April 2003
Remarkably Preliminary
Do not quote or cite

The Diffusion of Health Care Technology

Jonathan Skinner
Dartmouth College and NBER

Douglas Staiger
Dartmouth College and NBER

Abstract

There are a variety of models that seek to explain the process of technological innovation, but less is understood about the nature of innovation and diffusion in health care. A standard Bayesian learning-by-doing model may imply that different hospitals or regions may evolve different strategies for treating specific diseases, but the productivity of their strategies – that is, their quality-adjusted price – should exhibit convergence over time. More complicated models where regions or hospitals differ systematically in barriers to adopting new technology, can imply non-convergence. Using Medicare data on 2.6 million heart attack patients during 1989-2000, we find first that the rapid technological gains occurring during the late 1980s and early 1990s have flattened by the late 1990s; indeed the quality-adjusted price has risen since 1995. Second, in considering regional differences in mortality costs and outcomes, we found no evidence of convergence with respect to productivity during this period. As well, the evidence favors long-term differences across states in their propensity to adopt new technology; states likely to adopt hybrid corn in the 1930s and 1940s were most likely to adopt the use of highly effective and low cost $-blockers in the 1990s.

We are grateful to seminar participants at UCLA, Rand, and the University of Pennsylvania, and to Amitabh Chandra, James Feyrer, Amy Finkelstein, Elliott Fisher, David Weil and Jack Wennberg for helpful comments and discussions. Weiping Zhou and Daniel Gottlieb provided superb data analysis. This research was funded by the National Institute on Aging.

1. Introduction

The process by which new innovations are adopted and diffused has been the subject of a variety of studies both theoretical and empirical. Griliches (1957) described the differential adoption of hybrid corn, with some states adopting early and others later; ultimately technological progress converged across states so that nearly all farmers in each region were using the more productive corn. Recently, Jovanovic and Nyarko (1995, 1996) developed models of learning-by-doing where individuals or firms gradually improve their productivity in performing functions or procedures, and face the options of dropping the old (and better-understood) technology and adopting a new technology. While these models can generate wide variation in productivity and slow convergence, depending on the complexity of the learning process, they too ultimately exhibit convergence in productivity gains. The process of convergence is accelerated in competitive markets, where firms that do not improve productivity are ultimately driven out of the market, thus leading to convergence through attrition.

In this paper we ask whether there has been convergence in the productivity of medical care for heart attacks across regions in the United States. The diffusion of technological innovation in health care is of interest for at least two reasons. First, with health care expenditures at 14 percent of GDP, it is the single largest industry in the U.S. As well, the output of health care, health, is conventionally valued in dollar terms far in excess of GDP; this is because the typical quality-adjusted life year is priced at between $50,000 - $100,000 or more, considerably larger than per capita GDP. Second, there is a widespread perception that health care is a quite different output from hybrid corn or

other marketed commodities. It is more difficult for the physician or health professional to discern the signal from the noise in treatment strategies; sample sizes of patients of given characteristics with a given disease tends to be small relative to the dimensionality of the learning process and if physicians are not actively using a given treatment, they may be unable to learn about its effectiveness (Bikhchandani et al, 2002). In theory, medical literature is supposed to help diffuse the relevant information, but in practice physicians often do not translate the academic studies into actual practice (Berwick, 2003). Finally, there are no competitive national markets in health care to put pressure on inefficient health care providers. Instead, there are incentives for hospitals to adopt technologies with the greatest profitability for providers, even if these specific technologies are not the most cost-effective.

We adopt the Jovanovic and Nyarko (1995, 1996) models of learning-by-doing in a simplified model of technology diffusion and adoption to develop our testable implications. These models illustrate how optimal learning about complex decisions can lead to large and potentially long-lasting differences across firms in productivity during periods of technological innovation. Moreover, it is entirely possible for different physicians, and indeed different regions, to adopt varying technologies for health care treatments, i.e. we could expect to observe divergence in the adoption of treatment strategies. A strong implication of this class of models, however, is convergence with regard to productivity. Cutler et al (1998) developed an ideal measure for measuring productivity through their "cost of living" price index for heart attack patients, since it captures both the monetary price of the treatment and the quality (as measured by expected life years) of the treatment. Regardless of the model of technological diffusion,

long-run efficiency requires that this quality-adjusted price of health care should *converge* over time. This is the measure we consider in the empirical analysis below.

Alternative models tend to imply slower convergence or even non-convergence. Different hospitals or regions may vary with regard to their adoption of new technologies (as in Jovanovic and Nyarko, 1996), and if those variations are permanent – for example, because of systematically different discount rates for innovation or other persistent barriers to adoption – then some regions could lag permanently behind. Other models loosen the implausible assumption in Jovanovic and Nyarko (1995) that the "correct" treatment is revealed after the physician treats the patient. Fournier et al (2002) consider a learning-by-doing model in which random realizations cause physicians in a region to iterate down either one branch of treatment (say surgery) while other regions converge towards a different strategy (nonsurgical drug treatments), with either being potentially efficient. In Bikhchandani et al. (2002), Bayesian learning with informational cascades may lead to less efficient technology adoption in some areas – they aren't using the more efficient treatment and hence can't learn of its favorable effects, leading to an inefficient (and potentially non-converging) equilibrium.

Our empirical work uses a sample of 2.6 million elderly patients with heart attacks (acute myocardial infarction, or AMI) from the Medicare claims data during 1989-2000. Technological improvements in the treatment for heart attacks have been well established at least through the early 1990s (Cutler et al, 1998) and have figured prominently in the debate over the "money's worth" of high-technological medical spending (Cutler and McClellan, 2001). Both costs and risk-adjusted mortality rates are considered separately for each of the 306 Hospital Referral Regions (HRRs) defined in

the Dartmouth Atlas of Health Care (Wennberg and Cooper, 1999). These are regions that were determined according to the migration patterns of individual patients to hospitals; each HRR contains at least one hospital that provides surgical services for heart attack patients. By considering the data at the regional level, we thereby can partially sidestep problems of small sample sizes and case-mix issues encountered at the individual physician or hospital level of analysis.

