994), Baicker and Chandra (2004a, 2004b) and the *Dartmouth Atlas of Health Care* [Wennberg and Cooper (1999) and Wennberg and Birkmeyer (2000)].

²In the context of heart attack treatments, see for example, Anderson et.al (2003) and Jacobs (2003). A review of over twenty-three trials by Keeley, Boura and Grines (2003) noted the demonstrated superiority of the intensive intervention over alternative therapies for the management of heart attacks.

explained by a “flat of the curve” argument, where physicians perform the intervention until the marginal return is zero [Enthovan(1980), Fuchs (2004)]. This is the interpretation that is used by McClellan et.al (1994) and McClellan and Newhouse (1997), to explain the small returns from more intensive treatments for heart attacks. This explanation emphasizes inefficiency of medical spending, suggesting that welfare improvements may be realized by reducing spending in high use regions. Indeed, Fisher et al. (2003a,b) use this logic to argue that a 30 percent reduction in Medicare spending, such that spending in high use regions is reduced to that of low intensity areas, would not have any deleterious effects on patient outcomes or satisfaction.

While the diminishing returns model is intuitively appealing, it has a number of problems. First, there is no reason to expect wide variation in the use of treatments across areas that are similar, without making additional assumptions such as area norms or supplier-induced demand. Second, such models still predict a positive relationship between medical spending in a region and patient outcomes unless all areas are in the range of zero or negative marginal benefits; this has never been found in the literature. A more fundamental problem with the diminishing returns model is that it predicts that the marginal benefit from more intensive patient treatment is lower in areas that are more aggressive. But US-Canada comparisons (Beck et al., 2003 and Pilote, McClellan, et al., 2003) suggest the opposite: the marginal benefit from more technologically intensive treatment in heart attack patients is larger in the US, where management of heart attacks is much more aggressive. These are facts that the flat of the curve model cannot reconcile. If it is the incorrect model, then embracing its policy implications of cutting spending could result in deleterious outcomes for many patients.

In this paper we develop a simple model of specialization in healthcare that uses a prototypical Roy model of treatment choice as its principal building block. In our model, patients are treated medically or intensively depending on their clinical appropriateness for these treatments, but we allow productivity spillovers to also affect the returns from each treatment choice. These spillovers may arise from knowledge spillovers, where physicians build experience by learning from one and other, but may also arise from other sources, such as the selective migration of the best intensive physicians to certain areas, and the corresponding migration of physicians who specialize in less intensive treatments to other areas. With these spillovers, as the proportion of patients in an area that are treated intensively increases, productivity spillovers increase the return to intensive treatment, while simultaneously reducing the return to the competing medical treatment. Thus, this model naturally generates higher returns to receiving intensive treatment in intensive areas, yet, because of the negative externality on patients receiving the competing treatment, does not necessarily generate any relationship between specialization and overall health outcomes. As we discuss in the next section, the model generates a rich set of sharp predictions that we test using detailed data on a sample of Medicare beneficiaries who had a heart attack (clinically referred to as an Acute Myocardial Infarction (AMI)). Several of these predictions are consistent with other models such as the flat of the curve model, but others are unique to the Roy model with productivity spillovers. We find strong empirical support for our model, and reject alternative explanations such as “flat of the curve medicine,” or supplier induced demand, that are commonly proposed to explain geographic variation in medical care.

In Section II, we generalize a prototypical Roy model of specialization to allow for productivity spillovers, and derive a wide range of testable predictions of this model. In Section III we justify our focus on heart-attacks to test this model, and describe our data and estimation strategy. Section IV tests the implications of our model. Section V concludes and discusses the policy implications of our work including its implications for the interpretation of randomized clinical trials. In the data appendix to this paper we provide details of our estimation and sample.

II. A Roy Model of Heart Attack Treatments with Productivity Spillovers

We consider a simple Roy model of patient treatment choice in the presence of productivity spillovers. In this model, an individual patient chooses between two alternative treatments: non-intensive management (emphasizing medical management, denoted with subscript 1), versus intensive intervention (emphasizing surgery, denoted with subscript 2). The physician chooses the treatment option for each patient that maximizes utility, based on the expected survival rate ($Survival_1, Survival_2$) and cost ($Cost_1, Cost_2$) associated with each option. The survival rate from a given treatment option depends on patient characteristics (Z), but is also positively related to the fraction of all patients that receive the same treatment ($P_1, P_2=1-P_1$).³ Similarly, the cost from a given treatment option also depends on patient characteristics, and may be negatively related to the fraction of all patients that receive the same treatment. Thus, a patient's choice of a specific treatment has a positive productivity externality, raising survival and reducing costs for all other patients receiving the same treatment.

The fundamental assumption of productivity spillovers in our model is quite plausible in health care. Much of physician learning about new techniques and procedures occurs from direct contact with other physicians ("see one, do one, teach one"), which leads to natural Marshallian knowledge spillovers from practicing in an area in which physicians have specialized in a particular style of practice. In their seminal paper on knowledge spillovers in medicine, Coleman, Katz and Menzel (1957) found that doctors who were the most integrated in a social-network were the first to adopt a new drug. More recently, a randomized control trial found that providing information to "opinion leaders" in a hospital resulted in large increases in the use of appropriate medications following heart attacks and decreases in the use of outdated therapies (Soumerai et al., 1998). Strong evidence of knowledge spillovers comes from Meltzer (2006), in which hospitalists (physicians specialize in hospital care, as opposed to a specific disease) were randomly assigned to physicians groups. Physicians in the randomized groups exhibited social network effects by learning from the experience of the hospitalists that they were randomized to.

An alternative mechanism generating productivity spillovers may operate through the availability of support services in area hospitals (cath labs, cardiac surgeons, cardiac care units, nurse staff), which will depend more on the overall practice style in the area rather than on any individual patient's needs. This second mechanism is analogous to the increased availability of software as more users adopt a given computer operating system, which has been emphasized in the literature on network externalities (see for example, Goolsbee and Klenow (2002)). Finally, productivity spillovers may occur through the matching process of physicians to areas, since physicians who are more skilled at a particular treatment may self-select into areas in which use of this treatment is more common. Note that this last mechanism has different welfare implications, since selection of physicians into one area imposes externalities on other areas. We will return to this point in our discussion of welfare implications.

A. The Model

More formally, conditional on the fraction of patients receiving the treatment, let the survival rate and cost associated with each treatment option take the simple form:

$$Survival_i = \beta_i^s Z + \alpha_i^s P_i + \varepsilon_i^s \quad \text{for } i = 1, 2. \quad (1)$$

³This is similar to how network externalities are modeled by Katz and Shapiro (1985, 1986). The social network literature commonly uses the fraction of the population adopting a particular practice as a proxy for information flows, as it captures the density of adopters in a network (c.f. Rogers, 2003). An alternative model would have the production externality depend on the number (rather than proportion) of patients in an area that receive the i th treatment. In this alternative model larger areas are better at the use of both interventions, an implication that we test and reject (see Table 8).

$$\text{Cost}_i = \beta_i^c Z + \alpha_i^c P_i + \varepsilon_i^c \text{ for } i = 1, 2. \quad (2)$$

The subscript i indexes the treatment option, and we omit the patient subscripts for convenience. We assume that each patient's indirect utility function (U) depends on both the price and quality of care, as captured by the expected survival rate and costs associated with each treatment option:

$$U_i = \text{Survival}_i - \lambda \text{Cost}_i = \beta_i Z + \alpha_i P_i + \varepsilon_i \text{ for } i = 1, 2. \\ \text{where } \beta_i = (\beta_i^s - \lambda \beta_i^c), \alpha_i = (\alpha_i^s - \lambda \alpha_i^c), \text{ and } \varepsilon_i = (\varepsilon_i^s - \lambda \varepsilon_i^c). \quad (3)$$

The parameter λ represents the value of life (survival per dollar), and captures the tradeoff being made by the patient and physician between improved survival and increased costs. A typical value of λ used in cost effectiveness studies would place the societal value of a life year at around \$100,000 (Cutler, 2004), although physicians and patients may use $\lambda=0$ if they do not pay directly for the cost of treatment (for example, because of insurance). The term $\beta_i Z$ represents an index of how appropriate a given patient is for each treatment based on medically relevant characteristics (Z). Patients who are more appropriate have greater utility from the treatment. The second term ($\alpha_i P_i$) captures the productivity spillover, which is positive if $\alpha_i > 0$: area specialization in a given treatment improves survival and/or reduces costs for all patients receiving that treatment. The final term (ε_i) represents factors that influence survival and costs and which are known to the patient (or physician) at the time of choosing a treatment option but unobserved to the econometrician.

An individual patient receives the intensive treatment ($i=2$) only if their utility is higher from the intensive than the non-intensive treatment ($U_2 > U_1$). Therefore, the probability that an individual patient receives the intensive treatment is given by:

$$\begin{aligned} \Pr\{\text{intensive treatment}\} &= \Pr\{i = 2\} = \Pr\{(U_2 - U_1) > 0\} \\ &= \Pr\{(\alpha_1 + \alpha_2)P_2 - \alpha_1 + (\beta_2 - \beta_1)Z > (\varepsilon_1 - \varepsilon_2)\} \\ &= \Pr\{\alpha P_2 - \alpha_1 + \beta Z > \varepsilon\}, \text{ where } \alpha = (\alpha_1 + \alpha_2), \beta = (\beta_2 - \beta_1), \varepsilon = (\varepsilon_1 - \varepsilon_2). \end{aligned} \quad (4)$$

Integrating equation 4 over the distribution of patients (Z) in the population yields the market demand curve for the intensive treatment.

