THE EVOLUTION OF PHYSICIAN PRACTICE STYLES:
EVIDENCE FROM CARDIOLOGIST MIGRATION

David Molitor

The Evolution of Physician Practice Styles: Evidence from Cardiologist Migration
David Molitor
NBER Working Paper No. 22478
August 2016
JEL No. H51,I11,I18

ABSTRACT

Physician treatment choices for observably similar patients vary dramatically across regions. This paper exploits cardiologist migration to disentangle the role of physician-specific factors such as preferences and learned behavior versus environment-level factors such as hospital capacity and productivity spillovers on physician behavior. Physicians who start in the same region and subsequently move to dissimilar regions practice similarly before the move, but each percentage point change in practice environment results in an immediate 2/3 percentage point change in physician behavior, with no further changes over time. This suggests environment factors are twice as important as physician-specific factors for explaining regional disparities.

David Molitor
University of Illinois at Urbana-Champaign
340 Wohlers Hall
1206 S. Sixth Street
Champaign, IL 61820
and NBER
dmolitor@illinois.edu

1 Introduction

Health spending per capita varies dramatically across U.S. regions. For example, age, race, sex, and price adjusted spending in Medicare's traditional fee-for-service program in 2012 was \$13,596 per enrollee in the Miami, FL region compared with \$7,998 in the Minneapolis, MN region.[1] These spending disparities arise primarily from regional differences in the types and quantities of services patients receive (Skinner and Fisher, 1997; Gottlieb et al., 2010). Spawned by Wennberg and Gittelsohn's (1973) classic study finding ten-fold differences in tonsillectomy rates across Vermont towns, an enormous literature has consistently documented widespread variability in cross-regional rates of hundreds of medical interventions within a variety of patient populations and institutional contexts (Phelps, 1992).[2]

Despite extensive research documenting regional variations in health care delivery, relatively little is known about their causes. Direct adjustments to reflect apparent differences in average levels of patient illness, socioeconomic status, or preferences typically resolve little of the variations (Barnato et al., 2007; Zuckerman et al., 2010). Moreover, a variety of evidence suggests that the quality of care and health outcomes in high-use regions are little better or even worse than in low-use regions (Fisher et al., 2003a,b; Baicker and Chandra, 2004; Sirovich et al., 2006). A common interpretation of this fact is that additional health spending yields little or no health benefit, implying that moving high-use regions to behave like low-use regions could lower overall spending by 30 percent without sacrificing quality of care (Wennberg et al., 2002). But in order to address how to change patterns of care—or to assess whether changes are even desirable—it is essential to understand what drives these regional variations. Because patterns of care ultimately arise from the accumulation of decisions individual physicians make about which procedures to prescribe their patients, a more

---

[1]Price adjustments remove regional differences in Medicare reimbursement rates, such as higher payments to hospitals with medical training programs. Skinner et al. (2011) describe the Medicare regional spending measurement methodology. Regional spending for 2012 accessed from http://www.dartmouthatlas.org.

[2]Regional disparities are also prevalent in Medicaid, the centrally budgeted Veterans Affairs health system, and the private sector (Martin et al., 2007; Congressional Budget Office, 2008; Philipson et al., 2010). The study of medical practice variations began with Glover (1938), who analyzed regional tonsillectomy rates of British school children.

fundamental question is what drives physician treatment decisions.

This paper explores the role of the physician versus his practice environment in explaining regional differences in how physicians treat similar patients. Environment-specific factors such as financial and legal incentives, hospital capacity, and productivity spillovers extend influence across local groups of physicians, and therefore may drive practice style differences across practice settings. On the other hand, physician-specific factors such as preferences, training, and experience may cause physicians to treat patients differently even under similar environments. Consistent with this possibility, physicians practicing in the same local health care market often exhibit large and persistent "style" differences in their tendency to prescribe certain treatments and utilize medical resources (Phelps, 2000; Grytten and Sørensen, 2003; Epstein and Nicholson, 2009). These styles exist even when physicians have access to the same hospital facilities and ancillary staff and when the patients are randomized to physician teams (Doyle et al., 2010). If physicians agglomerate geographically based on individual-level factors that drive practice styles (e.g. physicians practicing close to where they were trained, or physicians in the same region accumulating similar experiences), physician-specific factors could drive practice style differences across regions.

At least two conceptual issues have hampered empirical investigations attempting to separate effects of the environment from those specific to the physician. The first is that physician factors such as training and experience may form endogenously in response to the physician's environment. This issue can be at least partially resolved by looking at factors such as residency training that pre-date the current environment (Dranove et al., 2011). However, even when historical physician information is available, a second and potentially more substantial identification issue is that physicians may choose a practice setting based on their individual practice style or otherwise correlated with physician-specific determinants of practice style. Failure to account for such "positive matching" may yield estimates that overstate the effect of the environment on physician behavior.

The primary contribution of this paper is to exploit an empirical context providing varia-

tion in a physician's environment while also allowing explicit controls for physician selection. Using 15 years of Medicare patient claims, I construct histories of treatment decisions for individual physicians and identify a set of physicians who move across geographic regions. I then trace out how migrant behavior changes over time with respect to the move as a function of the change in environment experienced across the move. Selective migration is identified by the extent to which physicians who move to higher or lower intensity regions have differential levels or trends in pre-move behavior relative to their peers. The full environment effect is identified by the change in physician behavior across the move as well as by the subsequent time-pattern of behavior relative to the move date.

Using this approach, I begin by testing two polar scenarios. First, I test whether physician practice styles are fully ingrained once physicians have completed medical training and taken up clinical practice. If this is the case, then changes in a physician's practice environment should not affect how the physician treats similar patients. Second, I test the other extreme of whether physicians completely conform to changes in their environment regardless of their training or past experiences. Full convergence would point to steady-state differences in regional practice styles arising from differences in the contemporaneous influences under which physicians operate. If physicians do not completely conform to environment changes, however, then physician behavior is persistent and small changes in their early training or experience could have long-run effects. Finally, an additional advantage of the empirical approach I employ is that it not only allows me to test whether either of these polar scenarios holds true but also provides an estimate for where reality lies between the two.

The specific context of my study is cardiologists treating heart attack patients. The data include 19,945 cardiologists treating patients over the period 1998-2012. Of these, 3,089 (15.5%) are observed to move their practice location across geographic medical markets. Cardiologists may choose to treat heart attack patients with an "aggressive" approach marked by early patient receipt of an invasive procedure called cardiac catheterization, or they may follow a "conservative" approach using medical management (drugs). Consistent

3

with previous studies on geographic variations (e.g. Gatsonis et al., 1995), I find that the share of heart attack patients receiving aggressive treatment over the sample period varies considerably across geographic regions, averaging 0.43 with an interquartile range of 0.10.

In my key empirical analysis, I find that for cardiologists who move, a change in a physician's practice environment results in a significant and rapid change in the physician's individual practice style. Specifically, if O and D represent the fraction of patients treated aggressively in a physician's respective origin and destination practice regions, then the physician's individual propensity to treat aggressively changes by $2/3(D - O)$ on average across the move. Moreover, this change in behavior occurs immediately after a physician's move with no additional changes over time, suggesting that further learning or adaptation is limited. Finally, I fail to find evidence of physician selection—cardiologists who move to more-aggressive regions appear no more aggressive than their peers prior to the move. These results reject both polar views discussed above: physicians respond to changes in their practice environment, but do not completely conform to these changes. The estimated change in physician behavior implies that both the environment and the physician influence treatment choices, with the environment playing twice as large a role as physician-specific factors.

Next, I explore the nature of physician behavior changes in greater detail to shed light on the mechanisms underlying regional heterogeneity in practice styles. One predominant theory used to explain the existence and persistence of regional practice variations is the Phelps and Mooney (1993) "schools of thought" model of information diffusion in which physician practice styles initially form during training and evolve over time according to a Bayesian learning process as physicians are exposed to new environments. In contrast to the implications of this model, I find that physician behavior responds discretely to changes in their environment with no further convergence over time. I also find that physicians who move later in their career respond about the same to changes in their environment as those who move early in their career, suggesting physician practice styles remain elastic over time.

Finally, I explore whether physicians respond asymmetrically to changes in their en-

vironment and find that physicians moving from a more-intensive region retain more of their previous practice style (i.e. change behavior less) than physicians moving from a less-intensive region. This suggests that hard technological capacity constraints, such as lacking a catheterization laboratory, are not the key driver of physician treatment choices in this context. Further supporting this view, 89% of heart attack patients in the sample are admitted to hospitals with cath labs, and the estimates of physician response across a move change little when limiting the analysis to this subsample.