We have three primary findings. First, the sharp decline in quality-adjusted prices for AMI found by Cutler et al (1998) reversed itself by the mid-1990s, with quality-adjusted prices rising since about 1995. Second, we find no evidence of convergence across regions with regard to mortality, cost, or the quality-adjusted price of AMI treatment. Indeed, we find some evidence of divergence across some regions: Regions that were relatively early adopters of beta-blockers (a cost-effective drug treatment diffusing during this period) experienced much larger declines in mortality and quality-adjusted prices than did regions with later adoption of beta-blockers, while regions with a high supply of cardiologists per capita experience a larger increases in costs and quality adjusted prices than did regions with a low supply of cardiologists. Our final finding, returning to the original Griliches (1957) data on the adoption of hybrid corn, is that states adopting the hybrid varieties earlier in the 1930s and 1940s also tended to be those states adopting effective treatment technologies for AMI in the 1990s. That is, states appear to differ in the long term with regard to the adoption of *any* technology. This last finding suggests that productivity differences may the result of persistent and pervasive barriers to technology adoption rather than transitory information differences or particular inefficiencies of the health care market. More generally, the results found for heart attack

treatments suggest that, consistent with macro-level studies, the factors affecting technological adoption (and total factor productivity) are more important than the overall level or growth in the use of capital and labor in determining differences in productivity over time and across areas.

2. A Model of Technological Innovation and Diffusion in Health Care

We base our model largely on Jovanovic and Nyarko (1995). Assume initially that the quality choice consists solely of choosing a single output, survival, and a single input, for example the correct dosage of a drug. Let quality of care q_i for patient i is given by

$$q_i = A[1 - (y_i - z_i)^2] \qquad (1)$$

where $y_i = \theta + w_i$; θ is the expected value of the best possible care, and w is an iid random variable that reflects the idiosyncratic characteristics of the patient. After treating each patient, the physician observes q and y, but must chooses z_i before observing y or q. Optimally, the physician sets $z_i = E_{i-1}(\theta|\Omega_{t(i)})$, where z_i is the expectation of the optimal treatment θ conditional on past experience with the previous i-1 patients, as well as the general information set $\Omega_{t(i)}$ that was conveyed through what the physician has absorbed from the medical literature available at time t(i) or through informal spillovers associated with other physicians practicing in the same region or hospital. For every additional patient, the physician updates z through standard Bayesian updating; as well updating may occur through the improved quality of the information set $\Omega_{t(i)}$. A good example of this approach would be where the physician is trying to control hypertension using ACE inhibitors; the outcome (blood pressure) is observed every time a new patient arrives or

even when the dosage of the drug is changed.[1] Clearly, if every provider begins with different expectations of the appropriate value for z (that is, $x_0 = E_0(\theta) - \theta$ at time 0 may differ across different physicians), they will gradually converge to $z = \theta$; some physicians will get there earlier because of random patient outcomes, while others will take longer.[2] Because of the randomness in w, there will be randomness in individual patient outcomes q_i, but that randomness will quickly average out over a modest number of patients within a hospital or region, leading to convergence. In other words, average outcomes across areas will converge as beliefs about the optimal decision, $z_i = E_{i-1}(\theta|\Omega_{t(i)})$, converge.

The model becomes more complex when there is more than one input in the treatment of the disease. Consider the generalization of the single variable model above;

$$q_i = A\Pi_{s=1} [1 - (y_{si} - z_{si})^2] \quad (2)$$

where there are s = 1,…,N additional inputs into the quality "production function." Jovanovic and Nyarko (1995) demonstrate that a larger N (more complex decisions) generates more steep learning curves, more variation in productivity (q) at any point in time, and slower convergence over time. Nevertheless, convergence is still implied, particularly toward the end of the learning process when productivity gains are slowing and beliefs about optimal treatments are converging on the correct values.

The model is easily extended to consider competing technologies. Suppose there are two strategies to treating heart attack patients, say nonsurgical (as above, in Equation 2), and surgical, denoted by astericks:

$$q^*_i = A^* \Pi_{s=1} [1 - (y_{si}^* - z_{si}^*)^2] \quad (3)$$

---

[1] These are pharmaceuticals that inhibit the production of certain hormones that constrict blood vessels, and thereby attenuate high blood pressure.

[2] As shown in Jovanovic and Nyarko (1995), the updating rule is $x_i = \sigma_w^2\sigma_\theta^2/(\sigma_w^2 + \sigma_\theta^2 i)$.

Because of the gains in efficiency based on learning-by-doing, it is possible to observe one region using predominately surgical methods, and another using predominately non-surgical methods, as in Fournier et al, 2002. Assume for the moment that there is no switchover from one technology to another. However, this does not necessarily imply the absence of convergence; to this point we have not introduced costs or prices, and it may be that one technology may be more expensive, but may also yield better results.

To see this point, suppose that expected survival under the two regimes (after convergence) is S and S*, and the value of a life-year is φ. Suppose further that costs of the two regimes are C(Z) and C(Z*). If $(S-S^*)\varphi = C(Z) - C(Z^*)$, then we would judge either strategy of equal value, and there should be no reason why homogeneity should be encouraged – indeed, were one region to adopt the other's technology, the quality of care would decline because of their lack of experience with the newly adopted technology.

Cutler et al (1998) developed a quality-adjusted price index for the treatment of AMI that exactly captures this productivity index. They define quality-adjusted price index (relative to the numeraire at time 0) as

$$P_t = 1 + [C(Z_t) - C(Z_0) - \varphi(S_t - S_0)]/L$$

where subscripts refer to the time period. The denominator L captures not just the initial cost of the heart attack treatment, $C(Z_0)$, but instead is designed to reflect the "cost of living," that is, the dollar cost of stayin'alive, which they assume to be measured by the present value of remaining lifetime resources.[3] Using alternative measures of L is an issue of scaling, but in fact the use of lifetime resources provides a convenient metric that is related most closely to equivalent variation as a fraction of remaining income. Thus if

[3] One could argue that staying alive actually costs less than the average present value of future resources, but as noted above, using a different number involves just a scalar transformation.

survival improves by enough to lower the real price to $P_t = 0.95$ (for example), then we might interpret the technological innovations as if consumers (those with AMI in any case) would be willing to pay 5% of their income in order to enjoy the time t technology compared to the time 0 technology.

It is straightforward to adopt this valuable mechanism to the evaluation of different strategies for treating AMI. That is, treating P_0 as the numeraire, the equivalent measure of price for strategy * is given by

$$P_t^* = 1 + [C(Z_t^*) - C(Z_0) - \varphi(S_t^* - S_0)]/L$$

This can be expressed as two components, $P_t^* - P_0^*$, the change over time in the quality-adjusted price of health care treatment, and $P_0^* - 1$, the static cross-sectional difference in the effectiveness of the * strategy compared to the nonsurgical strategy. (In our empirical calculations, the unit of observation will be the hospital referral region, or HRR, rather than the type of treatment *per se*. There is a substantial degree of variation across regions in treatment strategies, however.) We suspect that the former (secular) component of quality-adjusted price changes will be measured with greater accuracy than cross-sectional measures, given the greater stability in underlying unmeasured health status over time compared to that in the cross-section.