Among patients that choose the intensive treatment, the expected utility gain is given by:

$$E[(U_2 - U_1) | (U_2 - U_1) > 0] = \beta Z + \alpha P_2 - \alpha_1 + E[\varepsilon | (U_2 - U_1) > 0]. \quad (5)$$

Thus, the expected utility gain among patients receiving the intensive treatment (the treatment on the treated) depends on a patient's relative appropriateness for the treatment, the proportion of all patients that receive the intensive treatment, and a term representing the selection effect. Because equation (5) has a simple tobit structure, the truncated mean will increase with the mean of the underlying variable for any log concave distribution for ε (Heckman and Honore, 1990). Thus, patients that choose the intensive treatment will have a higher expected utility gain from the treatment (in terms of higher survival or lower costs) if they are more appropriate (higher βZ) or in more intensive regions (higher αP_2).

B. Equilibrium

Equations (3)–(5) represent a standard Roy model except that, in equilibrium, P_2 must be equal to the demand for the intensive treatment according to equation (4). Letting $f(Z)$ represent the distribution of Z in the population, this implies the following equilibrium condition must hold:

$$P_2 = \int_Z \Pr(\alpha P_2 - \alpha_1 + \beta Z > \varepsilon) f(Z) dZ \\ \equiv G(P_2). \quad (6)$$

Thus, equilibrium in this model is defined as a solution to equation (6) (a fixed point). In other words, in equilibrium the proportion of patients choosing the intensive treatment must generate benefits to intensive treatment that are consistent with this proportion actually choosing the treatment. For example, if all patients choose the intensive treatment but this generates utility benefits such that some patients would prefer the non-intensive treatment, then rates of intensive treatment will decline. As fewer patients choose the intensive treatment, costs will rise and/or survival will fall (through the spillover mechanism), further lowering the benefits to intensive treatment. This process continues until we reach an equilibrium in which only the most appropriate patients are left receiving the intensive treatment. Thus, cost and quality of care determine demand for each treatment option, and adjust to equilibrate the market.

C. Empirical Implications

This simple model has a number of very strong empirical implications. First, as illustrated in Figure 1, Equation (6) can generate variation across areas in the proportion of patients receiving the intensive treatment for two different reasons. For a given distribution of patients (which fixes the function G), there can be multiple equilibria because of the productivity spillovers. For example, panel A illustrates a situation with two equilibria: An intensive equilibrium in which most patients receive intensive treatment and the returns to this treatment are high, and a non-intensive equilibrium in which few patients receive intensive treatment and the returns to this treatment are low (the middle crossing point represents an unstable equilibrium). Alternatively, differences across areas in the distribution of patients (which changes the function G) will also lead to different equilibria. For example, panel B illustrates a situation in which most of the patients in one area are more appropriate for intensive treatment leading to an intensive equilibrium, while most of the patients in the non-intensive area are not appropriate for intensive treatment. Even in the single equilibrium scenario the productivity spillover increases the differences across areas: Having more patients appropriate for intensive treatment in an area (a shift upward in G) increases the return to intensive treatment for other patients, which in turn leads to a further increase in intensive treatment (a move to the right along G). Thus, small differences in the distribution of patients can potentially generate large equilibrium differences in specialization across areas.

A second empirical implication of the model that follows immediately is that identical patients will be more likely to be treated intensively in an area where more patients are appropriate for intensive management of the heart-attack (based on Z). Any shift in the distribution of patients (Z) that results in a higher proportion being treated intensively (P_2) results in a higher probability of choosing intensive treatment for each patient as a direct consequence of Equation (4). This implication is similar to that tested by Goolsbee and Klenow (2002) for individuals' purchase of home computers, and is a unique implication of spillovers: Preferences of the population have spillovers on the choice of an individual. While this implication of the model is always true in the single equilibrium case, it may not hold in the multiple equilibrium case if a shift in the distribution of Z affects the choice among the equilibrium. Unfortunately, our static model is silent on what determines the choice among multiple equilibria.⁴

A number of additional implications of the model are illustrated in Figure 2. For simplicity, we ignore unobserved patient characteristics (ε) in this figure in order to focus on the intuition of the model; under standard regularity assumptions about ε (see Heckman and Honore

⁴Arthur (1989) emphasizes the importance of "lock-in" by historical events. In our application, two regions might have differed in terms of their distribution of initial patient types. This difference is what caused the choice of present specialization.

(1990)) this results in no loss of generality, but significantly simplifies the exposition. The top panel plots patient utility (survival net of costs) as a function of a patient's appropriateness for intensive management (which depends on Z) for each of the two treatment options (intensive and non-intensive). Patients to the left of the intersection of these two curves are treated non-intensively, while patients to the right are treated intensively. Patients further to the right are both more likely to receive intensive treatment, and experience higher returns to the treatment if treated. Thus, the model predicts that the return to intensive treatment (treatment on the treated) is highest among patients with the highest probability of treatment. This must occur through a combination of larger survival benefits or lower treatment costs among patients with the highest probability of treatment (the "more appropriate" patients).⁵

The bottom panel of Figure 2 illustrates four additional implications of the model by comparing patient utility in intensive and non-intensive areas. First, the utility associated with non-intensive management is worse in areas that are intensive. This suggests that there should be a negative relationship between the proportion of patients that receive surgically intensive treatment and indicators of the quality of care for patients treated medically – i.e., intensive surgical treatment crowds out good medical management. A second implication of the model follows directly: In intensive areas patient utility will be higher among patients that are most appropriate for intensive management, but lower among patients that are more appropriate for medical management. In other words, the area's specialization in intensive management helps patients that are appropriate for this type of care through higher survival or lower costs, but patients who require less intensive management are harmed by the area's specialization in intensive management. For example, to the extent that older age is a crude proxy for requiring less intensive management, our model predicts that very-old patients will have lower survival and higher costs in an intensive area, while those who are younger benefit through higher survival and lower costs from their areas' specialization in intensive management. On net, however, average benefits across all patients may be higher or lower in surgically intensive areas.

A third implication of our model is that the marginal patients receiving the intensive treatment in intensive areas will be less appropriate for the treatment than the average patient receiving the intensive treatment. As shown in the top panel of Figure 2, patients are given the intensive surgical treatment if their appropriateness is above the point where the non-intensive and intensive utility curves intersect. As seen in the bottom panel of Figure 2, this intersection is at a lower appropriateness level in intensive areas. Therefore, the additional patients receiving the intensive treatment in these intensive areas are less appropriate than the patients receiving the intensive treatment in non-intensive areas.

A final implication of our model is that among those receiving intensive treatment, the benefit to receiving intensive treatment in the intensive area is larger than the benefit in the non-intensive area. In other words, the treatment effect on the treated will be larger in more intensive areas. As can be seen in the bottom panel of Figure 2, this higher return is the net result of a higher benefit if patients are treated intensively and a lower benefit if patients are treated non-intensively in the intensive areas. Thus, the high return to intensive treatment in intensive areas is the result of both a positive productivity externality on intensive treatment and a negative productivity externality on medical treatment.

⁵More precisely, our model makes unambiguous predictions about the utility gains to a patient receiving the intensive treatment ($U_2 - U_1$). Utility can increase in three ways: (1) survival increases and/or cost declines, (2) survival increases that outweigh cost increases, or (3) cost declines that outweigh survival declines. Our discussion focuses on the first case as the most relevant, but the others are possible.

All of the model's implications are consistent with the key empirical regularities summarized in the introduction. In particular, our productivity spillovers model can generate: (1) substantial differences across areas in the use of intensive procedures that is unrelated to average patient outcomes, (2) a negative correlation between surgical intensity and the quality of medical management of a condition, and (3) large returns to receiving the intensive intervention, particularly in high-intensity areas. Thus, a very simple equilibrium model of productivity spillovers can rationalize all of the main stylized facts in the literature.

Some of the empirical implications of our model are the result of productivity differences across areas, rather than spillovers *per se*. In our model, productivity differences across areas arise solely through the productivity spillover term ($\alpha_i P_i$) in equation (3). But there may be other reasons for productivity differences across areas that are independent of the fraction of patients receiving each treatment. If this were the case, then areas that happened to be relatively good at the intensive treatment (with high patient utility relative to non-intensive treatment) would have high returns to intensive treatment and also be more likely to use the intensive treatment for any given patient – just as predicted by our productivity spillover model.

However, there are two empirical implications of our model that depend on productivity spillovers. First, spillovers imply that a patient is more likely to be treated intensively in an area where other patients are more appropriate for intensive treatment (based on Z). This implication would not arise without a direct link between productivity and the fraction of patients being treated intensively. Second, spillovers imply that specialization in the intensive treatment must help patients that are most appropriate for intensive treatment, but at the same time must harm patients who are least appropriate by reducing specialization in the non-intensive treatment. This implication would not necessarily arise without the negative relationship between the productivity of intensive and non-intensive treatments that is implied by spillovers. For example, suppose that productivity in both the intensive and non-intensive treatments was driven by a latent quality factor in the area, but that this latent quality had a larger impact on productivity for the more complex intensive treatment. Then areas with higher latent quality would treat more patients intensively, but all patients in these areas would have better outcomes.

Most of the empirical implications of our model would not be predicted by the standard flat-of-the-curve explanation for area variation in treatment rates. In our modeling framework, the flat-of-the-curve argument has no spillovers ($\alpha=0$). Instead, areas vary in terms of the tradeoff they make between survival and costs (λ varies): some areas are willing to incur greater cost per life saved (lower λ), perhaps because they have higher income, have excess hospital capacity, or are paid on a cost-plus (e.g. fee-for-service) basis. In this flat-of-the-curve model, the areas that have the lowest λ place the least weight on costs in making treatment decisions and, therefore, are more likely to use the intensive (and more costly) treatment for patients who receive little survival benefit. But such a model predicts that the average survival benefit of intensive treatment will be *lower* for patients receiving the treatment in high-intensity areas, while the incremental costs will be *higher*. That is, the flat-of-the-curve story fundamentally implies that high-intensity areas simply choose to treat additional patients who yield relatively few lives saved per dollar. Moreover, since externalities play no role in the flat-of-the-curve argument, this model does not imply a lower quality of medical management or worse outcomes among patients least appropriate for intensive treatment in high-intensity areas. Finally, the appropriateness of other patients in the area play no role in determining treatment of a given patient in the flat-of-the-curve model, while this is a fundamental feature of models with productivity spillovers. Instead, the flat-of-the-curve model suggests that area factors such as income and hospital reimbursement structure will determine the area intensity.