My approach in this paper is closely related to a growing literature that uses migration patterns to isolate the effects of culture and past experiences from the current environment on consumer preferences and choices (Fernández, 2008, provides a review). Finkelstein et al. (2014) exploit migration of Medicare patients to isolate the role of patient demand in driving geographic variation in health care utilization. Outside the health care context, Ichino and Maggi (2000) use worker movements across branches in a firm to identify the impact of group interactions on shirking behavior; Song et al. (2010) use patient migration across geographic regions to identify regional diagnosis propensities separately from patient characteristics; and Chetty et al. (2011) track tax payers across a move to identify local neighborhood effects on worker response to the EITC. My approach is also closely related to the "brand capital" model of Bronnenberg et al. (2012), in which consumer purchase choices depend not only on contemporaneous supply factors but also on brand exposure in the past. Analogously in my context, physicians accumulate "treatment capital" that may influence their treatment choices holding constant the practice environment. To the best of my knowledge, this paper is the first to exploit physician migration patterns to separately identify the role physician-specific and environment-level factors play in determining physician practice styles.

The organization of the remainder of the paper is as follows. Section 2 describes the empirical context and key data elements, and Section 3 lays out the empirical strategies and results. In Section 4 I briefly discuss potential mechanisms, and Section 5 concludes.

2 Setting and Data

2.1 Context: Heart Attack Treatment

Each year nearly 1 million Americans suffer a heart attack, resulting in more than 130,000 deaths.[3] Heart attacks, referred to clinically as acute myocardial infarction (AMI), occur when part of the heart's blood supply is blocked, starving the heart of oxygen and causing muscle cells to die. Heart attacks are an emergency condition and require immediate hospitalization. While there are a variety of heart attack treatments, all essentially amount to reducing the heart's demand for oxygen and increasing blood supply to the muscle. To increase blood supply, doctors may either use medical management (drugs) or take an invasive approach. In the medical approach, thrombolytic "clot-busting" drugs are used to dissolve blood clots blocking coronary arteries and are typically most effective when administered within 3 hours after the heart attack occurs. The primary invasive techniques to restore blood flow to the heart are angioplasty (balloon dilation of the blocked artery, with or without stenting) and open-heart bypass surgery (artery graft to "bypass" the blockage).

To determine whether a patient is a candidate for an invasive procedure, the doctor must identify the precise location and severity of blockages. This can be accomplished through a diagnostic technique called angiography. This procedure is usually included as part of a cardiac catheterization (often referred to simply as a "cath") in which a thin catheter is threaded into the coronary arteries. Contrast dye is injected through the catheter into the blood stream, while x-ray video cameras track the flow of dye to reveal areas where the coronary arteries are severely restricted or blocked. In this role, cardiac catheterization is commonly used and well-understood as a marker for invasive heart attack treatment (see e.g. McClellan and Newhouse, 1997; Chandra and Staiger, 2007).

The empirical work in this paper focuses on AMI treatment for four reasons. First, heart attack treatment is characterized by two competing management approaches: an "early inva-

---

[3]Source: CDC, http://www.cdc.gov/NCHS/data/nvsr/nvsr58/nvsr58_19.pdf. Death count from 2007.

sive" approach marked by patient catheterization shortly after hospital admission regardless of the patient's receipt of or response to thrombolytic therapy, and a conservative "wait-and-see" approach in which patients are first given thrombolytic drugs and receive cardiac catheterization only if symptoms persist. Both approaches have been heavily analyzed and debated in the medical literature (see e.g. Keeley and Grines, 2004; Brophy and Bogaty, 2004; Scanlon et al., 1999). Since early versus delayed cardiac catheterization is typically defined with reference to a 12- to 48-hour time window (Kushner et al., 2009), I use receipt of catheterization within 2 days of AMI hospital admission as the measure of early invasive management. This dichotomy allows both regions and physicians to be characterized by the management style choices (i.e. cath rate) used for their patients.

Second, the rate of invasive heart attack management varies significantly across regions, exposing cardiologists who move to potentially large changes in their practice environment. Third, the emergency nature of heart attacks generally inhibits patients from traveling long distances to seek care, making it possible to define geographically distinct markets for AMI treatment in which physicians practice. Finally, the emergency nature of heart attacks also plausibly limits the degree to which patients most appropriate for a particular type of treatment are sorted to cardiologists who specialize in that treatment.

2.2 Data Description

The primary data for the analysis is Medicare administrative and claims records for the Medicare fee-for-service population over the period 1998-2012. The data include a 100% sample of hospital admissions records, which are used to identify over 4 million patients with new heart attack episodes (at least one year since any previous heart attack) based on a principal diagnosis of AMI (ICD-9-CM codes 410.x). While these records cover the universe of fee-for-service beneficiaries over this period, I cannot observe treatment outcomes for patients in Medicare Advantage plans which are reimbursed on a capitated basis.[4] AMI

---

[4]Over the sample period 1998-2012, Medicare Advantage covered 19.9% of Medicare enrollees, although the share fluctuated over that period. See the Appendix for additional sample details.

patient hospital records are matched to physician claims to identify the physicians treating the patient, which limits the sample to the 20% of beneficiaries for whom physician claims are available. Medicare claims are available beginning in 1992, but I exclude years prior to 1998 both because physician claims are only available for 5% of Medicare beneficiaries those years and because fewer hospitals had cath labs during that period.

2.2.1 Cardiologist catheterization rates

Behavior of individual cardiologists over time and across practice settings is identified in the data using a physician's Unique Physician Identification Number (UPIN) on billing claims. A UPIN is given to each physician who treats patients in the Medicare program and remains with the physician throughout his or her career.[5] I link the universe of Medicare UPINs to the American Medical Association (AMA) Physician Masterfile and identify cardiologists as those who have completed a 3-year fellowship in cardiovascular disease.[6]

I measure the cath behavior of cardiologists over time by assigning AMI patients to the first cardiologist treating the patient. While the first cardiologist's decision is only one of many in the hospital setting that may affect patient treatment, identifying a patient with the first cardiologist minimizes concerns of selective sorting of patients to cardiologists—a typical emergency room protocol is to initially assign a confirmed or suspected AMI patient to the cardiologist on call. Moreover, due to the emergency nature of heart attacks and high time-sensitivity of the relative benefits of different treatment paths, the initial cardiologist is likely to have an important impact as a "gatekeeper" to subsequent care the patient receives, whether or not this cardiologist actually performs the services.

To implement the assignment of patients to cardiologists, I focus on the sample of Medicare fee-for-service beneficiaries for whom physician claims are available. I then identify patients with a new AMI episode who see a cardiologist within two days of hospital admis-

---

[5]Beginning in 2007, Medicare transitioned from UPINs to the National Provider Identifier (NPI) standard. I match NPIs to UPINs using a crosswalk developed by the NBER, available at http://www.nber.org/data/npi-upin-crosswalk.html, and supplemented with Medicare claims that contain both fields.

[6]The AMA Physician Masterfile includes current and historical data on virtually every Doctor of Medicine (MD) ever trained or licensed to practice in the United States, regardless of physician AMA membership.

sion. For each patient, I identify the cardiologist(s) who treat the patient first. Because claims only identify the day of service, some patients (34%) match to multiple "first" cardiologists. Seeing more than one cardiologist on the first day may itself depend on the initial physician's treatment choice. I therefore use all patient episode-physician pairs in the baseline analysis and focus on the 66% of patients that see a unique first cardiologist in robustness checks. Over the analysis period 1998-2012, I observe 19,945 cardiologists treating 669,397 patient heart attacks (see Table 1).

2.2.2 Cardiologist migration

To identify movers, I focus on cardiologists who move their practice across Hospital Referral Regions (HRRs), geographic units developed by the Dartmouth Atlas of Health Care and commonly used as the regional unit of analysis for heart attack treatment (e.g., Skinner et al., 2006; Chandra and Staiger, 2007). HRRs partition ZIP codes into 306 regions based on where the majority of Medicare beneficiaries are referred for tertiary health care services, and each HRR contains at least one hospital performing major cardiovascular procedures.

I base a cardiologist's practice location at a point in time on the dates and hospital of admission for the physician's patients. I define "practice episodes" to be the first and last date a cardiologist practiced in a given HRR during 1998-2012 and limit to episodes where the cardiologist treated two or more AMI patients. I mark the practice episode during which that physician treated the most patients as the cardiologist's "primary" episode. Similarly, I further define a cardiologist's "secondary" practice episode to be the largest episode (in terms of patients treated) that does not overlap the primary episode, if such an episode exists. Movers are those with both primary and secondary practice episodes. Of the 19,945 cardiologists in the baseline analysis file, 3,089 (15.5%) are identified as movers.[7]

As reported in the last column of Table 1, approximately 1-2% of cardiologists observed

---

[7]This definition of a move requires a clean split in time between the origin and destination HRR. If a cardiologist practices in HRR A from dates $d_1 - d_2$ and HRR B from dates $d_3 - d_4$, this would be considered a move as long as $d_2 \leq d_3$. However, if $d_2 > d_3$, which could happen if the cardiologist returns to practice in HRR A after first switching to HRR B, this would not be marked as a move.

each year are also identified as moving that year. Table 2 summarizes the migration patterns in the sample: nearly 80% of migrants move across states, and over 45% move across Census Regions; over half the moves occur within the Midwest and South.