To this point, we have not considered technology switching or adoption. Suppose that the decision to adopt new technology is a function of the discounted value of the new innovation, both with regard to the stream of quality, as well as the stream of profits from adopting the new technology.[4] For example, suppose we write the expected gains from adopting the * technology, relative to the existing technology, as

[4] While Jovanovic and Nyarko (1996) develop generalized models of switching technology, we depart from their approach by considering explicit discounting and costing factors.

$$\Delta_{qi} = \Sigma_{k=i} (q_k^* - q_k)\lambda^{-t(k)} \qquad (4a)$$

$$\Delta_{mi} = \Sigma_{k=i} (m_k^* - m_k)\lambda^{-t(k)} \qquad (4b)$$

where the subscript k reflects the number of future patients (predicted with certainty), t(k) is the time, relative to the current time, where the kth patient will appear (so if volume is 300 heart attack patients per year, t(301) = 1 year), and $0 < \lambda < 1$ is the discount factor used to make tradeoffs both with regard to future increases in quality of care, and future net revenue from insurance and patient payments m_k (nonsurgical) and m_k^* (surgical) for the kth potential patient.[5] The probability of adopting the new technology is a positive function of both Δ_{qi} and Δ_{mi}.[6]

Why might hospitals or regions differ with regard to the adoption of new technology? One factor would be the expertise of the stock of physicians with the existing technology; to the extent they are closer to the optimal θ, the relative gains of shifting to the new technology (and the learning process associated with that), the less likely they are to switch (as in Jovanovic and Nyarko, 1996). A second factor would be the volume of new patients. The higher the volume, the more rapidly the physician learns about the new technology, and the lower is the discount rate for the future gains. (In other words, the future benefit arising from treating the 200th patient is discounted by 4 years when volume is 50 per year, compared to just one year when the volume is 200 per year.) A third factor in switching is the degree of scientific information Ω available on the new methods of treatment. Where scientific evidence is good about the effectiveness

---

[5] We ignore the impact on patient populations of changing technologies.

[6] We do not want to posit here the potential of substitutability between high quality care and money income, which is a common approach in the supplier-induced demand literature. (Newhouse (2003) refers to it as "stinting" although it could just as well be reflected as overtreatment as undertreatment.)

of a specific treatment, the expected variance of x (the difference between perceptions of θ and true θ) is smaller, making the putative costs of switching less. By contrast, when the quality of scientific evidence is weak, physicians end up relying entirely on learning-by-doing. Finally, differences in discount rates (λ) will influence the adoption of new technologies. When the future is highly discounted (small λ), physicians will be unwilling to incur the up-front costs of learning new technologies (in the form of temporarily lower productivity).

Bikhchandani et al (2002) has emphasized the poor convergence properties of learning-by-doing where physicians can potentially avoid effective treatments or adopt ineffective treatments because of a cascading effect where several randomly good outcomes (or randomly poor outcomes) induce fundamental choices of treatment (such as surgery or β blocker use). This alternative approach illuminates a fundamental weakness with the Jovanovic and Nyarko approach, the assumption that each value of y_i (appropriate treatment for patient i) is revealed after the patient has been treated. Of course, there is no *deux ex machina* to reveal to physicians after the fact that they should have administered β blockers to their patient. Nor is there any reason for these physicians to converge to the appropriate use of β blockers in a pure learning-by-doing model with a realistic specification of how knowledge is acquired.[7]

3. The Data

The primary dataset is a 20% sample of the Medicare Part A (hospital) claims data for all heart attack (AMI) patients age 65 and over in the U.S. during 1989 – 1992,

[7] There are also models of heterogeneity in the acquisition and acceptance of new technology, see Berwick (2003) for a discussion and application to health care.

and a 100% sample from 1993 through 2000, with information from the denominator file (i.e., information on mortality) through December 31, 2001, N = 2,573,775.[8] The Medicare claims data includes detailed information on comorbidities (i.e., preexisting conditions) such as vascular disease, pulmonary disfunction, dementia, diabetes, liver disfunction, renal failure, cancer, and metastatic cancer. As well, there are also variables measuring the type of heart attack: anterior, inferior (denoting the location of the blockage, with anterior the most serious) and non-Q wave infarction (one in which the blockage of the heart muscle is less extensive and thus with much better clinical prognosis).

The primary measure of health outcomes is the one-year mortality rate; thus if the primary event – the heart attack – was on December 31, 1999, we require information on mortality through December 31, 2000. (We also report the 30-day mortality rate, since medical technologies for heart attacks are most likely to influence mortality in the first 30 days.) As a first step, regression analysis was used to estimate risk-adjusted mortality rates, where the covariates were 5-year age intervals (and the category of age 85+), sex, and race (black and nonblack). A separate dummy variable was created for each possible classification by age, sex, and race (i.e., a dummy variable for nonblack males age 70-74). The individual comorbidities were also included in the regression, along with year dummy variables (and no constant term). All comorbidities and age-sex-gender variables were demeaned, so that the year dummy variable coefficients are the year-specific mean values for a hypothetical individual with the average (over 1989-2000) set of covariates.

[8] We have performed a preliminary merge of the 5% Part B data sample from 1992-1997 (except for 1993 which is not available), and the 20% sample from 1998-2000, but these have not yet been included in the analysis.

Individuals with prior AMIs (117,763) and those with HMO membership in the current year (51,849) were excluded from the sample.

To measure HRR-year-level adjusted mortality, we estimated the simple comorbidity adjusted model, and then calculated mean residuals by HRR and year to create the required mortality measures for the HRR and year cells. The data analysis considered below is therefore at the HRR and year level, leading to approximately 3060 observations.

Price data was constructed using just the Part A (hospital) expenditures, deflated by the GDP deflator to 2000 dollars. To adjust for the fact that Medicare enrollees in some regions may be older or sicker than others, the Part A expenditures were adjusted (as for mortality) by regressing expenditures on the full set of comorbidities and demographic information, and using the residual values to average over HRR-year cells. Spending variables have not been adjusted, however, for differences in prices across regions; for example in some regions hospitals receive additional reimbursements because they are teaching hospitals, or because of the Disproportionate Share Hospital (DSH) program designed to provide more payments to hospitals with a larger fraction of low income or Medicaid patients. As well, Medicare reimbursements may differ across areas because of "outlier" payments for patients who have run up very large hospital bills. Note that we are measuring Medicare reimbursements, and not hospital costs. Unfortunately, hospital costs are nearly impossible to measure from administrative data, even if one could decide what measure of average (or marginal) cost one wished to measure.