D. Welfare Implications

Any change in the fraction of patients receiving the intensive treatment has positive externalities on some patients and negative externalities on others. Therefore, pareto improvements over any equilibrium are not possible in our model without side payments from winners to losers. We focus instead on whether any deviation from a given equilibrium will result in higher average utility in the population, implicitly assuming that either appropriate side payments could be made or that the societal objective is to maximize average utility. Finally, for this to be an appropriate measure of welfare, we must assume that all spillovers are captured within the model and, in particular, that specialization in one market has no impact on other markets (e.g., the spillover mechanism is not working through attracting quality physicians from other areas).

Because each patient's choice of treatment affects the outcomes of all other patients, there is no reason to expect that equilibrium in our model is optimal. Let P^* be the proportion of patients receiving the intensive treatment in equilibrium (a solution to equation (6)). Small increases in P^* away from its equilibrium value, accomplished by changing the treatment for the marginal patient, will not affect the marginal patient who is indifferent between the two treatments in equilibrium. Thus, the effect of such an increase in P^* on expected utility depends solely on the externality it has on other patients – the positive externality on patients receiving the intensive treatment, balanced by the negative externality on patients receiving the non-intensive treatment. Let U be the average utility in the population given the proportion of patients receiving the intensive treatment (P_2), so that $U = \int_Z \max(U_1, U_2 | P_2 = P) f(Z) dZ$. Then it is straightforward to show that:

$$\left. \frac{\partial U}{\partial P_2} \right|_{P_2=P^*} = \alpha_2 P^* - \alpha_1 (1 - P^*) = -\alpha_1 + \alpha P^*, \text{ where } \alpha = (\alpha_1 + \alpha_2) \quad (7)$$

Equation (7) is quite intuitive. The impact on expected utility of a small increase in intensive treatment away from its equilibrium value is simply the weighted average of the positive externality (α_2) on intensive treatment and the negative externality ($-\alpha_1$) on non-intensive treatment, where the weight is the proportion of patients receiving the intensive treatment in equilibrium. For $P^* > \alpha_2/(\alpha_1 + \alpha_2)$, increasing intensive treatment above its equilibrium value will have a positive effect on expected utility, while for $P^* < \alpha_2/(\alpha_1 + \alpha_2)$ the opposite is true. Thus, expected utility of patients is not in general maximized in equilibrium.

Note that the right hand side of equation (7) is exactly the same as the term representing the impact of spillovers on patient choice of treatment in equation (4). Thus, if an area does more intensive treatment in equilibrium than it would in the absence of spillovers (implying the right hand side of equation (7) is positive), then there are additional welfare gains from doing even more intensive treatment. Similarly, if an area does less intensive treatment in equilibrium because of spillovers, then there are additional welfare gains from doing even less. From a welfare perspective, there will be too little area variation in treatment rates if the marginal person ignores the externality that their treatment decision has on others. Said a different way, if physicians maximize social welfare (rather than patient welfare) and take this externality into account in making treatment decisions, then they will optimally do some harm (in a utility sense) to the marginal patient because of the externality on other patients, and the result will be more variation across areas in specialization.

In the case of multiple equilibria, the equilibria can be ranked in terms of expected utility. Thus, shifts from one equilibrium to another can improve expected utility. Unfortunately, the model has no general implications regarding which equilibrium will have the highest expected utility (and recall that no equilibrium is pareto dominant without side payments). However, if all areas

had identical patient distributions (Z) and the variation in treatment rates across areas arose from multiple equilibrium, then one could compare patient survival and costs across areas to determine the optimal equilibrium rate of intensive treatment. This was the approach taken by Fisher et al. (2003b), who argue that a 30 percent reduction in Medicare spending (such that spending in high intensity regions is reduced to that of low intensity areas) would not have any deleterious effects on patient outcomes or satisfaction. Our model suggests that this conclusion is justified only if the observed variation represents multiple equilibria for regions with identical patient distributions. In the single equilibrium case our model suggests the opposite: when equilibrium treatment rates differ across regions because of differences in the underlying patient distribution, then expected patient utility in more aggressive regions would be raised by increasing – not decreasing – intensive treatment rates above their equilibrium values.

III. Data and Estimation

Data

We focus our empirical work on the treatment of AMI for four reasons. First, cardiovascular disease, of which AMI is the primary manifestation, is the leading cause of death in the US. Because mortality post-AMI is high (survival rates at one year are less than 70 percent), a well-defined endpoint is available to test the efficacy of alternative interventions. Second, markets for heart attack treatment are geographically distinct. The acute nature of this condition requires immediate treatment and generally precludes a patient from traveling long distances to seek care. Therefore, mobility is limited and it is possible to observe production in many distinct local markets. Third, the treatment of heart attacks is characterized by the choice between two competing interventions, both of which restore blood flow to the coronary arteries—intensive management characterized by the use of angioplasty or bypass, and non-intensive management characterized by the use of fibrinolytics therapy (also known as thrombolytics), that dissolve clots that may have formed in a blood vessel. We measure the use of intensive therapy by focusing on the use of cardiac catheterization since it is a well-understood marker for surgically intensive management of patients (our decision follows McClellan et.al (1994) and McClellan and Newhouse (1997)).⁶ Patients who received bypass or angioplasty are included in the set of persons receiving catheterization, and therefore, intensive treatment. Fourth, some geographic areas specialize in the intensive management of heart attacks, while others specialize in non-intensive management, but unlike other markets, neither production choice can completely dominate a market because some patients are always more appropriate for a particular intervention (for example, 95 year old patients do not benefit from surgery, and must be treated non-invasively). Thus, the productivity of both intensive and non-intensive management can be observed in all markets.

Regardless of whether a patient is treated intensively or using thrombolytics, patients should also be prescribed beta-blockers during their hospital stay that reduce the uptake of adrenalin and thereby slow the heart. Beta-blockers have been shown to improve outcomes for the majority of patients, their use is substantially below what most experts believe is appropriate, and the rate of beta-blocker use among AMI patients is a widely used measure of the quality of medical care [Jencks (2003), and Baicker and Chandra (2004a)]. Because beta-blockers are a form of medical management of patients, their use serves as a marker of the quality of non-intensive medical management in an area.

⁶In performing cardiac catheterization, a cardiologist inserts a thin plastic tube (catheter) into an artery or vein in the arm or leg, from where it is advanced into the chambers of the heart and into the coronary arteries. The catheter measures the blood's oxygen saturation and blood pressure within the heart. It is also used to get information about the pumping ability of the heart muscle. Catheters are also used to inject dye into the coronary arteries which can then be imaged to assess arterial stenosis using an x-ray. Catheters with a balloon on the tip that is inflated in order to compress the atherosclerotic plaque to improve blood flow are referred to as angioplasty.

Because acute myocardial infarction is both common and serious, it has been the topic of intense scientific and clinical interest. One effort to incorporate evidence-based practice guidelines into the care of heart attack patients, begun in 1992, is the Health Care Financing Administration's Health Care Quality Improvement Initiative Cooperative Cardiovascular Project (CCP). The CCP developed quality indicators that were based heavily on clinical practice guidelines developed by the American College of Cardiology and the American Heart Association. Information about more than 200,000 patients admitted to hospitals for treatment of heart attacks was obtained from clinical records. The CCP is considerably superior to administrative data (of the type used by McClellan et.al (1994)) as it collects chart data on the patients—detailed information is provided on laboratory tests, the location of the myocardial infarction, and the condition of the patient at the time of admission. In the data Appendix we provide a detailed account of the estimation sample used in this paper.

Defining Geography

To construct local markets for health care we exploit insights from The *Dartmouth Atlas of Health Care*. The Atlas divides the U.S. into 306 Hospital Referral Regions (HRR), with each region determined at the zip code level by the use of an algorithm reflecting commuting patterns and the location of major referral hospitals. The regions may cross state and county borders, because they are determined solely by migration patterns of patients. For example, the Evansville, Indiana hospital referral region encompasses parts of three states because it draws patients so heavily from Illinois and Kentucky. HRRs are best viewed as the level where tertiary services such as cardiac surgery are received (although they are not necessarily the appropriate geographical level for primary care services).

Analysis at the HRR level is preferable to analysis at the city, or state level since it uses the empirical pattern of patient commuting to determine the geographic boundaries of each referral region, rather than assuming that the arbitrary political boundaries of states and cities also define the level at which the health care is delivered. Furthermore, for the purpose of studying geographic productivity spillovers, an analysis at the HRR level is superior to one at the level of individual hospital for two reasons. First, patients can be assigned to HRR based on their residence, rather than based on the hospital at which they received treatment (which may be endogenous). In addition, geographic productivity spillovers are likely to operate at a broader level than that of a given hospital, e.g. these spillovers are expected to reach beyond the boundary of the firm to affect productivity at all firms in a region. Physicians often have operating privileges in multiple hospitals and interact (socially and professionally) with other doctors who may or may not practice in their hospital, and patients are commonly referred to other hospitals within the HRR for treatment. Through such interactions, the entire area (as measured by HRRs) is more appropriate for analyzing equilibrium implications of productivity spillovers.⁷

Estimation

The key estimating equations in this paper take the following form:

$$Outcome_{ijk} = \beta_{0k} + \beta_{1k}Intensive\ Treatment_i + X_i\Pi_k + u_{ijk} \quad (8)$$

⁷As an alternative to the use of HRRs we have also re-estimated our analysis using the U.S. Census Bureau's "Metropolitan Statistical Area" (MSA) designation of urban areas. In general, HRRs may be thought of encompassing the relevant MSA with the addition of surrounding rural areas whose patients travel to the MSA to seek care. Using the MSA sample results in a loss of sample size, and introduces some noise into the IV analysis since within an MSA (such as Manhattan or Boston) there is much less variation in differential distance. The results reported in this paper are robust to the use of MSAs, with the exception of the IV estimates which, while not statistically different from what we report, are very imprecise.