Table 3 compares the migrant physician population to other physicians, by region and overall. The first two columns describe the proportion of cardiologists in each census region that move either out of (emigration) or into (immigration) an HRR in that region, relative to the total number of cardiologists who ever practice in that region. Even when measured as a proportion of total physicians, the Midwest and South continue to show the most migration activity. The highest net emigration occurs out of the Midwest (3.1%), and the highest net immigration occurs into the West (4.5%). As shown in the last four columns, migrants are also slightly more likely to be female and foreign-born than their non-migrant counterparts.

Table 3 also compares the time lapse between a migrant's observed move and completion of a cardiology disease (CD) fellowship, relative to the distribution of time since CD fellowship completion for non-migrants in the sample. Half of migrants move within 8 years of finishing CD training; comparatively, the median number of years since CD training for non-migrants in the sample is 14 years. The number of years since CD training can also be roughly converted to physician age, since cardiologists who go straight through training typically finish at age 32-33 (after college, physicians must complete 4 years of medical school, 3 years internal medicine residency, and a 3 year CD fellowship). Thus, the median age in the sample is approximately 41 for migrants (at the time of the move) and 47 for non-migrants.

2.2.3 HRR and hospital catheterization rates

The central aim of this paper is to understand how individual physician actions respond to changes in their environment, where the environment is characterized by the average action (2-day cath) taken by physicians in that region. As highlighted first by Manski (1993) and more recently by Angrist (2014), individual actions are highly correlated with group average actions, potentially leading to spurious conclusions that a causal connection exists.

To address these concerns, I define regional 2-day cath rates experienced for each cardiolo-

gist using a leave-out average of cath choices that omits the cardiologist's own patients. This eliminates any mechanical correlation between a migrant cardiologist's treatment choices and experienced regional cath rates. To further minimize the impact any one physician is likely to have on practice patterns in the region, my preferred definition of a cardiologist's local practice region is the hospital market as defined by HRRs. Moreover, the key independent variable in the analysis is the change in regional cath rates experienced across a cardiologist's move. If a cardiologist's individual impact on the regional average is similar before and after move, the net effect after taking differences will be small.

Using HRRs to define local practice environments is preferable to political geographic divisions, such as states or cities, because HRR boundaries derive empirically from patient referral patterns. HRRs may also be preferable to finer definitions of a physician's relevant practice region such as the hospital for at least three reasons. First, cardiologists frequently hold operating privileges at multiple hospitals within a region at any point in time. Second, regional influences outside a physician's own hospital may also influence a physician's treatment behavior such as proximity to surgical backup, the ability to refer patients to nearby hospitals for treatment, and peer effects through professional and social interactions. Third, a broader definition of practice region such as the HRR minimizes the impact of any one physician on regional practice patterns and the scope for endogenous physician sorting.

While there are a number of advantages to defining physician practice region at the HRR level, as a complementary approach I also define hospital-level measures of the practice environment based on the hospitals where a physician's patients are admitted. While there may be less scope for cardiologist sorting at the HRR versus hospital level, using changes in the hospital environment across a physician's move may yield more precise estimates in the regression analysis compared to using changes in the more aggregate HRR environment measure. The empirical analysis below addresses additional trade-offs between using the HRR versus the hospital as the physician's relevant environment.

To account for potential regional differences in patient severity, I risk-adjust the raw

regional cath rates. As highlighted by (Song et al., 2010), a challenge in risk-adjusting regional treatment intensity is that diagnostic practices and the comorbidities they indicate may themselves be a function of regional treatment intensity. To avoid this issue, I adjust raw regional cath rates using indicators for patient age, race, sex, and first heart attack since the measurement of these characteristics is plausibly unrelated to regional treatment choices.

A key simplification I make is to use a time-invariant regional cath intensity measure over the period 1998-2012. As shown in Table 1, 2-day cath rates increased from an average rate of 34.0% in 1998 to 57.6% in 2012. While this secular trend in cath rates implies that cath rate *levels* are not directly comparable across years, the relative stability of the interquartile range over time implies that regional *differences* in cath rates are roughly comparable across years. To the extent that the intensity of an HRR relative to the secular trend remains stable over time, differences in HRR cath propensies over the pooled years will be the same as the difference in propensity in any given year. Time-invariant rates also have the advantage of being calculated over a larger sample, reducing sampling error in the estimates and further minimizing the extent to which the style of any doctor or group of doctors influences the cath rate in that region.

One disadvantage of using time-invariant regional cath rates is that if the relative intensity of a region changes over time (e.g. if HRRs have different growth rates in their cath propensities), then the average regional intensity ranking will be a noisy measure of the actual ranking in any given year. In the appendix, I provide additional details on the methodology for calculating both HRR and hospital 2-day cath rates, and also show that results based on time-invariant cath measures are similar to cath rates that vary by time.

Figure 1 maps the geographic distribution of 2-day cath rates across HRRs, with rates ranging from less than 41.1% in the lowest quintile of regions to more than 53% in the highest quintile. The change in physician j's practice environment experienced across a move is calculated as $\Delta_j = (destination\ region\ cath\ intensity)_j - (origin\ region\ cath\ intensity)_j$, where physician j's own patients are omitted when calculating the regional cath rates. Because

a large fraction of cardiologist moves occur in the Midwest and South where there is rich geographic variation in regional cath rates, the migrants in the sample face a wide spread of environment changes as shown in Figure 2. Panel A shows the change in HRR cath intensity across the move, while Panel B shows the change in hospital cath intensity. Both distributions of changes center close to zero, indicating that roughly equal numbers of physicians move to more- versus less-intensive regions. The spread in the distribution of hospital-level changes is roughly twice as large as that of HRR-level changes, consistent with a significant amount of within-region variation in cath intensity across hospitals.

3   Empirical Evidence

3.1   Difference-in-differences

3.1.1   Empirical specification

My primary empirical specification is a difference-in-differences "event study" of physician treatment decisions across a move. The key idea is to follow a physician in a long panel before and after move, and to trace out the level and time pattern of behavior with respect to the move. By comparing migrants who start out in the same region, I test for selective migration by observing whether the physicians who move to more intensive regions were already practicing more intensively than average prior to the move, or whether they started practicing more intensively shortly before the move. I then look at the change in physician behavior across a move to identify the effect of a change in environment on a physician's behavior. The key identifying assumption here is that nothing other than the environment changes simultaneously with the move that is correlated with the change in environment and also affects physician behavior.

To construct the event study, I measure each migrant cardiologist's cath behavior with respect to "event time" t, where t is the number of years since the physician's move. The event study is estimated using a regression where the dependent variable $(cath)_{ijt}$ is an indicator for whether heart attack patient i, treated by cardiologist j in event year t, received a cath within

2 days of hospital admission. The key interest is in how the migrant's change in environment Δ_j explains the physician's behavior over time. This is calculated by interacting Δ_j with the full set of event time dummies $\mathbf{1}(s = t)$. Each migrant's behavior is measured relative to baseline migrants in the same origin HRR by including a full set of physician origin HRR dummies and event time dummies $\mathbf{1}(s = t)$. Fixed effects for year of patient admission and patient age, race, sex, and first heart attack control for secular changes in cath propensity over time and observable differences in patient appropriateness for cath. Observations are limited to the treatment choices of migrant physicians within 8 years before or after move, yielding the regression equation

$$(cath)_{ijt} = \{origin\ HRR\ FEs\}_j + \sum_{s=-8}^{7} \left[\alpha_t \mathbf{1}(s = t) + \beta_t \Delta_j \mathbf{1}(s = t) \right]$$
$$+ \{calendar\ year\ FEs\}_i + \{patient\ risk\text{-}adjusters\}_i + \epsilon_{ijt}, \tag{1}$$

The main parameters of interest are the β_t coefficients. For a given value of t, β_t describes the difference between treatment styles of physicians t years since move per unit difference in Δ_j. If there is little selective migration, then physician styles prior to move should not differ systematically with Δ_j, and thus β_t should be close to zero for all $t < 0$. For $t \geq 0$, β_t describes how physician styles diverge in relation to Δ_j after a move. In combination with the information on selective migration uncovered in pre-move behavior, post-move behavior informs us how much physicians respond to changes in the environment. Namely, any break in the level of β_t across the move is informative about the extent of influence that the environment exerts on an individual physician's behavior. Moreover, the time pattern of any environment effect is informative about the mechanisms underlying this effect: immediate effects suggest that discrete factors such as the local availability of capital or peer effects are important determinants of physician style, whereas effects that increase over time suggest that "slow-moving" factors such as learning or adaptation play a role.