The cost and mortality data at the HRR level was also merged with other variables. For example, we used the HRR-level rate of β blocker use among ideal heart attack patients during 1994/95, based on the *Dartmouth Atlas of Cardiovascular Care* measures based in turn on the Cooperative Cardiovascular Project (CCP) data. These are pharmaceutical treatments for heart attack patients (whether treated surgically or nonsurgically) and are proven effective for appropriate patients. We focus here on β Blockers because their clinical effectiveness has been well documented in the medical literature, yet their adoption has varied widely across regions.

4. Empirical Results

We consider three sets of results; the first on overall technical progress in the treatment of heart attacks, the second on convergence (or divergence) in the treatment of heart attacks across regions, and finally some preliminary evidence on long-term state-level factors in the adoption of new technology.

The Real Quality-Adjusted Price of Heart Attack Treatments

Figure 1 shows trends in both 30-day and one-year mortality rates after controlling for demographic factors and comorbidities. Both measures of mortality exhibit the dramatic decline during the early 1990s documented in Cutler et al (1998) and in Cutler and McClellan (2001), but since the mid-1990s, there has been a flattening of 30-day rates and even a noticeable increase in one-year mortality rates. The increase in mortality rates is surprising, but may be the consequence of delaying the mortality of people with underlying heart disease; mortality would decline early in the period, but later as more 80+ year olds survive, they could be frailer than average and thus cause an

increase even in age-adjusted mortality rates.[9] It could also be the consequence of the increasing fraction of HMO enrollees, particularly in some areas of the U.S.; this could bias upward mortality rates if the sickest patients are the ones who remain in the fee-for-service plan. However, there was no correlation by HRR between the change in mortality rates during 1989-2000, and the change in the total number of heart attack patients (i.e., a measure of the decline in the size of the fee-for-service population).

Aggregate one-year costs are shown in Figure 2. There is a general trend upward (in constant 2000 dollars), although with a slowing of the trend during the late 1990s. These price estimates are lower bounds as they ignore Part B (outpatient) costs. Moreover, in terms of overall Medicare spending, Part B costs grew more rapidly than Part A (inpatient) costs during this period, suggesting that the rise in Part A costs may understate the total rise in costs.

We derive the quality adjusted price index (relative to 1989) for heart attacks from changes in mortality and cost using $P_t = 1+ [C_t – C_{89} – \varphi(S_t – S_{89})]/L$, where C is the cost of care and S is patient survival. We use the Cutler et al. (1998) conversion from changes in one-year mortality to changes in life expectancy to derive S.[10] We also update their values for the value of a life-year (φ) and average income (L) to year 2000 dollars using the GDP deflator. While they use as a benchmark value of life-year of $25,000 ($29,804 in $2000), they are being quite conservative and so we consider $50,000 ($59,609 in $2000) as a reasonable alternative benchmark value of one life-year.[11]

---

[9] We are grateful to Mark Pauly for this suggestion.

[10] They estimate a $1/6^{th}$ year increase in survival following a 3 percentage point decline in one-year mortality.

[11] To the extent that quality of life may be lower for heart-attack survivors, these numbers should not be compared directly to QALY or quality-adjusted life year values that often range around $100,000 per life year for someone in perfect health.

The quality-adjusted price of health care is shown in Figure 3. As in Cutler et al. (1998), we find a decline in the real price of health care during the early 1990s. Since 1995, however, the real price of heart attack treatment has increased. The larger is the value placed on a life-year, the steeper is the initial drop in quality-adjusted prices, as is the more rapid the subsequent increase during the latter 1990s. Costs have continued to rise while mortality has flattened out since 1995, largely reversing the price reductions of the early 1990s. In other words, technological progress for treating heart attacks no longer appears to justify the increasing costs.

Is There Convergence in the Quality-Adjusted Price of Heart Attack Treatment?

We next consider trends in heart attack mortality, costs, and quality adjusted price for each of the 306 HRRs. In particular, we address the question of whether there was convergence in these factors across HRRs. The aggregate trends suggest that technological progress in the treatment of heart attacks slowed in the late 1990s. Standard models of diffusion would predict that at the outset of the technological innovation there will be an increase in the cross-regional variation as some regions adopt and learn about the new technology while others do not, but eventually (when technological progress slows) there will be a decline in productivity variation – because new adopters have already realized the productivity gains available and the late adopters or slow learners are catching up. While these models can generate growing variation in productivity initially, they always imply convergence in the standard sense: low initial productivity will be associated with faster subsequent growth as these producers learn.

In Figure 4, we see that there was no apparent convergence in terms of the variance in adjusted mortality rates across HRRs over the period. Figure 4 shows the standard deviation of HRR-level mortality rates by year from 1989-1999. There is a one-time decline that is explained by the move from a 20% sample (in 1991) to a 100% sample (in 1992); otherwise there is no evidence of convergence.

An alternative test of convergence is to ask whether areas that differ initially in terms of mortality (or some other measure) converge over time, e.g. are high and low mortality areas in 1989 just as different in 1999? To answer this question, we identified HRRs in the bottom and top quintiles of mortality averaged over 1989-1991, and plot trends in mortality, cost, and quality adjusted price through 1999 for each quintile (Figure 5A-5C). Averaging over the years 1989-1991 (which are 20% samples) helps to minimize the measurement error in mortality rates. However, the remaining measurement error is still large enough to generate an obvious, one-time, mean reversion in mortality between 1989-91 and later years (Figure 5A). Note, however, that all of the apparent convergence occurs in 1992 (as would be expected if the initial mortality rates in 1989-1991 were measured with error), and post-1992 there is no evidence of convergence. Nor was there any real difference in costs (Figure 5B); they trended upward for both the high-mortality and low-mortality regions. It is not surprising, therefore, that the quality-adjusted price was lower in the low mortality regions and shows no evidence of convergence (Figure 5C).[12]

It is possible that the spread in mortality rates (as in Figure 5A) is simply the result of unmeasured health differences, rather than reflecting the quality of care *per se*.

[12] Regressing HRR-level changes in mortality and cost on lagged levels (instrumenting for lagged levels with further lags to account for measurement error) also finds no evidence of convergence.