Here, $Outcome_{ijk}$ refers to either survival or costs for patient i in HRR j ; k is an indicator variable that indexes the different groups for whom the effect of intensive management is sought (e.g., patients appropriate for intensive management versus not, or alternatively, high-intensity HRRs versus low intensity HRRs). Following McClellan et al. (1994) and McClellan and Newhouse (1996), *Intensive Treatment* is measured by the receipt of cardiac catheterization (including angioplasty or bypass) within 30 days of the AMI. Alternatively, in some specifications we use total spending (Medicare Part A and Part B charges) in the first year after the AMI as a proxy for intensive treatment. The vector X_i includes the entire set of CCP controls listed in the data appendix. These variables provide a relatively complete summary of the patient's condition at the time of admission, i.e. include all of the relevant clinical information available to the physician in the patient's chart at the time of the heart attack. Our tables report estimates of β_{1k} and their difference between different sub-samples of the data; the latter is the central parameter of interest for the purpose of testing our model. Standard errors are clustered at the level of each area.

Both our model and common sense suggest that the choice of intensive treatment for a patient will be endogenous. Even though we have excellent information on the patient's clinical condition at admission, the attending physician or cardiologist is likely to make the treatment choice based on information that is not observable in the CCP (for example, using information observed in the weeks following the initial admission). In particular, the selection problem that confounds OLS estimation of the above equation is that intensive treatment is recommended to patients who will benefit most, and these patients are typically in better health (e.g. did not die in the first 24 hours after the heart attack). This selection of healthy patients into treatment biases OLS estimates toward finding a large effect of intensive treatment. We follow the work of McClellan et.al (1994) and estimate equation (8) using instrumental variables. In particular, we use differential distance (measured as the distance between the patient's zip-code of residence and the nearest catheterization hospital minus the distance to the nearest non-cath hospital) as an instrument for intensive treatment, with a negative value of differential distance indicating that the nearest hospital is a cath hospital. We capped differential distance at ± 25 miles based on preliminary analysis that suggested little effect of differential distances beyond 25 miles on the probability that a patient receives catheterization.

In contrast to regressions where we study an individual's receipt of CATH (and where an instrumental variable is clearly indicated), we use Ordinary Least Squares (OLS) for tests of our model that require the use of area CATH rates as the key explanatory variable. Observed variation in area CATH rates would have to be driven by large differences in unobservable health status across areas to bias our results. We doubt that this is true, given the prior literature documenting that population health status is not correlated with intensive treatment at the area level (e.g. Fisher et al., 2003a, b). But even if it were, we would be biased toward observing a positive association between CATH rates and survival, which we do not.

To define a patient's clinical appropriateness for intensive management we estimate a logistic regression model for the probability of receiving cardiac catheterization within 30 days of the heart attack. Specifically we estimate:

$$Pr(\text{Cardiac Cath}_{ij}) = G(\theta_0 + \theta_j + X_i\Phi) \quad (9)$$

Here, θ_j is the risk-adjusted logistic index for the use of cardiac catheterization in HRR j ($j=1, 2, \dots, 306$). This equation is analogous to equation (4) in our model, with the HRR fixed effects capturing the externality that causes similar patients to be treated differently across areas. Fitted values from this regression that exclude the HRR fixed effect, $Pr(\text{Cardiac Cath}_{ij}) = \hat{G}(\theta_0 + X_i\Phi)$, are used as an empirical measure of clinical appropriateness for cardiac catheterization.

⁸ For simplicity, in much of the empirical work we split the fitted values at their median to

yield two equally sized groups; those above the median are appropriate for intensive management, and those below are not. In order to classify areas as being intensive versus not, we construct HRR level (risk-adjusted) catheterization rates by recovering the θ_j from the equation above.⁹

IV. Empirical Results

Our empirical results are organized into two sections. We begin by testing the basic implications of the Roy model – that patients are sorted into surgery versus medical management based on the returns to each type of treatment. We then turn to testing the model's implications that depend on the presence of productivity spillovers.

A. Testing Implications of the Basic Roy Model

Table 1 presents IV estimates of the impact of CATH (as a marker for intensive treatment) on 1-year survival and 1-year costs (in thousands of 1996 dollars). In the first panel of Table 1 we present the analysis for all patients, and then the remaining panels estimate separate regressions by the clinical appropriateness of the patient. In the third column we instrument costs with differential-distance to estimate the cost-effectiveness of spending on survival. The results in Table 1 suggest that the returns to more intensive management of patients are increasing in clinical appropriateness, as predicted by our model. Patients appropriate for catheterization who received the intervention saw an 18.4 percentage point survival gain relative to appropriate patients who did not receive intensive therapy. For the less appropriate group, the effect of receiving catheterization is not statistically different from zero.¹⁰ Column 2 of the table demonstrates that it also costs more to treat a patient who is clinically inappropriate for intensive management aggressively. One reason of this finding is the risk of iatrogenic (physician induced) complications that require additional rehabilitation days in the ICU or CCU. Column 3 uses total dollars spent on a patient as a measure of intensity. By this measure, increasing intensity (by spending an additional \$1000 on a heart-attack patient) raises survival by 3.8 percentage points for patients who are appropriate for intensive management, but by nothing for patients who are not good candidates for intensive management. The lower panel of Table 1 provides an alternative breakdown of patients into appropriateness groups based on age (65–80 versus over 80). Clinical guidelines recommend that older patients not be treated surgically, suggesting that the returns to such treatment should be negligible for the over 80 group. The last panel of Table 1 confirms that this is indeed the case. Together, the results in Table 1 provide support for our Roy-model characterization of specialization—patients most appropriate for intensive treatment benefit most from such therapy, whereas those least appropriate for it do not benefit from it.

Table 2 explores the validity of differential distance as an instrument in our sample. Following McClellan et al. (1994), we split the sample in half and compare average characteristics of the

⁸Our model suggests that omitting the HRR fixed effects will bias the coefficients on the comorbidities (X) because places with more appropriate patients will tend to do more CATH. However, in practice this choice is empirically irrelevant for our estimates. The correlation between our preferred measure of appropriateness (using all the CCP risk adjusters and HRR fixed effects) and one that excludes the HRR fixed effects is 0.997 (at the patient level). Similarly, the correlation between our preferred measure and simpler measures using only age-gender-race interactions is 0.75.

⁹To verify that more and less aggressive areas rank patients similarly, we estimated separate logistic regression models (by whether an area's risk-adjusted CATH rate was less than or greater than the median area's rate), of a patient's probability of receiving CATH on the all the CCP risk adjusters. For every patient in the data, we obtained predicted values from the two models and noted that their correlation was 0.986. This high correlation suggests that areas rank people similarly.

¹⁰Following suggestions made by our anonymous reviewers, we note that our results are robust to controlling for distance to the nearest hospital, as well as controlling for the distance to the nearest hospital and rural classification of the zip code of the residence. These results are available upon request. Briefly, after controlling for actual distance to the nearest hospital, the coefficient (and accompanying standard-error) on survival for appropriate patients was 0.182 (.033); for those less appropriate it was 0.051(.081). After, controlling for distance and rural location, these coefficients changed to 0.181(.046) and .084(.13) respectively.

sample above and below the median differential distance (−2.0 miles). The first two columns show that among all patients, there is a 6.1 percentage point difference in the CATH rate between the samples above and below the median, with higher differential distance to a CATH hospital associated with lower rates of CATH. A similar pattern is seen when the sample is restricted to patients who are most or least appropriate for CATH based on their propensity score or their age. Even among the patients with low propensity scores or over age 80, there is a 3–4 percentage-point decline associated with being further from a CATH hospital. These differences are all highly significant, even after controlling for the full set of patient controls from the CCP (the first-stage F-statistics on differential distance are over 50 for all specifications reported in Table 1).

The third and fourth columns of Table 2 report survival rates for patients above and below the median differential distance. If differential distance is unrelated to patient mortality risk, then these estimates can be combined with the difference in CATH rates from the first two columns to form a simple Wald estimate of the effect of CATH on survival.¹¹ Among all patients, there is a 0.9 percentage point decline in survival associated with being further from a CATH hospital, suggesting that the intensive treatment associated with CATH leads to improved survival. The Wald estimate suggests that CATH is associated with a 15 percentage-point improvement in survival (.009 gain in survival/.061 increase in CATH), an estimate quite close to the 14.2 percent estimate from Table 1 that controlled for the full set of patient controls. The remaining rows show results for other patient samples that are similarly consistent with the Table 1 estimates, with a larger association between differential distance and survival seen among patients that are more appropriate for CATH and no relationship seen among patients that are less appropriate for CATH. Given the similarity of the Wald estimates to IV estimates with a full set of patient controls, it appears that differential distance is uncorrelated with observable differences in mortality risk. The fifth and sixth columns show that, as expected, there is little difference in the average 1-year predicted survival rate (from a logit of survival on the patient controls) for patients above and below the median differential distance.

The final two columns of Table 2 compare average 30-day predicted CATH rates (the propensity to get CATH) for only those patients getting CATH in the areas above and below median differential distance. If the additional patients getting CATH in the low differential distance sample were less appropriate for CATH, we would expect to see that the average patient getting CATH in these areas would have a lower propensity. In contrast, we see little difference in the sample that is nearer to a CATH hospital. Thus, it appears that differential distance is an instrument that increases CATH rates among a sample of patients that is very similar (at least on observable factors) to the average patient being treated. Therefore, it appears reasonable to interpret the IV coefficients as estimates of the treatment effect in the treated population.