As a supplement to the event study, I also consider a traditional style difference-in-

differences (DD) estimate of how a change in environment affects physician behavior. This is implemented by replacing the event time dummies in Equation 1 with a single "after" dummy $\mathbf{1}(t \geq 0)$. The DD approach requires a parallel trends assumption that, absent a move, physician trends in behavior would have been the same for physicians who in fact moved to more-intensive regions as those who moved to less-intensive regions. The event study can boost the plausibility of this assumption by validating whether the assumption holds at least during the 8 years prior to a move (i.e. whether β_t is roughly flat for $t < 0$).

While the DD approach inherits its validity from the event study, there are at least two reasons for computing the DD estimate in addition to the event study. First, it provides a single summary measure of the effect of a change in environment. A second reason is that by lumping observations before and after move, the DD effectively computes the environment effect over a larger sample, yielding tighter estimates. Thus, the DD provides more power for adding additional controls, and makes it easier to compare the sensitivity of the results across different specifications.

For both the event study and traditional DD specifications, I include origin HRR fixed effects in the baseline specifications that measure changes in HRR cath rates across the move. A potential weakness of this specification is that it essentially measures physician behavior within groups defined by origin and destination HRR pairs, which potentially confounds interpreting measured changes in β_t over time as changes in individual physician behavior. This could occur, for example, if physicians tend to treat a higher volume of patients when the average practice intensity in their HRR is closer to the physician's individual preferred style. In this case, even if individual physicians changed their practice style little across a move, the average treatment used by a group of physicians who move between the same origin and destination HRRs could change due to a compositional shift in the fraction of patients each physician treats.

I resolve these concerns by considering alternative specifications that include physician fixed effects. When physician fixed effects are included, all changes in a physician's behavior

over event time are measured with respect to that physician's behavior in a baseline period, chosen to be the year immediately before the move (implying $\beta_{-1} = 0$). I also include physician fixed effects in all specifications that use changes in the hospital environment across the move, since a natural origin control group is difficult to define.[8] However, because physician fixed effects normalize $\beta_{-1} = 0$ mechanically, the β_t coefficients for years prior to the move no longer identify selective migration based on levels, though they will still capture differential trends. Thus, I prefer using origin HRR fixed effects to evaluate selective migration, but physician fixed effects to measure changes in physician behavior.

Finally, for both the event study and DD estimates, I compute two-way clustered standard errors at the physician and HRR levels. This accounts for potential serial correlation at the physician level and spatial correlation at the hospital market level.

3.1.2 Event study results

Figure 3 plots (solid black line) the sequence of β_t estimates from Equation 1, based on differences Δ_j^{HRR} in the HRR cath environment experienced across a move. The pattern highlighted in this figure is that the sequence of β_t estimates is roughly flat and close to zero before the move ($t < 0$) and then jumps discretely at $t = 0$ and thereafter remains roughly flat near 0.66. Error bars indicate 95% confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels.

I focus first on the β_t coefficients for $t < 0$. These estimates show whether there is any pre-move difference either in levels or trends between physician styles as a function of where the physician moves (as described by the change in environment Δ_j^{HRR}). In both cases the answer appears to be negative: the values of β_t for $t < 0$ show no particular trend and an F-test that all eight estimates are jointly equal to zero fails to reject with $p = 0.34$.

The lack of observed selective migration greatly facilitates interpreting the changes in physician behavior across the move for two reasons. First, the jump in physician behavior

---

[8]Since physicians may practice at multiple hospitals within the origin HRR, an "origin" hospital is not well-defined for many physicians.

across a move is the causal effect of the experienced change in environment under a parallel trends assumption that differences in migrant behavior would have remained unchanged absent the treatment. The lack of any trend in this difference in the years leading up to the move strongly boosts the plausibility of this assumption.

The pre-move estimates also speak to the possibility that different types of migrants may sort differentially to higher or lower intensity regions. This could raise the concern that, for a given treatment, the effect on a migrant who chose that treatment may not be the same for a migrant who did not choose that treatment. If, however, physicians who started in the same region and later moved to dissimilar regions practiced no differently before the move, it would rule out any (perhaps unobserved) sources of selective migration that are correlated with observed physician practice choices. This is in fact what the results in Figure 3 suggest, given that $\beta_t \approx 0$ prior to the move.

The change in β_t at $t = 0$, corresponding to the first year after a physician's move, rejects the null hypothesis that the environment has no effect on physician behavior (e.g. that physicians are "stuck" in their ways) and shows that there is a significant and immediate positive physician response to the new environment. The finding that there is no further physician response to the environment—β_t is flat for $t \geq 0$—suggests that the nature of the physician response is not about slow moving factors, such as skill development or learning. This stands in contrast, for example, to the hypothesis that physician styles evolve according to a Bayesian-learning process of adaptation (see e.g. Phelps and Mooney, 1993).

As discussed above, the measured change in behavior across the move could partly reflect a compositional shift in the fraction of patients treated by different physicians. However, this does not appear to be a major issue here: estimates controlling for physician fixed effects (dashed gray line, Figure 3) are very similar to the results that control only for origin HRR.

Finally, the physician response to a change in the environment is bounded away from unity, suggesting that physician behavior is not fully determined by the environment. If HRRs appropriately characterize each physician's practice environment, $\beta_t \approx 0.66$ (2/3) for

17

$t \geq 0$ implies that the environment matters about twice as much as the physician. However, if HRRs mismeasure a physician's relevant practice region, the estimated environment impact may only provide a lower bound. In Section 3.2 I explore this possibility in more detail using a cross-sectional approach to estimate a lower bound on the physician-specific effect.

An alternative and more direct approach to alleviating concerns that HRRs miscapture the relevant environment is to measure changes in environment based on differences Δ_j^{hosp} in a physician's hospital cath environment across the move. Figure 3 Panel A plots the estimates of β_t from Equation 1 based on differences Δ_j^{hosp} and controlling for physician fixed effects. The sequence of β_t estimates shows no apparent differential trend in behavior before the move ($t < 0$), with a discrete jump at $t = 0$ and remaining roughly flat thereafter near 0.78. The hospital environment results are qualitatively similar but notably more precise and slightly larger than those in Figure 3 based on changes in the more aggregate HRR environment measure.

The increase in precision when defining changes in the environment using hospital-level cath rates is important for detecting heterogeneity in physician response. For example, Panel B of Figure 3 plots the results of augmenting the regression behind Panel A to allow for separate effects β_t^{up} and β_t^{down} for physicians moving "up" to more-intensive hospitals ($\Delta_j^{hosp} > 0$) versus those moving "down" to less-intensive hospitals. While neither group of physicians shows differential trends in behavior prior to the move, physicians moving to less-intensive hospitals appear to retain more of their previous practice style (i.e. change behavior less) than do physicians moving to more-intensive hospitals. I explore the implications of this result in greater detail below, where I also show that estimating asymmetric responses based on HRR-level changes in cath rates yields a similar point estimate that is far less-precisely measured (see Table 6).

3.1.3 Difference-in-differences results

The event study results documented an absence of differential pretrends in physician behavior prior to a move, followed by a discrete change in the year immediately following

18

the move with little additional change over time. As described above, I summarize the event study estimates using a traditional style difference-in-differences estimate of the physician response to a change in the environment. The key advantage of the DD estimate is that it provides a single summary estimate of the physician response, which effectively increases the statistical precision of that parameter and makes it easier to compare the sensitivity of the results to different specifications.

Table 4 presents the difference-in-difference results. Each column reports the DD estimate from a separate regression. Columns (1-5) and (9-10) report estimates based on the change Δ_j^{HRR} in HRR environment, while columns (6-8) report estimates using the change Δ_j^{hosp} in hospital environment.

Column (1) reports the DD estimates of physician behavior with respect to a change Δ_j^{HRR} in HRR environment, relative to other migrants from the same origin HRR. The coefficient on Δ_j^{HRR} describes the degree of selective migration. If physicians who move to regions that are more intensive than their origin region practice more intensively than their peers prior to the move, the coefficient on Δ_j^{HRR} will be positive. In fact, the estimated coefficient on Δ_j^{HRR} is close to zero, consistent with the limited degree of selective migration based on observed treatment choices documented in the event study (Figure 3).

The DD environment effect, which measures the degree to which cardiologists alter their treatment decisions after a move in response to a change in regional norms, is captured by the coefficient on the interaction of Δ_j with an "after" dummy $\mathbf{1}(t \geq 0)$. The DD estimate of this effect is 0.63, implying that a 1 percentage point change in a physician's HRR cath environment corresponds to a 0.63 percentage point change in that physician's measured cath behavior. This estimate rejects both the polar view that physician practice styles are fully ingrained and do not respond to changes in the environment, and also the other extreme that physicians change their behavior 1-for-1 in response to a change in environment.