In this case, one could plausibly observe permanent non-convergent differences in mortality. However, if this were the case, we would not expect to observe divergence (or convergence for that matter) between any two groups of HRRs. Figures 6A-6C plot the mortality, cost, and quality-adjusted price of heart attack treatments for two groups of regions; those in the top quintile with regard to their adoption of β blockers in 1994/95 (for ideal patients) and those in the bottom quintile. These data come in the middle of the time-period considered in the sample, and thus may identify hospitals that adopted technology early in the period. Early and late adopters of beta blockers were initially similar in 1989 with regard to both mortality and costs, but mortality subsequently diverged while costs did not. During the 1990s, regions adopting β blockers experienced much sharper declines in mortality with no compensating increase in costs. By contrast, regions that did not adopt β blockers experienced relatively slow mortality improvements during the entire 1990s, so that by the end of the decade the differences in the value of treatments between the two groups of regions were nearly 10 percent of their remaining lifetime income.

A different pattern is shown in Figures 7A-7C, where HRRs were divided into quintiles based on the per-capita supply of cardiologists in 1992/93, from the *Dartmouth Atlas of Health Care 1996.* High supply and low supply areas differ little in mortality throughout the period. However, the cost and quality-adjusted price was initially higher where there was a high supply of cardiologists; in part this is the consequence of differential Medicare fees in urban areas because of DSH payments, regional cost adjusters, and graduate medical training costs. However, more importantly, these initial differences diverged during the period; by 1999, the difference in quality-adjusted prices

was nearly 15 percent of lifetime income. Thus, in contrast to areas with high beta blockers that were diverging on mortality, areas with a high supply of cardiologists were primarily diverging on costs.

Long-term Differences Across States in Technology Adoption

What can explain the non-convergence in the cost, mortality, and quality adjusted price of AMI during the 1990s? There are at least thee broad classes of potential explanations. One class of explanations would focus on chance events in learning and technology adoption that have accumulated over many years and dissipate slowly. Another class of explanations would focus on institutions particular to health care (e.g. lack of a market, specific reimbursement policies) that might discourage technology diffusion. A third class of explanations would focus on more generic barriers to technology adoption (e.g. size of the local market, extent of informational networks, or social barriers to change) that lead a given region to be persistently late adopters of all technologies. One such barrier suggested by our model is systematic differences in the discount rate λ, which may arise because of a lack of access to capital funding (either physical or human capital) or differences in hyperbolic discounting. Whatever the underlying reason, persistent barriers to technology adoption can generate long-term non-convergence, with some regions always being off of the productivity frontier.

As a simple test of these competing explanations, we compare regional rates of adoption of two dramatically different technologies: The use of β blockers (by state) for the treatment of heart attacks in 2001 (Jencks et al, 2003), and the adoption of hybrid corn by farmers in the 1930s and 1940s (Griliches, 1957). The adoption of corn hybrids is a useful comparison in that explanations focusing on chance events or institutions

specific to health care would not predict any relationship, whereas explanations focusing on persistent barriers to technology adoption would predict such a relationship. At the same time, the evidence on corn hybrids suggests that the basic diffusion process is similar in nature to β blockers: Both were relatively low cost technologies of proven effectiveness that nevertheless took decades to diffuse to all states. Despite well-documented under-use of β blockers (O'Conner et al 1997) and major education efforts, as of 2001 their use among ideal patients – which should be near 100 percent – was as low as 50 percent at the state level.

In Figure 8, we plot the state-level proportion of ideal patients receiving β blockers during 2001 (Jencks et al, 2003) along with the year in which each state attained at least 10 percent levels of use for hybrid corn varieties, based on the logisitics curves estimated in Griliches (1957). Despite more than one-half century that separates these adoption measures, and the very sharp differences in the nature of the outputs (and the markets for these outputs), the state-level adoption rates are strongly correlated (corr=-0.57); states that took longer to attain the 10 percent adoption level for hybrid corn (the horizontal axis) were also had much lower use of β blockers in 2001 (vertical axis). This evidence is quite suggestive that the persistent differences we see in the productivity of medical care across regions is the result of general barriers to technology adoption, not the result of chance events or factors specific to medical care.

5. Conclusion

We have considered the process of technological diffusion for the treatment of heart attacks in the elderly Medicare population. Most theories of technological diffusion imply the strong property of convergence; that over time, the lower quality regions will

adopt the newer technology and gradually catch up to the higher quality early-adopter regions. We have used the Cutler et al (1998) measure of quality-adjusted prices for heart attack patients to judge productivity gains over time. Using a panel of 2.6 million heart attack patients during 1989-1999, we found that the rapid declines in quality-adjusted prices found by Cutler et al (1998) during the late 1980s and early 1990s had largely ceased by the mid-1990s; since that time the quality-adjusted price of heart attacks has risen robustly.

Second, we did not find any empirical evidence that favored convergence; there were stable and persistent differences in mortality outcomes, and these could not be obviously attributed to unmeasured health status differences across regions. In some cases, as with the adoption of β blockers within an HRR, there was divergence in productivity across regions. The most likely explanation for this permanent difference in productivity is some form of persistent differences across regions in barriers to adopting innovations, as has recently been suggested in the medical literature (Berwick, 2003). Indeed, there was a strong correlation between states that were early adopters of hybrid corn in the 1930s and 1940s and states that were early adopters of β blockers in 2001.

There are strong parallels between the results from this paper and the evidence from cross-region or cross-country industry and firm studies that similarly suggest very slow (if any) convergence in productivity (Keller, 2001). We might expect even less convergence for health care markets, given the absence of national markets in their output. If export industries are unable to maintain productivity savings, they risk being priced out of the global market. In fact, shifting production to more efficient firms has been found to be an important channel through which market reforms affect productivity

growth in manufacturing (Pavcnik, 2002; Tybout, forthcoming). By contrast, there are no such forces operating currently in the health care system, given the difficulty in monitoring quality across disparate groups of individuals in different states. Thus it might be expected that convergence in health care may not be expected soon.

References

Berwick, Donald M., "Disseminating Innovations in Health Care," JAMA 289(15) (April 16, 2003): 1969-1975.

Bikhchandani, S., A Chandra, D. Goldman, and I. Welch, "The Economics of Iatroepidemics and Quackeries: Physician Learning, Informational Cascades, and Geographic Variation in Medical Practice," mimeo (2002).

Cutler, David M., Mark McClellan, Joseph P. Newhouse, and Dahlia Remler, "Are Medical Prices Declining? Evidence from Heart Attack Treatments," *Quarterly Journal of Economics* 93(4): 991-1024.

Cutler, David M., and Mark McClellan, "Is Technological Change in Medicine Worth It?" *Health Affairs* 20(5) (2001).

Fournier, Gary M., Kislaya Prasad, and Mary A. Burke, "Physician Social Networks and Treatment Variations in Coronary Inpatient Care," mimeo, Florida State University (May, 2002).