Finally, the Roy model predicts that as areas increase their intensity the additional patients receiving the intensive management should be less clinically appropriate for the intervention. In contrast, a non-economic model would predict that intensive areas simply perform more intensive treatments on all patients; there is no triage of patients into the treatment based on their appropriateness. To test this insight we estimate:

$$\text{Appropriateness}_{ij} = \mu_0 + \mu_1 \ln(\text{Cath Rate})_j + \epsilon_i \quad (10)$$

The dependent variable is a measure of the appropriateness of a patient for cardiac cath (e.g., the propensity score), and the equation is estimated at the patient level *only for individuals*

¹¹Further evidence that the differential distance instrument is uncorrelated with other determinants of receiving intensive treatment are provided in McClellan et al. (1994), McClellan and Newhouse (1997) and McClellan and Noguchi (1998).

receiving CATH. The explanatory variable of interest is the natural logarithm of the risk-adjusted area CATH rate. Following the logic of Gruber, Levine and Staiger (1999), the coefficient μ_1 measures the difference in mean Appropriateness between the average and the marginal patient receiving the invasive treatment. If average appropriateness among patients getting CATH is declining as the CATH rate rises ($\mu_1 < 0$), we infer that the marginal patient getting CATH was less appropriate than the average patient (analogously to deriving marginal cost curves from average cost curves).

The results of estimating equation (10) are reported in Table 3. As the area's intensity increases, the marginal patient is 4.5 percentage points less likely to be appropriate for intensive treatments. As a robustness check, we included two alternative measures of clinical appropriateness: age and a measure of clinical *ineligibility* for CATH as defined by the American College of Cardiologists (AHA) and the American Heart Association (AHA). Table 3 indicates that by both measures, the marginal patient is significantly less likely to be appropriate for cardiac catheterization.¹²

Figure 3 provides a graphical illustration of this relationship. For each of the 306 HRRs we graph the average propensity to receive cardiac catheterization (amongst patients who actually received it) against the log of the area risk-adjusted CATH rate. Using local-regression, we estimated the relationship between the average propensity and the risk-adjusted CATH rate and the slope of this line at each point (which we also smoothed). These estimates were then used to plot average and marginal patient receiving treatment. As seen in the figure, average appropriateness of patients getting CATH declines in more intensive areas. The average appropriateness can decline only if the marginal patient is less appropriate— as the lower line in Figure 3 indicates, the appropriateness of the marginal patient appears to be below the average patient and also declines as areas become more surgically intensive.

B. Testing Implications of Productivity Spillovers

The empirical results in Tables 1–3 provide support for the basic assumptions of the Roy model – that patients are sorted into surgery versus medical management based on the returns to each type of treatment. The remaining tables focus on testing for the presence of productivity spillovers. Table 4 tests the prediction that the quality of medical management is worse in areas that are surgically intensive. To measure the extent of intensive treatment in an area, we use the risk-adjusted 30-day CATH rate in each HRR. The risk-adjusted CATH rate reflects variation across areas in the probability that observationally identical patients will receive CATH, and therefore measures the productivity externality through its role of increasing the probability of receiving the intensive treatment. To measure the quality of medical care in an area, we use the risk-adjusted rate at which patients received beta-blockers in the HRR. Use of beta-blockers is a widely used marker for the quality of medical care. Not surprisingly, the risk-adjusted CATH rate is positively correlated with other risk-adjusted rates of intensive surgical treatment such as bypass (CABG) and angioplasty (PTCA). More importantly, the negative correlation between the risk-adjusted CATH and beta-blocker rates supports the view that the quality of medical management is worse in surgically intensive areas.

The remaining rows of Table 4 report the correlation between risk-adjusted CATH rates and other area-level characteristics of interest. Risk-adjusted CATH rates are positively associated

¹²Both measures are included as covariates in our estimation of the empirical appropriateness for cardiac catheterization. As such, these results are not three separate confirmations of the same prediction. The ACC/AHA measures are based on an evaluation of the patients chart characteristics and classify each patient as being ideal, appropriate, and inappropriate for catheterization; see Scanlon et.al (1999) for details. This measure, however, is not our preferred measure because, as others have noted, expert panels have been shown to exhibit enormous variability, particularly for procedure use classified as inappropriate, and to be greatly influenced by the composition of the panel (Shekelle et.al (1998)).

with cardiovascular surgeons per capita (physicians who perform cardiac surgery), and the number of CATH labs per capita. These correlations are consistent with the view that a higher level of support services available in high-intensity areas may contribute to the externalities. The last three rows of this table demonstrate, perhaps surprisingly, that risk-adjusted CATH rates are not strongly associated with demographic characteristics of the HRR such as population, income, or education.

A more fundamental implication of productivity spillovers is that the characteristics of other patients in the population will influence the treatment choice of an individual: all else equal, a patient will be more likely to receive a CATH if she lives in an area where the *average* patient is more appropriate for CATH. In the first column of Table 5 we regress a patient's receipt of CATH on the average propensity to get CATH in each HRR and patient risk-adjusters. The coefficient is statistically significant and implies that for every 1 percentage point increase in the expected area CATH rate (based on patient characteristics), there is an additional 0.53 percentage point increase in the probability that an individual receives a CATH because of spillovers. In the second column we control for area demographics which could potentially confound this relationship, but they are insignificant and do not materially change the coefficient.

A key implication of our model (and any model in which productivity differences across areas generate specialization) is that the return to intensive management should be higher in high-intensity areas versus low intensity ones. This prediction provides a sharp test of our model against the flat-of-the-curve model, which predicts the opposite. To test this prediction, we use our estimates of HRR level intensity from the estimation of equation (9) to classify patients as being treated in high or low intensity regions (as measured by whether the risk-adjusted CATH rate is above or below the median rate). In Table 6, we report IV estimates of the effect of receiving intensive management within each of these two groups. The results of this analysis are reported in a manner that mirrors the estimates in Table 1. The survival returns to intensive management in intensive areas is roughly three times the return observed in low intensity areas, while there is no statistically significant difference in the costs associated with the different areas.¹³ In the second panel of Table 6, we further split the sample within high and low intensity regions according to whether patients were more or less appropriate for CATH (based on a CATH propensity above or below the median). As predicted by our model, the most appropriate patients in high intensity areas have the highest survival returns (and lowest costs) associated with intensive management, while the lowest survival returns (and highest costs) are seen among the less appropriate patients in low intensity areas.

Table 7 tests a unique implication of productivity spillovers: Patients most appropriate for intensive treatments are better off being treated in high-intensity areas, while patients who are least appropriate for intensive treatments should be worse off in such areas. To test this prediction we split patients into 3 equal-sized groups of appropriateness for intensive treatments. Within each group, we report the relationship between survival and the area CATH rate in the first column. We also report the relationship between spending and the area rate in the second column. The first panel of Table 7 replicates the central finding from the area-variations literature—area intensity is associated with costs but not associated with improved outcomes among patients as a whole. However, as the second panel demonstrates this finding masks significant heterogeneity by patient appropriateness. Patients appropriate for intensive managements clearly benefit from being treated in intensive areas. However, as the productivity externality predicts, patients least appropriate for intensive treatments are harmed

¹³Analogous to note 10, controlling for distance to the nearest hospital barely changed our coefficients. In these specifications, the IV return to CATH for more and less intensive areas were 0.264 (0.059) and 0.094 (0.058) respectively. Further controlling for rural designation of a patients zipcode yielded estimates of 0.301(0.098) and 0.111(0.077) respectively.

as a result of being treated in intensive areas. The size of the negative externality is large—increasing area CATH intensity by ten percentage points (0.1) would reduce the survival of patients least appropriate for CATH by 0.75 percentage points. This finding holds true (although is less significant) when we split the data by alternative indicators of appropriateness for intensive treatment, including age and ACC/AHA guidelines.¹⁴

Our model suggests that patients who are inappropriate for CATH are worse off in intensive areas because (1) the quality of medical management in these areas is worse than in other areas, and (2) few of these patients receive intensive treatment, even in the more intensive areas. The last two columns of Table 7 explore these two dimensions of care directly by estimating the relationship between area CATH intensity and the rate at which patients receive beta blockers (as an indirect marker of quality of medical management) and CATH (as a direct marker of intensive management). Beta-blocker use is lower among all patient groups in the high intensity areas, suggesting that quality of medical management is generally poor in these areas. At the same time, CATH rates among those patients least appropriate for CATH rise much less in high intensity areas than for other patients. Thus, worse outcomes for this group appear to result from worse medical management in areas with high CATH intensity, as our model predicts.

Finally, our analysis has assumed that productivity spillovers depend on the proportion of patients receiving a given treatment, rather than the absolute number. Alternatively, productivity spillovers could be more prevalent in larger areas; larger HRRs such as Los Angeles or Manhattan may excel at both the intensive and non-intensive delivery of care. To explore this hypothesis further, we estimate the relationship between alternative measures of HRR size and patient survival in Table 8. We regress 1-year survival on patient risk-adjusters, the risk-adjusted CATH rate and the log of the resident population (first four columns) and log of AMIs per hospital (last four columns). Table 8 demonstrates that the size measures are largely insignificant—larger areas do not result in improved survival. On the other hand, the “externality” (as measured by the HRR CATH rate) is always protective of patients who are appropriate for CATH and harmful for inappropriate patients.

V. Conclusion

A very simple equilibrium model of patient treatment choice with geographic productivity spillovers appears to rationalize the main stylized facts concerning variation across areas in the use of technologically intensive medical care. The model yields a range of additional empirical implications, for which we found uniform support in our analysis of treatment for heart attacks. Alternative models, such as those based on “flat-of-the-curve” medicine, have fundamentally different implications that are clearly rejected by the data. Thus, there appears to be strong empirical support for the presence of productivity spillovers in medical care.