To account for potential compositional shifts in the fraction of patients treated by different physicians across a move, the regression result reported in Column (2) includes physician

fixed effects. The resulting DD estimate of 0.65 is very similar to the estimate in Column (1) that controls only for origin HRR. Modifying the specification from Column (2) using changes $\Delta_j^{hospital}$ in the hospital environment across the move yields a DD estimate of 0.80. Like the results based on changes in the HRR cath intensity, the estimate based on changes in hospital intensity across a move suggest that physician practice styles are highly responsive to changes in their environment.

3.1.4   Difference-in-differences robustness

While the DD estimates above are consistent with the interpretation that the environment has a large impact on the treatment of a physician's patients, it is possible that this effect is driven in part or whole by cases where multiple specialists treat a heart attack patient, each independently making decisions whether to refer the patient for an early cath. In this case, even if a migrant cardiologist never changes his or her propensity to recommend an early cath after a move, their patients would nevertheless receive treatment more in line with the destination region. This would confound the interpretation of results as a change in a physician's practice style.

To investigate whether the multiple specialist hypothesis appears to be a key driver of the observed change in the treatment of a physician's patients across a move, I first estimate the DD effect over the 66% of patients who are treated by only a single cardiologist on the first day. If the first cardiologist is not significantly changing cath recommendations following a move, the DD estimate over this set of patients should be lower. In fact, as reported in column (3) of Table 4, the estimated DD effect of 0.71 is actually slightly larger than in the baseline specification.

A primary concern with limiting the sample of patients to those who see a single cardiologist is that the number of cardiologists seen on the first day may depend on the treatment choices of the first cardiologist seen. Following a strategy closely related to that of Doyle (2016), I therefore consider an "only-specialist-there" check which limits the sample to patients admitted on days when there is only one cardiologist admitting patients at that

20

hospital.[9] Because I am measuring two-day cath rates, I further include patients for whom only one cardiologist admits any patients at that hospital within two days of the patient's admission. The DD effects estimated over this sample, reported in column (4) for changes in HRR environment and in column (7) for changes in hospital environment, change little relative to the baseline HRR and hospital estimates in columns (2) and (6), respectively. This reinforces the interpretation of changes in a cardiologist's patient treatment across a move as a change in the physician's practice style.

Another potential concern is that the DD estimates may simply reflect the fact that some patients are admitted to hospitals that do not have catheterization facilities. If a cardiologist moves from a region with a high share of hospitals with cath labs to a region with a lower share, the change in the cardiologist's cath behavior may simply reflect that it is hard to do caths when the initial hospital has no cath lab. The lack of cath facilities seems unlikely to play a major role over the period of analysis in this paper. 84% of heart attack patients in 1998 were admitted to a hospital with cath facilities, and this share only grew over time.

As a direct test, I estimate the DD regression limited to patients admitted to hospitals with cath lab facilities. To do this, I first measure whether a hospital has a cath lab facility in a given year based on whether at least two Medicare patients admitted to the hospital that year (for any condition) received a cardiac catheterization. I also define modified versions of Δ_j^{HRR} and Δ_j^{hosp} based only on patients admitted to hospitals with cath facilities. The DD effects estimated over the cath-lab sample, reported in column (5) for changes in HRR environment and in column (8) for changes in hospital environment, decrease only slightly relative to the baseline HRR and hospital estimates in columns (2) and (6), respectively. This suggests that lacking a cath lab is not the key factor driving changes in physician treatment choices across a move, at least during this period.

---

[9]The set of admitting physicians at a hospital on a given date is based on the date of admission for patients with any diagnosis, not just AMI patients, and the associated attending physicians. Hospital admissions are observed for a 100% sample of Medicare fee-for-service beneficiaries using the MedPAR data. While these data do not report the admitting physician, prior literature has documented that the admitting physician is also usually the attending physician for heart attack patients (Jollis et al., 1996).

As a final robustness check, I re-estimate the HRR environment specification using both migrant and non-migrant physician data. This has two primary benefits. First, it allows HRR, calendar year, and comorbidity fixed effects to be estimated using all data, providing more efficient estimates under the condition where these effects are the same for both migrants and non-migrants. The second benefit is that it allows a direct comparison of migrant and non-migrant behavior. In particular, in a specification without physician fixed effects, I will estimate whether migrants are more or less intensive on average than non-migrants.

To estimate this regression, I define the origin and destination HRRs for non-migrants to be equal to the current HRR in which they are observed practicing and also set their event time at $t = -1$ in all periods. Thus, by definition, $\Delta_j^{HRR} = 0$ for all non-migrants.

Results estimated over the full sample of cardiologists are reported in the last two columns of Table 4. In column (9), which includes both origin and destination HRR fixed effects, the DD estimate is very similar to the same specification estimated over movers only (column 2). As shown by the coefficient on the "mover" dummy, migrant behavior is very similar to non-migrants on the whole. This suggests that selection into the migrant sample is not driven by factors that are correlated with a physician's practice intensity level. Column (10) reports results controlling for physician fixed effects; the resulting estimate of 0.65 is essentially identical to estimating the same regression over movers only (column 2).

3.2 Cross-section

3.2.1 Measurement error and estimate bounds

While the difference-in-differences results (Section 3.1.3) suggest that physicians are highly responsive to changes in their environment, the results also show less than 1-to-1 conformity to these changes, rejecting the null hypothesis that physician behavior is fully characterized by the physician's current practice environment. A valid concern with this conclusion, however, is whether mis-measurement exists in the key independent variable—namely, the measured change in a physician's practice environment across a move—and if so

whether this biases the estimated physician response. Importantly, if the estimated response is biased toward zero, then we must interpret the estimate as a lower bound of the true physician response to a change in environment and we can no longer reject the possibility that physicians fully converge to the new environment. Thus, the goal of this section is to lay out the conditions under which measurement error does or does not bias the estimated physician response, and to provide a framework that provides both upper and lower bounds of the estimates in the presence of pernicious measurement error.

Throughout the analysis, Δ_j has denoted the change in environment a physician experiences across a move. Suppose that instead of observing Δ_j directly, we can only measure $\Delta_{g(j)}$ where $g(.)$ is the potentially limited set of information available to the econometrician about a physician's change in environment. For example, $g(.)$ may include only the origin and destination HRRs corresponding to a physician's move. Then linear estimates of physician response to Δ_j as in (1) can be consistently estimated by replacing Δ_j with $\Delta_{g(j)}$ if and only if the consistency condition $\Delta_{g(j)} = E[\Delta_j|g]$ is satisfied.[10]

In other words, the consistency condition states that the measured change in environment Δ_g need not equal the physician's actual change in environment Δ_j, but it must equal the conditional expected value of Δ_j. An important case satisfying this condition occurs if the scope of geography relevant for measuring a physician's practice environment is smaller than an HRR (e.g. a ZIP code or specific hospital). Then as long as migrant physicians do not positively select into sub-environments conditional on the choice of HRR (which may not be an unreasonable assumption given the previously observed limited scope of selection into HRRs), the expected change in the physician's own environment across a move is just the average change in environment across the origin and destination HRRs. For this reason, misclassifying geographic practice regions too broadly need not bias the resulting estimates.

There are, however, at least two potential sources of measurement error that could arise

---

[10]To see this, suppose the conditional expectation function (CEF) of y given x is given by $E[y|x] = x\beta$. Then if g contains less-specific information than x, the law of iterated expectations implies $E[y|g] = E[x|g]\beta$. Thus, OLS using $E[x|g]$ in place of x produces consistent estimates of β. This result is analogous to the well-known result that linear CEFs can be estimated by OLS over grouped means.

and violate the consistency condition. The first is that even if physicians do not positively select into more- or less-intensive HRRs as suggested by the main results, there still may be some scope for them to positively select into sub-environments conditional on HRR. This could occur, for example, if physicians choose a destination HRR based on reasons unrelated to practice style, but then sort into a specific practice location within the HRR, such as the city or hospital that most closely resembles their previous practice setting. In this case, the measured change in environment based on HRR is systematically larger than the physician's expected change ($\Delta_g > E[\Delta_j|g]$), and estimates of physician response based on Δ_g would be biased toward zero. A second potential source of mismeasurement is classical measurement error, which through attenuation bias would also lead to under-estimates of the true physician response to a change in environment. This type of measurement error might arise, for example, if the HRR practice environments vary over time.[11]

Because these types of error in measured environment changes raise the possibility that estimates from the baseline analysis should be interpreted as a lower bound on the true physician response to changes in environment, I adopt an alternative approach that instead provides an upper bound on the physician response. The key idea underlying this approach is rather than estimate the degree β to which physicians conform to a (potentially mismeasured) change in environment as in the event study and DD specifications, instead I estimate the degree θ to which physicians fail to conform to their new environment, and thus these two parameters are related by $\beta = 1 - \theta$. Importantly, obtaining a lower bound on θ provides an upper bound on β.