Griliches, Zvi, "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica* 25(4) (October 1957): 501-522.

Klausen, Liv Marit, Trond E. Olsen, and Alf Erling Risa, "Technological Diffusion in Primary Health Care," *Journal of Health Economics* 11 (1992): 439-452.

Jencks, Stephen F., Edwin D. Huff, and Timothy Cuerdon, "Change in the Quality of Care Delivered to Medicare Beneficiaries, 1998-99 to 2000-2001," *JAMA* 289(3) (January 15, 2003): 305-312.

Jovanovic, Boyan, and Yaw Nyarko, "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," *Brookings Papers: Microeconomics* (1995) :

Jovanovic, Boyan, and Yaw Nyarko, "Learning by Doing and the Choice of Technology," *Econometrica* 64(6) (November 1996): 1299-1310.

Phelps, Charles E., "Information Diffusion and Best Practice Adoption," in A.J. Culyer and J.P. Newhouse (eds.) *Handbook of Health Economics* Volume 1, Elsevier Science (2000): 223-264.

Trajtenberg, Manuel, "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners," *Journal of Political Economy* 97(2): 444-479.

Figure 1: One-Year and 30-Day Mortality Following AMI

[Adjusted for Age/sex/race and comorbidities *and* type of AMI]

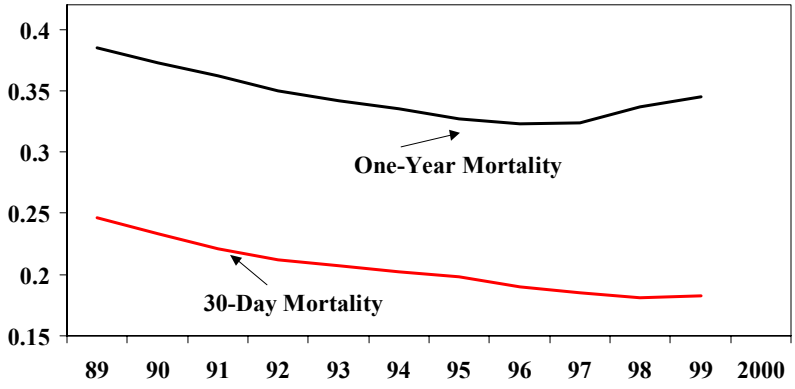

Figure 2: One-Year Costs (Part A) Following AMI

[Adjusted for Age/sex/race, comorbidities, and type of AMI]

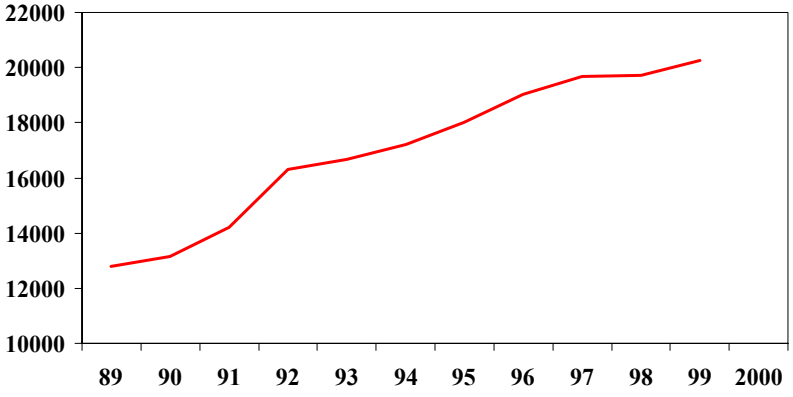

Figure 3: Quality-Adjusted Price of AMI Treatment

[Life-Year = $29,804 or $59,609]