Our findings suggest that productivity spillovers play an important role in explaining geographic specialization in production, as first proposed by Marshall (1890). A large medical literature has documented the important role of social networks in physician adoption of new technologies, suggesting that knowledge externalities are the source of the productivity spillovers. While knowledge spillovers are the most natural interpretation of our model and empirical results, alternative mechanisms, such as the migration of specialized inputs into regions that favor one type of production, could generate the productivity spillovers we observe. These alternative mechanisms share the key feature that specialization in one sector

¹⁴In results available from the authors, we have verified that our results are similar (though less precise) in simpler specifications where we replace the risk-adjusted HRR rate with the difference between actual and expected rates based on only age, race and gender. For patients under age 80, the coefficient on the alternatively defined HRR rate is 0.021 (0.029), whereas for those over the age of 80 it is -0.041(0.038).

improves productivity in that sector while reducing productivity elsewhere, thereby reinforcing the tendency to specialize.

The presence of productivity spillovers may justify a policy intervention to ameliorate the deleterious effects of the externality. However, evaluating the welfare implications of any proposed intervention is not a trivial exercise. With multiple equilibria we may be able to identify areas stuck in sub-optimal equilibrium, and one-time interventions that “shock” the system to a different equilibrium would be called for (for example, see O’Connor et al. (1996)). While other commentators (Fisher et al., 2003a,b; Phelps, 2000) have emphasized the welfare cost of excess geographic variation in practice style, our theoretical model (with a single equilibrium) suggests that there is too little variation from a welfare perspective – if aggressive areas became even more aggressive, average patient welfare would increase. Thus, understanding the optimal policy response hinges importantly on whether the variations observed in the data are the consequence of single or multiple equilibria.

Our results also raise important questions about what can be learned from randomized controlled trials in medicine. While randomized trials are considered the gold standard for determining the effectiveness of a given medical treatment, they are designed to provide a partial-equilibrium estimate of the treatment effect in a well-defined population. But with productivity spillovers, the general equilibrium effect of adopting a new treatment could be smaller or larger than the partial equilibrium estimate of treatment effectiveness, because of the negative externality imposed on patients who are more appropriate for an alternative treatment and the positive externality on patients who are more appropriate for the new treatment. This general-equilibrium effect is not identified in a typical randomized trial, but could potentially be estimated in a trial that randomized across areas rather than individuals. In addition, the effectiveness of the treatment depends on where the trial is conducted (just as our IV estimates depended on intensity in the area): Surgical interventions may perform poorly in an area that specializes in medical management of its patients and perform well in a more surgically intensive area. As such the external validity of a randomized trial is compromised.

The implications of our findings extend beyond the treatment of heart attacks. To the extent that productivity spillovers are a common feature of many sectors, our results provide compelling evidence that such spillovers are an important feature of geographic specialization. Moreover, our results provide some of the first direct evidence of the negative externalities imposed on a subset of the population because of equilibrium pressures towards specialization. Such negative externalities are a central component of arguments for government intervention. Finally, our model has interesting empirical implications when applied to the more general question of human capital externalities. For example, our model would suggest that people living in areas with higher ability populations would be more likely (holding ability constant) to go to college, the return to going to college in these areas would be higher, but wages of low ability people in these areas would be lower. In principal, these are all testable implications.

References

- Arthur WB. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal* 1989;99:116–131.
- Andersen HR, Nielsen TT, Rasmussen K, et al. A comparison of coronary angioplasty with fibrinolytic therapy in acute myocardial infarction. *New England Journal of Medicine* 2003;349:733–742. [PubMed: 12930925]
- Beck, Christine A.; Penrod, John; Gyorkos, Theresa W.; Shapiro, Stan; Pilote, Louise. Does Aggressive Care Following Acute Myocardial Infarction Reduce Mortality? Analysis with Instrumental Variables to Compare Effectiveness in Canadian and United States Patient Populations. *Health Services Research* 2003;386:1423–40. [PubMed: 14727781]

- Coleman JS, Katz E, Menzel H. The Diffusion of an Innovation Among Physicians. *Sociometry* December;1957 20(4):253–270.
- Baicker, Katherine; Chandra, Amitabh. Medicare Spending, The Physician Workforce, and The Quality Of Health Care Received By Medicare Beneficiaries. *Health Affairs* April;2004a :184–97.
- Barnato, Amber, et al. Medical Care. 43.4. Apr. 2005 Hospital-Level Racial Disparities in Acute Myocardial Infarction Treatment and Outcomes; p. 308–319.
- Becker, Gary; Murphy, Kevin. The Division of Labor, Coordination Costs, and Knowledge. *Quarterly Journal of Economics* November;1992 107(4):1137–1160.
- Enthoven, Alain C. *Health Plan*. Reading, MA: Addison-Wesley; 1980.
- Fisher, Elliott, et al. The Implications Of Regional Variations In Medicare Spending. Part 1: The Content, Quality, And Accessibility Of Care. *Annals of Internal Medicine* February 18;2003a 138(4):273–87. [PubMed: 12585825]
- Fisher, Elliott, et al. The implications of regional variations in Medicare spending. Part 2: Health outcomes and satisfaction with care. *Annals of Internal Medicine* February 18;2003b 138(4):288–98. [PubMed: 12585826]
- Fuchs, Victor F. More Variation In Use Of Care, More Flat-Of-The-Curve Medicine. *Health Affairs*. October 7;2004 Web Exclusive
- Glover, Allison F. The Incidence of Tonsillectomy in School Children. *Proceedings of the Royal Society of Medicine* 1938;31:1219–1236.
- Goolsbee, Austan; Klenow, Peter. Evidence on Learning and Network Externalities in the Diffusion of Home Computers. *Journal of Law & Economics* October;2002 45:317–344.
- Gruber, Jonathan; Levine, Phillip B.; Staiger, Douglas. Abortion Legalization and Child Living Circumstances: Who is the Marginal Child? *Quarterly Journal of Economics* 1999;263–292.
- Heckman, James J.; Honore, Bo. The Empirical Content of the Roy Model. *Econometrica* Sept;1990 58 (5):1121–49.
- Jacobs Alice K. Primary Angioplasty for Acute Myocardial Infarction -- Is It Worth the Wait? *New England Journal of Medicine* 2003;349:798–800. [PubMed: 12930933]
- Jencks SF, et al. Quality of Medical Care Delivered to Medicare Beneficiaries. *Journal of the American Medical Association* October 4;2000 284(13)
- Katz ML, Shapiro C. Network externalities, competition, and compatibility. *American Economic Review* 75;1985:424–440.
- Katz ML, Shapiro C. Technology adoption in the presence of network externalities. *Journal of Political Economy* 1986;94:822–841.
- Keeley EC, Boura JA, Grines CL. Primary angioplasty versus intravenous thrombolytic therapy for acute myocardial infarction: a quantitative review of 23 randomised trials. *Lancet* 2003;361:13–20. [PubMed: 12517460]
- McClellan, Mark; McNeil, BJ.; Newhouse, JP. Does more intensive treatment of acute myocardial infarction reduce mortality? *Journal of the American Medical Association* September;1994 272(11): 859–66. [PubMed: 8078163]
- McClellan, Mark; Newhouse, JP. The marginal cost-effectiveness of medical technology: a panel instrumental-variables approach. *Journal of Econometrics* March;1997 77(1):39–64.
- McClellan, M.; Noguchi, H. Validity and Interpretation of Treatment Effect Estimates: Using Observational Data. Department of Economics, Stanford University; Mimeo: 2001.
- Marshall, Alfred. *Principles of Economics*. New York NY: Macmillan; 1890.
- Marciniak TA, Ellerbeck EF, Radford MJ, et al. Improving the quality of care for Medicare patients with acute myocardial infarction: results from the Cooperative Cardiovascular Project. *Journal of the American Medical Association* 1998;279:1351–1357. [PubMed: 9582042]
- Meltzer, David. Unpublished manuscript. University of Chicago; 2006. Diffusion of Innovation through Social Network Connections among Physicians.
- O'Connor, Gerald T., et al. A regional intervention to improve the hospital mortality associated with coronary artery bypass graft surgery. The Northern New England Cardiovascular Disease Study Group. *Journal of the American Medical Association* 1996;275(11):841–6. [PubMed: 8596221]

- O'Connor, Gerald T., et al. Geographic Variation in the Treatment of Acute Myocardial Infarction: The Cooperative Cardiovascular Project. *Journal of the American Medical Association* 1999;281:627–633. [PubMed: 10029124]
- Phelps, CE. Information Diffusion and Best Practice Adoption. In: Culyer, AJ.; Newhouse, JP., editors. *Handbook of Health Economics*. 1. Elsevier Science B.V.; 2000. p. 224-264.
- Pilote L, McClellan M, Saynina O, Lavoie F. Cardiac procedures use and outcomes in elderly patients with acute myocardial infarction in the United States and Canada, 1988–1994. *Medical Care* 2003;41(7):813–822. [PubMed: 12835605]
- Rogers, Everett M. *Diffusion of Innovations*. 5. New York: The Free Press; 2003.
- Scanlon PJ, Faxon DP, Audet AM, et al. ACC/AHA guidelines for coronary angiography. *Journal of the American College of Cardiology* 1999;33(6):1756–824. [PubMed: 10334456]
- Soumerai, Stephen B., et al. Effect of Local Medical Opinion Leaders on Quality of Care for Acute Myocardial Infarction: A Randomized Controlled Trial. *Journal of the American Medical Association* 1998;279:1358–1363. [PubMed: 9582043]
- Shekelle, Paul G., et al. The reproducibility of a method to identify the overuse and underuse of medical procedures. *New England Journal of Medicine* 1998;338:1888–95. [PubMed: 9637810]
- John, Wennberg; Cooper, MM., editors. *The Quality of Medical Care in the United States: A Report on the Medicare Program, The Dartmouth Atlas of Health Care*. Chicago: American Health Association; 1999.
- Wennberg, David E.; Birkmeyer, JD. *The Dartmouth Atlas of Cardiovascular Health Care*. Chicago: American Hospital Publishing, Inc; 2000.