To measure the degree θ of non-conformance, I adopt a "cross-sectional" approach and estimate non-conformance as the degree to which migrant physicians retain their original practice style after a move, relative to other physicians practicing in their post-move environment. When practice environments are measured without error and physicians do not

---

[11]While we can perform the analysis using time-varying rates directly as in Appendix B.1, this approach trades one type of classical measurement error for another as time-varying rates are effectively measured over smaller samples, potentially introducing statistical noise into the environment measure.

positively sort into practice environments, lack of conformance θ is given by

$$(cath)_{ij} = \theta\big[(\textit{origin HRR cath intensity})_j - (\textit{current HRR cath intensity})_j\big]$$

$$+ \{\textit{environment FEs}\}_i + \{\textit{year FEs}\}_i + \epsilon_{ijt}. \tag{2}$$

To estimate this regression, I include patient-physician observations for both post-move migrants and non-migrants, where the origin HRR is defined to be the current HRR for non-migrants. The coefficient θ describes the degree to which physicians in the same practice environment (as specified by the choice of environment dummies) differ in their behavior per unit difference in their previous practice environment. Importantly, the same potential sources of downward bias in the estimation of β will also result in downward bias in the estimation of θ. If there exists positive matching of physicians to practice regions, the estimate of θ will be a lower bound on the true value of the parameter; similarly, any classical measurement error results in attenuation bias of the estimate. Thus, estimating θ from equation (2) provides a lower bound on the degree of non-conformance of physician behavior to changes in practice environment, which as previously noted gives us an upper bound on the degree of conformance by subtracting this estimate from 1.

The results from this regression are shown in Panel A of Table 5. Column (1) shows results when then environment is defined as the HRR, while column (2) defines the environment as the hospital. Standard errors are calculated using two-way clusters in the cardiologist's origin and current HRRs. In both columns (1-2), the point estimates on $(\textit{origin HRR cath intensity})_j$ are significantly larger than zero, allowing us to reject at the 0.01 level the null hypothesis that a physician behavior does not vary based on the intensity of the physician's previous practice environment. This in turn rejects the hypothesis that physician practice styles are fully characterized by the current practice environment.

Finally, I convert this estimate into an upper bound $\bar{\beta}$ on the degree of physician response to a change in practice environment. To be conservative, I take the estimate of $\hat{\theta} = 0.17$ from

the hospital-level specification in column (2) of Table 5, from which I calculate $\bar{\beta} = 1 - 0.17 = 0.83$. This result is can be compared to the estimates of $\hat{\beta} \approx 0.66$ implied by the event study specification shown in Figure 3 which as previously discussed can be interpreted as a lower bound on the true value of β. These bounds of 0.66 and 0.83 also include the event study estimates $\hat{\beta} \approx 0.78$ based on hospital-level changes in cath intensity shown in Figure 4. That the upper and lower bounds are relatively similar suggests that the difference-in-differences results are quite reliable as an estimate of the degree to which physicians conform to changes in their practice environment.

3.2.2   Doctor-patient sorting

The cross-section results just described show that physicians systematically differ (even within the same hospital) in their treatment decisions based on prior experience. However, a potential concern is that these results are driven by patients being sorted to doctors based on (potentially unobservable) clinical appropriateness in a way that is correlated with a physician's background. In this section, I test the plausibility of this concern by evaluating whether physicians from different backgrounds see patients with *observably* different levels of clinical appropriateness for intensive heart attack management.

To do this, I first construct an index of patient clinical appropriateness for intensive management. Similar to Chandra and Staiger (2007), I define clinical appropriateness using logistic regression of patient catheterization within 2 days of a heart attack, as

$$Pr\left(cath_{iht}\right) = G\left[\theta_h + \theta_h t + X_{it}\Phi\right]. \tag{3}$$

Here, θ_h is an indicator for the HRR h in which patient i was treated. This indicator enters directly and also interacted with continuous calendar year t to allow for arbitrary linear trends by HRR. X_{it} includes calendar year dummies, patient comorbidities, and comorbidities interacted with calendar year. The empirical index of patient clinical appropriateness is obtained as the fitted values from (3) evaluated at a baseline year and HRR.

Given the index of patient clinical appropriateness, I test whether two physicians with different backgrounds but currently practicing in the same hospital systematically see patients with different levels of appropriateness. Specifically, I estimate the same regressions reported in Table 5, except to replace the original dependent variable $(cath)_{ij}$ with $\hat{Pr}(cath)_{ij}$.

The results are reported in Panel B of Table 5. When using HRRs to define the current environment, the resulting estimate of θ (column 1) is small but statistically significant at the 5% level, indicating some scope for selective matching of patient and physician types within HRRs. Specifically, the point estimate of $\hat{\theta} = 0.025$ implies that for a 10 percentage point difference in the cath backgrounds of two physicians now practicing in the same HRR, the physician with the more intensive background sees patients whose empirical propensity to be cathed is 0.25 percentage points higher, on average. Controlling for the hospital rather than the HRR as the current environment reduces the estimate of θ to 0.018 (column 2), which is not statistically different from zero. The reduction in θ when controlling for the hospital environment suggests that some sorting within an HRR is due to differences in patient and physician characteristics across hospitals within the HRR.

As a final specification, column (3) repeats the specification from column (2), but limited to the same "only-specialist-there" check described in Section 3.1.4. This check limits the sample to patients admitted when there is only one cardiologist admitting patients at that hospital. Plausibly, any within-hospital sorting of patients to cardiologists is likely to be minimized on such days. In fact, no sorting is detected: $\hat{\theta} = -0.0002$.

In combination, the limited or no sorting results in columns (2-3) support two conclusions. First, cardiologists currently practicing in the same hospital do differ in their treatment choices based on their prior environment, and this difference does not appear to be driven by patient sorting within the hospital. Second, sorting of AMI patients is plausibly negligible on days when only one cardiologist is admitting patients at the hospital. This latter conclusion reinforces the usefulness of the "only-specialist-there" robustness check of the DD estimates discussed and reported in Section 3.1.4.

3.3 Characterizing Physician Behavior Changes

In this section, I aim to characterize in more detail the nature of physician response to changes in the physician's environment. This is useful for evaluating prevailing theories of how physician styles are formed by testing the distinct implications these theories have for how physician behavior should respond to changes in the environment.

3.3.1 "Schools of Thought" Theory

A predominant theory used to explain the existence and persistence of regional practice variations is the Phelps and Mooney (1993) model of information diffusion and physician learning. In this model, uncertainty and complexity regarding the efficacy of various medical interventions ultimately lead to regional "schools of thought" concerning what constitutes best practice. Physicians form initial practice styles based on where they train in medical school. Over time, these practice styles evolve according to a Bayesian learning process, as physicians update their beliefs based on local community norms. In this model, variations in health care delivery arise from incomplete information. Deviations from the fully informed provision of care either through over- or under-provision result in welfare losses. Determining whether this learning model appears to explain variations in AMI care is thus important for whether we should rely on its welfare implications.

An obvious implication of the learning model is that physicians will change behavior following a move across environments (as I find), but the model has two further implications which I am able to test directly in my empirical context. First, migrant behavior should evolve smoothly over time, eventually converging to the new school of thought regardless of where they came from. However, the patterns of behavior following a physician's move as measured by the event study in Figure 3 show that physicians partially conform their behavior to a new environment nearly immediately. This very rapid change in behavior across the move, with no further convergence even after 8 years, together are difficult to explain in a learning context.

A second implication of this learning model is that physicians who move later in their career should change their behavior less than physicians who move early in their career. I test this by testing for heterogeneity in the DD estimator from Section 3.1.1 based on whether the physician was more than 8 years post-fellowship completion (the median among migrants) at the time of the move. The specific regression I estimate takes the form

$$(cath)_{ijt} = \{physician\ FEs\}_j + \beta\{\Delta_j \times \mathbf{1}(t \geq 0) \times \mathbf{1}(tsf_j > 8)\}$$
$$+ \{main\ effects\ and\ two\text{-}way\ interactions\ of\ \Delta_j, \mathbf{1}(t \geq 0), \mathbf{1}(tsf_j > 8)\}_j$$
$$+ \{calendar\ year\ FEs\}_i + \{patient\ risk\text{-}adjusters\}_i + \epsilon_{ijt}. \tag{4}$$

The estimates of β when Δ_j is based on HRR and hospital changes in cath environment across a move are reported in Table 6 columns (1) and (4), respectively. In both cases, $\hat{\beta}$, is negative but small, and not significant at the 5% level. The small point estimates imply that physicians who move later in their career respond about the same to changes in their environment as those who move early in their career, suggesting physician practice styles remain elastic over time.