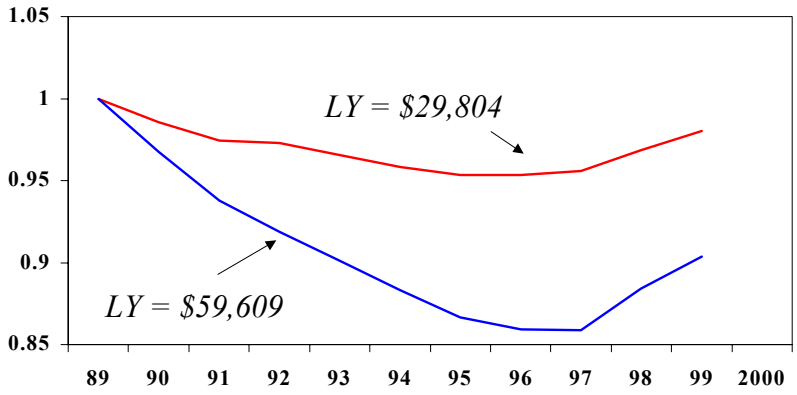

Figure 4: Standard Deviation in Mortality Rates
Across HRRs, 1989-1999

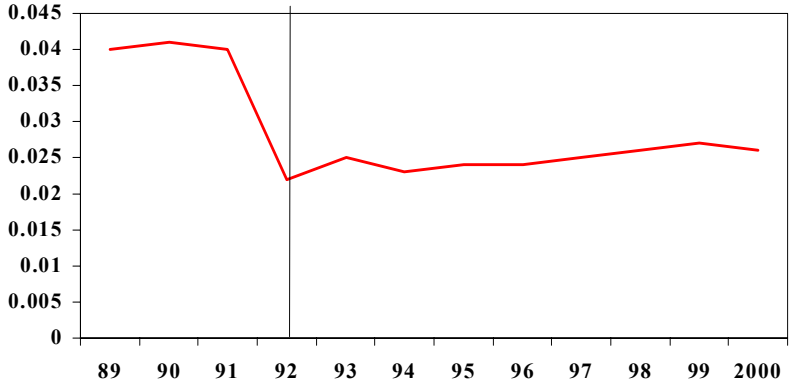

Figure 5A: Mortality By Lowest and Highest Predicted 1993 Mortality Quintile

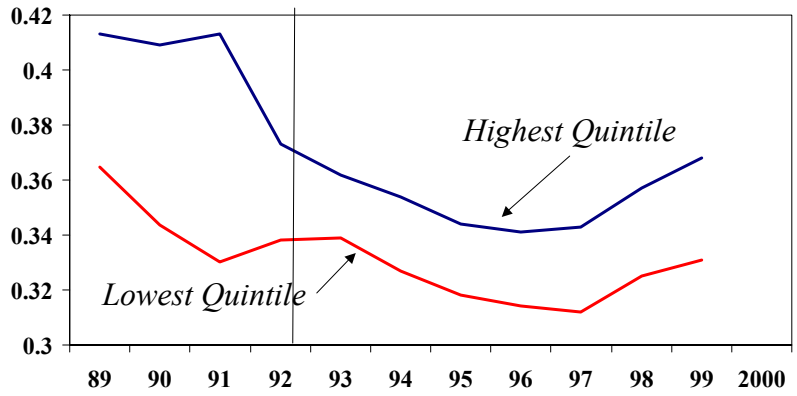

Figure 5B: Costs by Highest and Lowest Predicted 1993 Mortality Quintile

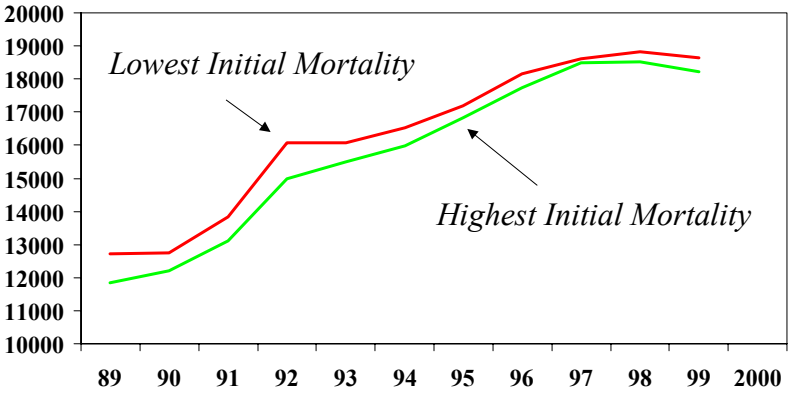

Figure 5C: Quality Adjusted Price
By Highest and Lowest 1993 Mortality Quintile

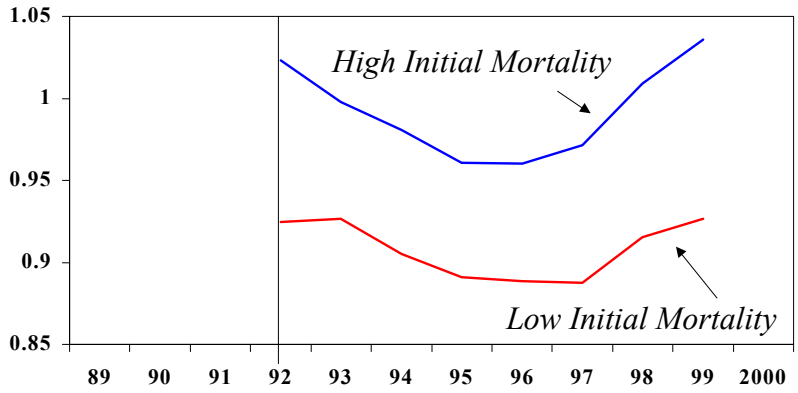

Figure 6A: One-Year Mortality Following AMI: By Quintile of Beta Blocker Use
[Adjusted for Age/sex/race, comorbidities, and AMI type]

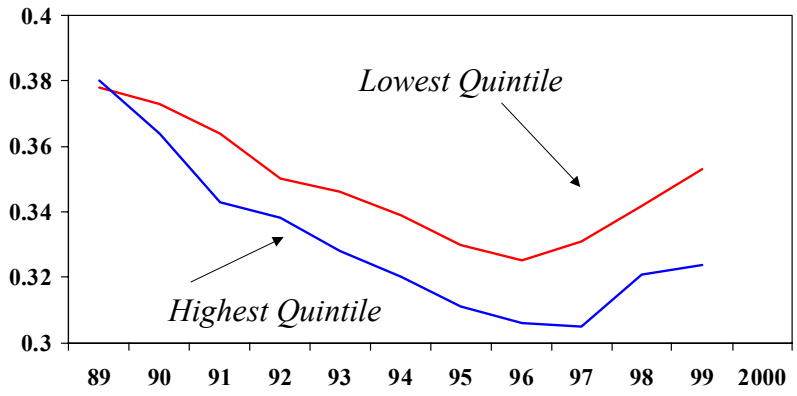

Figure 6B: One-Year AMI Costs, by Quintiles of Beta Blocker Use (1993)

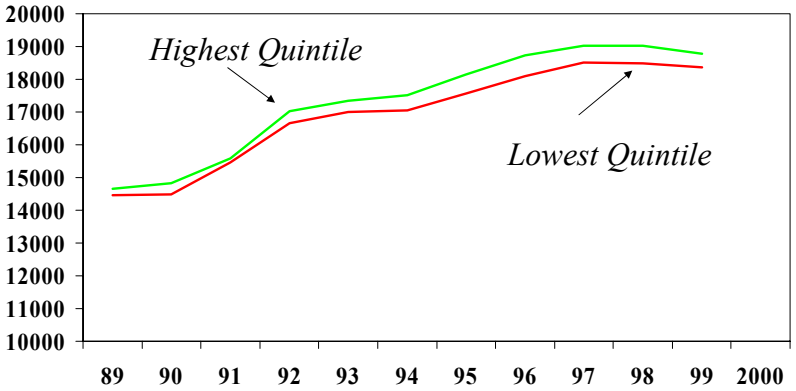

Figure 6C: Quality-Adjusted Price of AMI Treatment
By Regional Use of Beta Blockers in 1994/95
[Life-Year = $59,609]

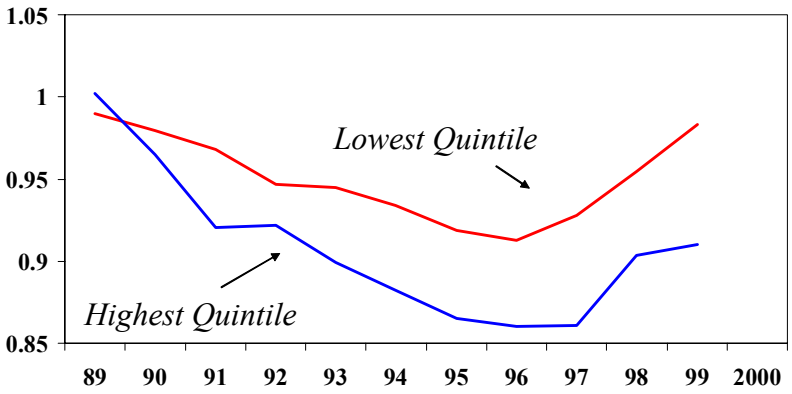

Figure 7A: One-Year Mortality Following AMI: By Quintile of Cardiologists

[Adjusted for Age/sex/race, comorbidities, and AMI type]