Data Appendix

I. Construction of Hospital Referral Regions

Hospital Referral Regions are constructed using a two part algorithm.

Step 1

All acute care hospitals in the 50 states and the District of Columbia were identified from the American Hospital Association Annual Survey of Hospitals and the Medicare Provider of Services files and assigned to a location within a town or city. The list of towns or cities with at least one acute care hospital (N=3,953) defined the maximum number of possible Hospital Service Areas (HSA). Second, all 1992 and 1993 acute care hospitalizations of the Medicare population were analyzed according to ZIP Code to determine the proportion of residents' hospital stays that occurred in each of the 3,953 candidate HSAs. ZIP Codes were initially assigned to the HSA where the greatest proportion (plurality) of residents were hospitalized. Approximately 500 of the candidate HSAs did not qualify as independent HSAs because the plurality of patients resident in those HSAs were hospitalized in other HSAs. The third step required visual examination of the ZIP Codes used to define each HSA. Maps of ZIP Code boundaries were made using files obtained from Geographic Data Technologies (GDT) and each HSA's component ZIP Codes were examined. In order to achieve contiguity of the component ZIP Codes for each HSA, "island" ZIP Codes were reassigned to the enclosing HSA, and/or HSAs were grouped into larger HSAs. This process resulted in the identification of 3,436 HSAs, ranging in total 1996 population from 604 (Turtle Lake, North Dakota) to 3,067,356 (Houston) in the 1999 edition of the Atlas. Intuitively, one may think of HSAs as representing the geographic level at which "front end" services such as diagnoses are received.

Step II

Hospital service areas make clear the patterns of use of local hospitals. A significant proportion of care, however, is provided by referral hospitals that serve a larger region. Hospital referral regions were defined in the *Dartmouth Atlas* by documenting where patients were referred for major cardiovascular surgical procedures and for neurosurgery. Each hospital service area was

examined to determine where most of its residents went for these services. The result was the aggregation of the 3,436 hospital service areas into 306 hospital referral regions that were named for the hospital service area most often used by residents of the region. Thus if a Medicare enrollee living in Hartford, CT was admitted to a hospital in Boston, MA, the utilization would be attributed to Hartford, and not to Boston. This assignment avoids the serious shortcoming of unusually high utilization rates in large referral centers such as Boston or Rochester, MN.

II. Construction of CCP Estimation Sample

The CCP used bills submitted by acute care hospitals (UB-92 claims form data) and contained in the Medicare National Claims History File to identify all Medicare discharges with an International Classification of Diseases, Ninth Revision, Clinical Modification (ICD-9-CM) principal diagnosis of 410 (myocardial infarction), excluding those with a fifth digit of 2, which designates a subsequent episode of care. The study randomly sampled all Medicare beneficiaries with acute myocardial infarction in 50 states between February 1994 and July 1995, and in the remaining 5 states between August and November, 1995 (Alabama, Connecticut, Iowa, and Wisconsin) or April and November 1995 (Minnesota); for details see O'Connor et.al (1999). The Claims History File does not reliably include bills for all of the approximately 12% of Medicare beneficiaries insured through managed care risk contracts, but the sample was representative of the Medicare fee-for-service (FFS) patient population in the United States in the mid-1990s. After sampling, the CCP collected hospital charts for each patient and sent these to a study center where trained chart abstracters abstracted clinical data. Abstracted information included elements of the medical history, physical examination, and data from laboratory and diagnostic testing, in addition to documentation of administered treatments. The CCP monitored the reliability of the data by monthly random reabstractions. Details of data collection and quality control have been reported previously in Marciniak et.al (1998). Finally, the CCP supplemented the abstracted clinical data with diagnosis and procedure codes extracted from Medicare billing records and dates of death from the Medicare Enrollment Database.

For our analyses, we delete patients who were transferred from another hospital, nursing home or emergency room since these patients may already have received care that would be unmeasured in the CCP. We transformed continuous physiologic variables into categorical variables (e.g., systolic BP < 100 mm Hg or $\geq$ 100 mm Hg, creatinine <1.5, 1.5–2.0 or >2.0 mg/dL) and included dummy variables for missing data. We used date of death to identify patients who did or did not survive through each of three time points: 1-day, 30-days, and 1-year after the AMI. For all patients, we identified whether they received each of 6 treatments during the acute hospitalization: reperfusion (defined as either thrombolysis or PCI within 12 hours of arrival at the hospital), aspirin during hospitalization, aspirin at discharge, beta-blockers at discharge, ACE inhibitors at discharge, smoking cessation counseling, and each of 3 treatments within 30-days of the AMI: cardiac catheterization, PCI, or CABG. We used the ACC/AHA guidelines for coronary angiography to identify patients who were ideal (Class I), uncertain (Class II), or inappropriate (Class III) for angiography; these details are provided in Scanlon et.al (1999).

For each AMI patient we computed Medicare Part A and Part B costs within 1 year by weighting all Diagnosis Related Groups (DRGs) and Relative Value Units (RVUs) nationally. This measure of costs abstracts from the geographic price adjustment in the Medicare program.

III. Construction of clinical appropriateness index, HRR catheterization rates, and HRR beta-blocker rates

To compute this index we estimate a logistic regression model for the probability of receiving cardiac catheterization within 30 days of the heart-attack. We estimate:

$$Pr(\text{Cardiac Cath}_{ij}) = F(\theta_0 + \theta_j + X_i\Phi)$$

$$Pr(\text{Beta Blockers}_{ij}) = F(\theta_0 + \Pi_j + X_i\Phi)$$

Fitted values from this regression, $Pr(\text{Cath}|X)$, are used as an empirical measure clinical appropriateness for cardiac catheterization. The HRR fixed effects for the use of cardiac catheterization and beta-blockers are obtained as θ_j and Π_j respectively. Note that these fixed-effects are not used to obtain our empirical measure of appropriateness for cardiac catheterization. Included in this model are the following risk-adjusters:

Age, Race, Sex (full interactions)	hx angina missing (ref=no)	peak ck gt 1000
previous revascularization (1=y)	hx terminal illness (1=y)	non-ambulatory (ref=independent)
hx old mi (1=y)	current smoker	ambulatory with assistance
hx chf (1=y)	atrial fibrillation on admission	ambulatory status missing
history of dementia	cpr on presentation	albumin low(ref>=3.0)
hx diabetes (1=y)	indicator mi = anterior	albumin missing(ref>=3.0)
hx hypertension (1=y)	indicator mi = inferior	bilirubin high(ref<1.2)
hx leukemia (1=y)	indicator mi = other	bilirubin missing(ref<1.2)
hx ef <= 40 (1=y)	heart block on admission	creat 1.5-<2.0(ref=<1.5)
hx metastatic ca (1=y)	chf on presentation	creat >=2.0(ref=<1.5)
hx non-metastatic ca (1=y)	hypotensive on admission	creat missing(ref=<1.5)
hx pvd (1=y)	hypotensive missing	hematocrit low(ref=>30)
hx copd (1=y)	shock on presentation	hematocrit missing(ref=>30)
hx angina (ref=no)	peak ck missing	ideal for CATH (ACC/AHA criteria)

The choice of variables was based on those selected by Fisher et al. (2003a,b) and Barnato et al. (2005). With the exception of two variables that are both measured by blood-tests, albumin and bilirubin (where the rates of missing data were 24 percent), we do not have a lot of missing data (rates were less than 3 percent). Furthermore, there is no relationship between the presence of a missing values for albumin and bilirubin and the intensity of an area. Nor does having missing values for these variables affect the probability of receiving catheterization.

Figure 1. Multiple Equilibrium (Panel A), versus Single Equilibrium (Panel B) Characterizations of Area Variations.

Figure 2. Graphical Illustration of Roy Model with Productivity Spillovers

Panel A describes the relationship between two alternative ways to treat patients within an area. The survival production possibilities frontier describes the best treatment for a patient of given clinical appropriateness. The model predicts that the returns to intensive management are increasing in patients' appropriateness for such interventions. Panel B contrasts the care across two areas that differ in their surgical intensity. The productivity spillover results in patients appropriate for intensive management being better off in the surgically intensive areas, whereas patients appropriate for non-intensive management being worse off in such areas.

Figure 3. Relation Between Average Patient and Marginal Patient Receiving Cardiac Catheterization

For each of the 306 HRRs we graph the average propensity to receive cardiac catheterization (amongst patients who actually received it) against the log of the area risk-adjusted CATH rate. Using local-regression, we estimated the relationship between the average propensity and the risk-adjusted CATH rate and the slope of this line at each point. These estimates were then used to plot average (upper line) and marginal patient (lower line and estimated as the local difference in the average) receiving treatment.

Table 1

Instrumental Variable Estimates of Intensive Management and Spending on 1 Year Survival by Clinical Appropriateness of Patient

Sample:	IV Estimates of		
	Impact of CATH: on 1-Year Survival	on 1-year Cost (\$1000s)	Impact of \$1000: on 1-year Survival
1. All patients (n=129,895)	0.142 (0.036)	9.086 (1.810)	0.016 (0.005)
2. By CATH propensity			
a. Above the Median (n=64,799)	0.184 (0.034)	4.793 (1.997)	0.038 (0.017)
b. Below the Median (n=65,096)	0.035 (0.083)	17.183 (3.204)	0.002 (0.005)
Difference:	0.149 (0.090)	-12.39 (3.775)	0.036 (0.018)
3. By age			
a. 65-80 (n=89,947)	0.171 (0.037)	6.993 (1.993)	0.024 (0.009)
b. Over 80 (n=39,948)	0.016 (0.108)	16.026 (2.967)	0.001 (0.007)
Difference:	0.155 (0.114)	-9.033 (3.574)	0.023 (0.011)
4. By ACC/AHA Criteria			
a. Ideal for CATH (n=83,768)	0.145 (0.039)	8.807 (2.050)	0.016 (0.006)
b. Not Ideal for CATH (n=46,123)	0.124 (0.083)	8.895 (2.780)	0.014 (0.010)
Difference:	0.021 (0.092)	-0.088 (3.232)	0.002 (0.012)

Notes: CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters. Differential-distance (measured as the distance between the patient's zip-code of residence and the nearest catheterization hospital minus the distance to the nearest hospital) is the instrument. Each model includes all the CCP risk-adjusters and the standard errors are clustered at the level of each HRR.