3.3.2 Asymmetries

Depending on the primary environment-level mechanisms that drive physicians to change behavior across a move, physician responses to an increase in the environment's intensity may differ from the response to a decrease in intensity. For example, if hard capacity constraints such as lacking catheterization facilities are a predominant factor driving physician cath decisions, a move to a region that is less intensive because of restricted access to hospital capacity could plausibly have a larger impact on the physician's practice style than a move to a region that is more intensive because of expanded capacity. On the other hand, defensive medicine in the face of medical malpractice risk and locality rules may lead physicians to respond more to increases in the local diagnostic cath environment than to decreases.

Panel B of Figure 4 plots the results of the event study of physician behavior across a

29

move with respect to a change Δ_j^{hosp} in hospital cath intensity, with separate effects β_t^{up} and β_t^{down} for physicians moving "up" to more-intensive hospitals ($\Delta_j^{hosp} > 0$) versus those moving "down" to less-intensive hospitals. Averaging the β_t^{up} and β_t^{down} parameters separately for $t \geq 0$ implies that a physician moving to a 1pp more-intensive hospital changes behavior by about 0.21pp more than a physician moving to a 1pp less-intensive hospital.

To directly summarize this result for changes in HRR or hospital cath intensity, I estimate Equation 4 above, replacing $\mathbf{1}(tsf_j > 8)$ by an indicator for whether the physician moved to a more-intensive region ($\Delta_j > 0$). The results from HRR and hospital changes in cath environment across a move are reported in Table 6 columns (2) and (5), respectively. The point estimate based on changes in HRR cath rates is 0.20 (though not statistically significant), and the estimate based on hospital cath rates is 0.27 (significant at the 1% level). Both these estimates are very similar to the 0.22 estimate implied by the event study in Figure 3 and imply that physicians respond more to increases in environment intensity.

3.3.3 Long moves

A primary aim of this paper is to shed light on the role environmental factors play in shaping a physician's practice choices. Physician moves can be used estimate the role environment factors play on shaping migrant physician behavior, but an important consideration is whether these results are externally valid for the non-mover population. The previous analysis showed that migrants and non-migrants look similar in terms of pre-move behavior (as captured by approximately zero coefficient on the "mover" dummy in column 10 of Table 4). However, if changing one's practice style requires costly adjustments, it is possible that movers are selectively more elastic to changes in their environment.

Since by definition it is not possible to examine the change in behavior for non-migrants across a move, I instead consider whether physicians who move a long distance demonstrate a different response to changes in the environment compared with physicians moving a shorter distance. I do this by estimating Equation 4 above, replacing $\mathbf{1}(tsf_j > 8)$ by an indicator for whether the physician was among the 45% of migrants who moved across Census Regions

(Table 2). As reported in Table 6, columns (3) and (6), physicians who move across Census Regions respond about the same to changes in their environment as those who do not.

4   Mechanisms

The results in Section 3 provide evidence that the environment plays an important role in how physicians treat patients. There are a variety of ways in which a physician's environment could influence treatment decisions. Here, I briefly discuss three possible mechanisms that are likely to be important in the current context, including how the results in this paper shed light on how these factors are likely to play a role.

Because invasive treatment of a heart attack requires the use of a specialized laboratory setting and access to other hospital resources, the availability of these capital resources may be an important driver of how cardiologists treat heart attack patients. Evidence from Gatsonis et al. (1995) finds that states with more extensive on-site availability of cardiac catheterization have higher catheterization rates after adjusting for patient characteristics. However, while hard capacity constraints such as lack of catheterization facilities are likely to greatly limit early cath rates for heart attack patients admitted to those hospitals, this does not appear to be a primary mechanism driving the changes in physician behavior I observe in this setting where over 89% of heart attack patients in the sample are admitted to hospitals with cath labs and the estimates of physician response across a move change little when limiting the analysis to this subsample. An interesting question that merits further exploration is the role played by softer "intensive" capacity constraints, such as whether a hospital's cath lab is staffed nights or weekends. Since a two-day window for cath always overlaps with regular business hours on at least one weekday, these types of capacity constraints are unlikely to prohibit early caths, but may change the relative benefits of medical management if they increase the time to cath.

Chandra and Staiger (2007) find evidence that the environment may also influence physician decisions through productivity spillovers. These spillovers could occur at the regional

level, such as from knowledge spillovers across physicians practicing in the same region or by attracting physicians who have specialized in certain types of treatments. Spillovers could also occur at the physician level, through learning-by-doing and skill specialization (perhaps as a function of the underlying patient population). The current paper does not rule out physician-level spillover effects—physicians change their behavior less than one-for-one in response to a change in the environment, which could be a result of embedded habits or skills. However, the results in this paper do speak to how regional-level spillovers are likely to occur. First, given the very limited degree of physician sorting that I find, it appears that the attraction of specialized physicians to particular regions (at least at the level of the HRR) is quite small in this context.[12] To the degree that knowledge spillovers occur across physicians, these appear to occur in a manner that changes physician behavior immediately after a move with no further effect. Finally, it is worth noting that capacity constraints can be related to regional productivity spillovers: hospitals or regions may induce physician specialization by accumulating a stock of capital that targets a particular treatment.

Finally, I consider how the environment may influence physician behavior through "team" effects. Due to the emergency nature of heart attacks and the time-sensitivity of the relative benefits of different treatment paths, optimal patient treatment may depend on the speed and accuracy of preliminary diagnoses by the triage and emergency room staff. Moreover, the probability of complications as a result of invasive treatment and the ability to identify and cope with such complications could depend on the skill of cardiac catheterization lab technicians, hospital nurses, and surgical staff. Thus, the factors specific to the team of physicians and hospital staff involved in the care of the heart attack patient may also play an important role in a cardiologist's treatment behavior. Moreover, due to their discrete nature across practice settings, team factors are also consistent with the observed level-shift in physician behavior across a move.

---

[12]This limited selective migration applies to physicians moving later in their career. It is unknown whether the degree of selective migration differs for cardiologists moving directly out of their cardiology fellowship.

5 Conclusion

Cardiologists vary widely across U.S. regions in their propensity to intensively manage heart attacks, even after adjusting for apparent differences in average patient characteristics and illness severity across regions. Such variation could result from differences in local practice environments, such as access to hospital capacity, the availability of specialists, and medical malpractice exposure. On the other hand, the regional differences could be driven entirely by physician-specific factors such as training, preferences, and experience as a result of positive matching of physicians to other physicians with similar practice styles.

This paper attempts to identify the role of the environment on a cardiologist's behavior relative to physician-specific factors by exploiting changes in practice environment resulting from cardiologist migration. Using 15 years of Medicare data, I trace migrant treatment choices in a long panel before and after a move. Positive sorting is identified by the degree to which physicians starting in the same region and later moving to dissimilar regions already practiced dissimilarly before the move. The environment effect is identified by the change in physician behavior across the move. I find that both environment and physician-specific factors impact practice style, but the role of the environment is at least twice as large. Also, the pattern of physician behavior changes observed across a move is not consistent with the "schools of thought" model often used to describe regional differences in medical practice.

The results in this paper capture how individual physicians adapt to a new environment following a move. This environment consists of all things not embedded in that physician, including physical hospital capacity and systems processes, as well as the human capital of other (possibly non-randomly selected) medical providers also practicing in that environment. I find that physician practice styles even mid-career are highly "elastic" with respect to environment changes. Understanding the components primarily responsible for this environment effect is essential for policy makers challenged with changing provider behavior. While the analyses in the paper shed new light on the roles of certain environment-level factors, there is significant scope for more work to be done in this important area.

References

Angrist, Joshua D. The perils of peer effects. *Labour Economics*, 30:98–108, 2014.

Baicker, Katherine and Chandra, Amitabh. Medicare spending, the physician workforce, and beneficiaries' quality of care. *Health Affairs*, W4:184–197, 2004. URL http://content. healthaffairs.org/content/early/2004/04/07/hlthaff.w4.184.citation.

Barnato, Amber E.; Herndon, M. Brooke; Anthony, Denise L.; Gallagher, Patricia M.; Skinner, Jonathan S.; Bynum, Julie P. W., and Fisher, Elliott S. Are regional variations in end-of-life care intensity explained by patient preferences?: A study of the us medicare population. *Medical care*, 45(5):386, 2007.

Bronnenberg, Bart J; Dubé, Jean-Pierre H, and Gentzkow, Matthew. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102 (6):2472–2508, 2012.

Brophy, James M. and Bogaty, Peter. Primary angioplasty and thrombolysis are both reasonable options in acute myocardial infarction. *Annals of Internal Medicine*, 141(4):292–297, 2004.

Chandra, Amitabh and Staiger, Douglas O. Productivity spillovers in health care: Evidence from the treatment of heart attacks. *The Journal of Political Economy*, 115(1):103–140, 2007.