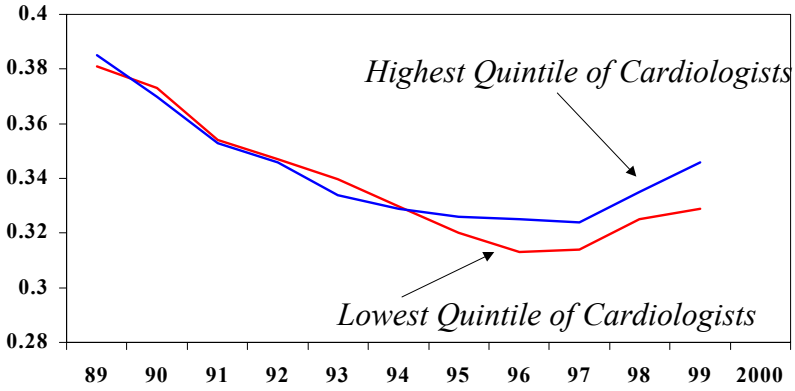

Figure 7B: One-Year AMI Costs, by Regional Per Capita Supply of Cardiologists (1993)

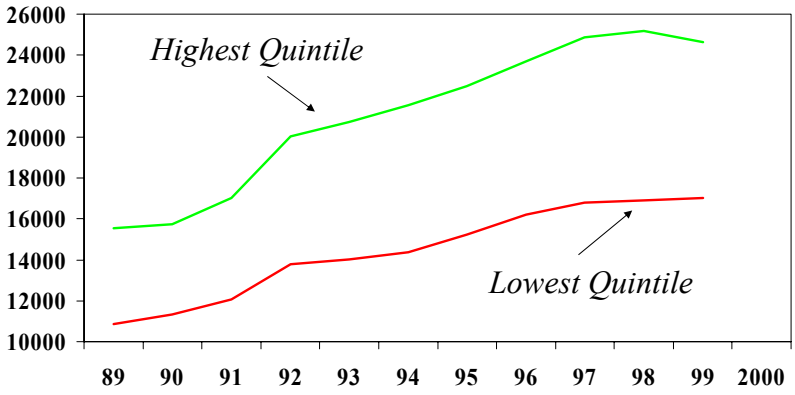

Figure 7C: Quality-Adjusted Price of AMI Treatment
By Regional Per Capita Supply of Cardiologists
[Life-Year = $60,943]

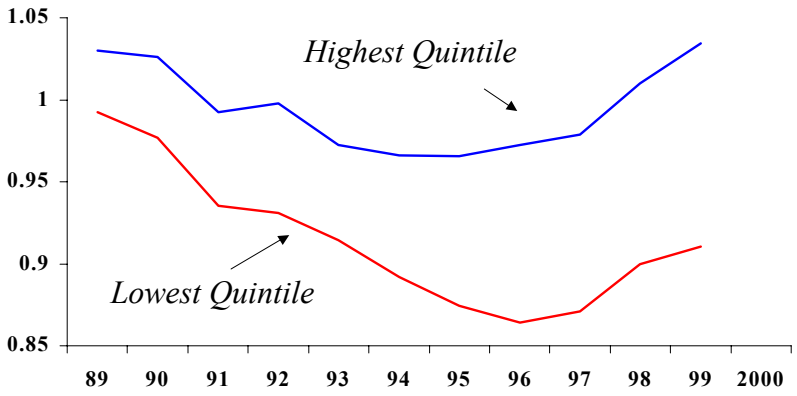

Figure 8: Year Hybrid Corn Use Attained 10% and Use of Beta Blockers Within 24 Hours of Acute Myocardial Infarction in 2001: By State

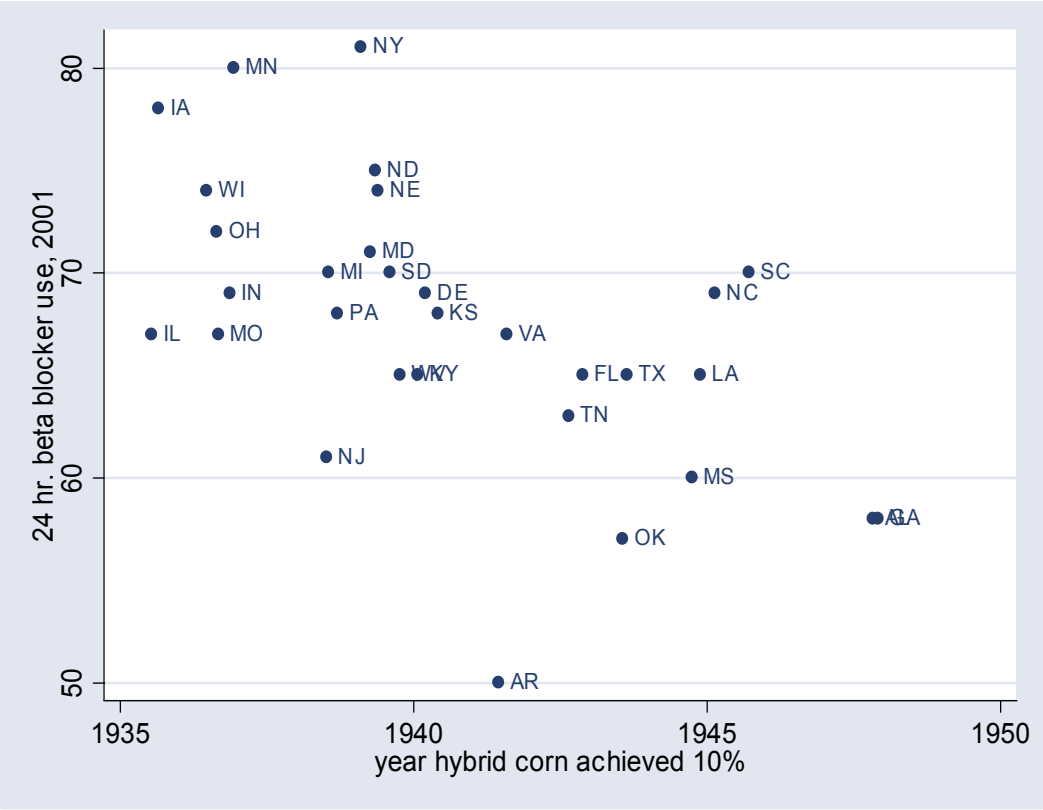

Source: Griliches (1957) and Jencks et. al. (2003).