Table 2
Relationship Between Differential-Distance and Probability of Catheterization and Survival, and Differential-Distance and Observable Characteristics

	30-day CATH rate		1-year Survival		1-year Predicted Survival		30-day predicted CATH rate for patients getting CATH	
	DD Below	Median	DD Above	Median	DD Below	Median	DD Below	Median
Sample:								
All patients (n=129,997)	48.9%		42.8%		67.5%		63.3%	
By CATH propensity								
Above	74.0%		67.1%		83.4%		72.6%	
the Median								
(n=64,733)								
Below	22.9%		19.5%		51.1%		32.3%	
the Median								
(n=65,244)								
By age								
65-80	61.1%		54.9%		73.9%		67.4%	
(n=90,016)								
Over 80	20.3%		16.5%		52.6%		34.6%	
(n=39,961)								

Notes: CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters. Differential-distance (measured as the distance between the patient's zip-code of residence and the nearest catheterization hospital minus the distance to the nearest hospital) is the instrument. Each model includes all the CCP risk-adjusters and the standard errors are clustered at the level of each HRR.

Table 3
Relationship between the Average and Marginal Patient Receiving Cardiac Catheterization

	Characteristic of average patient getting CATH across all areas	Difference Between average patient and marginal patient getting CATH in higher-CATH HRRs
Patient Characteristic		
1. CATH propensity	0.633 (0.002)	−0.045 (0.008)
2. Over age 80	0.125 (0.002)	0.063 (0.012)
3. Not eligible for CATH using ACC/AHA guidelines	0.028 (0.001)	0.010 (0.003)
N	303	303

Notes: CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters. Sample is restricted to patients receiving cardiac catheterization within 30 days of an AMI. ACC/AHA guidelines reflect a binary variable assigned to each patient in the CCP that measures whether the patient is ideal, appropriate, or not eligible for catheterization based on chart-review.

Correlation of HRR Level Measures of Intensive Treatment, Medical Management, Support of Medical Treatment and Demographic Characteristics

Table 4

HRR indicator	Mean	SD	10th percentile	90th percentile	Correlation with HRR CATH rate
Measures of intensive treatment					
Risk-adjusted 30-day CATH rate	46.3%	9.1%	34.5%	58.3%	1.00
Risk-adjusted 30-day PTCA rate	17.7%	5.1%	11.3%	23.6%	0.81
Risk-adjusted 30-day CABG rate	13.4%	2.9%	10.2%	16.9%	0.51
Risk-adjusted 12-hour PTCA rate	2.7%	2.6%	0.6%	5.8%	0.52
Measures of Quality of medical management					
Risk-adjusted β -blocker rate	45.6%	9.5%	34.2%	58.3%	-0.31
Support for intensive treatment					
Cardiovascular Surgeons per 100,000	1.06	0.27	0.70	1.40	0.33
Cath Labs per 10,000	2.40	0.76	1.50	3.30	0.39
Demographic characteristics					
log of resident population	13.96	0.89	12.72	15.18	-0.05
log of per capita income	9.55	0.20	9.31	9.85	0.02
Percent College graduates	19.3%	5.5%	13.1%	26.6%	-0.05

Notes: HRR surgical and medical intensity rates are computed as the risk-adjusted fixed effects from a patient level regression the receipt of CATH or beta-blockers on HRR fixed effects and CCP risk-adjusters. CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters.

Table 5

OLS Restimates of the Relationship between Probability of Receiving Catheterization and HRR Patient Characteristics

HRR-Level Independent Variables:	Probability of Receiving Catheterization	Probability of Receiving Catheterization
Average propensity to get CATH	0.529 (0.172)	0.575 (0.167)
Percent under age 65		0.150 (0.135)
log of resident population		-0.003 (0.005)
log of per capita income		0.024 (0.024)
N	138,873	138,873

Notes: Table reports OLS estimates of the relationship between a patient receiving Catheterization and the average appropriateness for Catheterization in an HRR. Regressions control for patient risk-adjusters, and standard-errors are clustered at the level of HRRs.

Table 6

Instrumental Variable Estimates of Intensive Management and Spending on Survival, by Surgical Intensity of Hospital Referral Region

Sample:	IV Estimates of		
	on 1-Year Survival	Impact of CATH: on 1-year Cost (\$1000s)	Impact of \$1000: on 1-year Survival
1. All patients			
HRR risk-adjusted CATH rate:			
a. Above the Median (n=63,771)	0.256 (0.061)	6.691 (3.510)	0.038 (0.021)
b. Below the Median (n=66,124)	0.09 (0.059)	9.835 (3.155)	0.009 (0.007)
Difference:	0.166 (0.085)	-3.144 (4.720)	0.029 (0.022)
2. Patients above the median CATH propensity			
HRR risk-adjusted CATH rate:			
a. Above the Median (n=32,388)	0.271 (0.064)	0.347 (4.370)	0.78 (9.820)
b. Below the Median (n=32,411)	0.168 (0.046)	4.962 (2.890)	0.034 (0.021)
3. Patients below the median CATH propensity			
HRR risk-adjusted CATH rate:			
a. Above the Median (n=31,383)	0.206 (0.129)	16.21 (5.130)	0.013 (0.009)
b. Below the Median (n=33,713)	-0.139 (0.165)	22.064 (6.870)	-0.006 (0.007)

Notes: HRR intensity rates are computed as the risk-adjusted fixed effects from a patient level regression the receipt of CATH or beta-blockers on HRR fixed effects and CCP risk-adjusters. Differential-distance (measured as the distance between the patient's zip-code of residence and the nearest catheterization hospital minus the distance to the nearest hospital) is the instrument. Each model includes all the CCP risk-adjusters and the standard errors are clustered at the level of each HRR.

Table 7

Relationship Between HRR Catheterization Rate, Survival and Costs, by Clinical Appropriateness for Intensive Management

Sample:	OLS Estimates of the Relationship Between HRR Risk-Adjusted CATH Rate and:			
	1-Year Survival	1-year Cost (\$1000s)	Beta Blocker in Hospital	Catheterization within 30 days
1. All patients (n=138,873)	0.007 (0.019)	8.093 (1.410)	-0.28 (0.073)	0.702 (0.004)
2. By CATH propensity				
a. Top Tercile (n=46,287)	0.052 (0.019)	10.012 (1.439)	-0.366 (0.073)	0.802 (0.032)
b. Middle Tercile (n=46,295)	0.03 (0.030)	11.154 (1.784)	-0.271 (0.082)	0.906 (0.021)
c. Bottom Tercile (n=46,291)	-0.075 (0.028)	2.763 (1.612)	-0.209 (0.073)	0.369 (0.021)
Difference (Top-Bottom):	0.127 (0.034)	7.249 (2.161)	-0.157 (0.103)	0.433 (0.038)
3. By age				
a. 65-80 (n=96,093)	0.023 (0.021)	9.616 (1.448)	-0.311 (0.072)	0.775 (0.012)
b. Over 80 (n=42,780)	-0.031 (0.028)	4.738 (1.603)	-0.215 (0.080)	0.531 (0.022)
Difference (Top-Bottom):	0.054 (0.035)	4.878 (2.160)	-0.096 (0.108)	0.244 (0.025)
4. By AHA/ACC criterion				
a. Ideal (n=89,569)	0.027 (0.023)	9.845 (1.599)	-0.302 (0.076)	0.769 (0.010)
b. Appropriate (n=31,800)	-0.002 (0.024)	6.174 (1.537)	-0.282 (0.080)	0.752 (0.026)
c. Not appropriate (n=17,504)	-0.08 (0.040)	2.958 (1.511)	-0.177 (0.065)	0.264 (0.021)
Difference (Top-Bottom):	0.107 (0.046)	6.887 (2.200)	-0.125 (0.100)	0.505 (0.023)

Notes: CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters. HRR surgical and medical intensity rates are computed as the risk-adjusted fixed effects from a patient level regression the receipt of CATH or beta-blockers on HRR fixed effects and CCP risk-adjusters.

Table 8
Relationship between Survival and alternative measures of HRR size, by clinical appropriateness of patient

HRR- Level Independent Variables:	All patients	By CATH propensity			All patients	By CATH propensity		
		Top Tercile	Middle Tercile	Bottom Tercile		Top Tercile	Middle Tercile	Bottom Tercile
Risk-adjusted CATH rate	0.006 (0.019)	0.052 (0.019)	0.031 (0.030)	-0.077 (0.028)	0.008 (0.019)	0.053 (0.019)	0.031 (0.030)	-0.075 (0.028)
log of resident population	-0.002 (0.002)	-0.001 (0.002)	0.002 (0.003)	-0.007 (0.003)	0.005 (0.004)	0.014 (0.005)	0.006 (0.007)	-0.007 (0.006)
log of AMIs per hospital								
N	138,873	46,287	46,295	46,291	138,873	46,287	46,295	46,291

Notes: CATH propensity is an empirical measure of patient appropriateness for intensive treatments. We define this measure by using fitted values from a logit model of the receipt of cardiac catheterization on all the CCP risk-adjusters. HRR surgical and medical intensity rates are computed as the risk-adjusted fixed effects from a patient level regression the receipt of CATH or beta-blockers on HRR fixed effects and CCP risk-adjusters.