Chetty, Raj; Friedman, John, and Saez, Emmanuel. Neighborhood and peer effects in tax policy: Evidence from the eitc and saver's credit. Seminar, National Bureau of Economic Research Summer Institute 2011, August 2011.

Congressional Budget Office, . Geographic variation in health care spending. February 2008. URL http://www.cbo.gov/ftpdocs/89xx/doc8972/02-15-GeogHealth.pdf.

Doyle, Joseph J. Physician characteristics and patient survival: Evidence from physician availability. Working Paper, 2016.

Doyle, Joseph J.; Ewer, Steven M., and Wagner, Todd H. Returns to physician human capital: Evidence from patients randomized to physician teams. *Journal of Health Economics*, 2010.

Dranove, David; Ramanarayanan, Subramaniam, and Sfekas, Andrew. Does the market punish aggressive experts? evidence from cesarean sections. *B.E. Journal of Economic Analysis & Policy*, 11(2), 2011.

Elixhauser, Anne; Steiner, Claudia; Harris, D. Robert, and Coffey, Rosanna M. Comorbidity measures for use with administrative data. *Medical Care*, 36(1):8–27, 1998.

Epstein, Andrew J. and Nicholson, Sean. The formation and evolution of physician treatment styles: An application to cesarean sections. *Journal of Health Economics*, 28(6):1126–1140, 2009.

Fernández, Raquel. Culture and economics. In Durlauf, Steven N. and Blume, Lawrence E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke, 2008.

Finkelstein, Amy; Gentzkow, Matthew, and Williams, Heidi. Sources of geographic variation in health care: Evidence from patient migration. Technical Report 20789, National Bureau of Economic Research, 2014.

Fisher, Elliott S.; Wennberg, David E.; Stukel, Thérèse A.; Gottlieb, Daniel J.; Lucas, F.L., and Étoile L. Pinder, . The implications of regional variations in medicare spending. part 1: the content, quality, and accessibility of care. *Annals of Internal Medicine*, 138(4): 273–287, 2003a.

Fisher, Elliott S.; Wennberg, David E.; Stukel, Thérèse A.; Gottlieb, Daniel J.; Lucas, F.L., and Étoile L. Pinder, . The implications of regional variations in medicare spending. part 2: health outcomes and satisfaction with care. *Annals of Internal Medicine*, 138(4): 288–298, 2003b.

Gatsonis, Constantine A.; Epstein, Arnold M.; Newhouse, Joseph P.; Normand, Sharon-Lise, and McNeil, Barbara J. Variations in the utilization of coronary angiography for elderly patients with an acute myocardial infarction: An analysis using hierarchical logistic regression. *Medical Care*, 33(6):625–642, 1995.

Gawande, Atul. The cost conundrum. *The New Yorker*, June 2009.

Glover, J. Alison. The incidence of tonsillectomy in school children. *Indian Journal of Pediatrics*, 5(4):252–258, 1938.

Gottlieb, Daniel J.; Zhou, Weiping; Song, Yunjie; Andrews, Kathryn Gilman; Skinner, Jonathan S., and Sutherland, Jason M. Prices don't drive regional medicare spending variations. *Health Affairs*, 29(3):537, 2010.

Grytten, J. and Sørensen, R. Practice variation and physician-specific effects. *Journal of Health Economics*, 22(3):403–418, 2003.

Ichino, Andrea and Maggi, Giovanni. Work environment and individual background: Explaining regional shirking differentials in a large italian firm*. *Quarterly Journal of Economics*, 115(3):1057–1090, 2000.

Jollis, James G.; DeLong, Elizabeth R.; Peterson, Eric D.; Muhlbaier, Lawrence H.; Fortin, Donald F.; Califf, Robert M., and Mark, Daniel B. Outcome of acute myocardial infarction according to the specialty of the admitting physician. *New England Journal of Medicine*, 335(25):1880–1887, 1996.

Keeley, Ellen C. and Grines, Cindy L. Primary percutaneous coronary intervention for every patient with st-segment elevation myocardial infarction: What stands in the way? *Annals of Internal Medicine*, 141(4):298–304, 2004.

Kushner, Frederick G.; Hand, Mary; Smith, Sidney C.; King, Spencer B.; Anderson, Jeffrey L.; Antman, Elliott M.; Bailey, Steven R.; Bates, Eric R.; Blankenship, James C.; Casey, Donald E.; Green, Lee A.; Hochman, Judith S.; Jacobs, Alice K.; Krumholz, Harlan M.; Morrison, Douglass A.; Ornato, Joseph P.; Pearle, David L.; Peterson, Eric D.; Sloan, Michael A.; Whitlow, Patrick L., and Williams, David O. 2009 focused updates: Acc/aha guidelines for the management of patients with st-elevation myocardial infarction (updating the 2004 guideline and 2007 focused update) and acc/aha/scai guidelines on percutaneous coronary intervention (updating the 2005 guideline and 2007 focused update). *Circulation*, 120(22):2271–2306, 2009.

Manski, Charles F. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531, 1993.

Martin, Anne B.; Whittle, Lekha; Heffler, Stephen; Barron, Mary Carol; Sisko, Andrea, and Washington, Benjamin. Health spending by state of residence, 1991-2004. *Health Affairs*, 26(6):w651, 2007.

McClellan, Mark and Newhouse, Joseph P. The marginal cost-effectiveness of medical technology: A panel instrumental-variables approach. *Journal of Econometrics*, 77(1):39–64, 1997.

Phelps, Charles E. Diffusion of information in medical care. *The Journal of Economic Perspectives*, 6(3):23–42, 1992.

Phelps, Charles E. Information diffusion and best practice adoption. *Handbook of Health Economics*, 1:223–264, 2000.

Phelps, Charles E. and Mooney, Cathleen. Geographic variation in health care: The role of private markets. *Competitive approaches to health care reform*, pages 139–175, 1993.

Philipson, Tomas J.; Seabury, Seth A.; Lockwood, Lee M.; Goldman, Dana P., and Lakdawalla, Darius N. Geographic variation in health care: The role of private markets. *Brookings Papers on Economic Activity*, pages 325–361, 2010.

Quan, Hude; Sundararajan, Vijaya; Halfon, Patricia; Fong, Andrew; Burnand, Bernard; Luthi, Jean-Christophe; Saunders, L. Duncan; Beck, Cynthia A.; Feasby, Thomas E., and Ghali, William A. Coding algorithms for defining comorbidities in icd-9-cm and icd-10 administrative data. *Medical Care*, 43(11):1130–1139, November 2005.

Scanlon, Patrick J.; Faxon, David P.; Audet, Anne-Marie; Carabello, Blase; Dehmer, Gregory J.; Eagle, Kim A.; Legako, Ronald D.; Leon, Donald F.; Murray, John A.; Nissen, Steven E.; Pepine, Carl J.; Watson, Rita M.; Ritchie, James L.; Gibbons, Raymond J.; Cheitlin, Melvin D.; Eagle, Kim A.; Gardner, Timothy J.; ArthurGarson, Jr; Richard O.Russell, Jr; Ryan, Thomas J., and Sidney C.Smith, Jr. Acc/aha guidelines for coronary angiography: A report of the american college of cardiology/american heart association task force on practice guidelines (committee on coronary angiography) developed in collaboration with the society for cardiac angiography and interventions. *J Am Coll Cardiol*, 33(6):1756–1824, 1999.

Sirovich, Brenda E.; Gottlieb, Daniel J.; Welch, H. Gilbert, and Fisher, Elliott S. Regional variations in health care intensity and physician perceptions of quality of care. *Annals of Internal medicine*, 144(9):641–649, 2006.

Skinner, Jonathan and Fisher, Elliot. Regional disparities in medicare expenditures: an opportunity for reform. *National Tax Journal*, 50(3):413–25, 1997.

Skinner, Jonathan S.; Staiger, Douglas O., and Fisher, Elliott S. Is technological change in medicine always worth it? the case of acute myocardial infarction. *Health Affairs*, 25(2): W34–W47, 2006.

Skinner, Jonathan S.; Gottlieb, Daniel J., and Carmichael, Donald. A new series of medicare expenditure measures by hospital referral region: 2003-2008. In Bronner, Kristen K., editor, *A Report of the Dartmouth Atlas Project*. The Dartmouth Institute for Health Policy and Clinical Practice, June 2011.

Song, Yunjie; Skinner, Jonathan; Bynum, Julie; Sutherland, Jason; Wennberg, John E., and Fisher, Elliott S. Regional variations in diagnostic practices. *NEJM*, 363(1):45–53, 2010.

Wennberg, John and Gittelsohn, Alan. Small area variations in health care delivery. *Science*, 182(4117):1102–1108, 1973.

Wennberg, John E.; Fisher, Elliott S., and Skinner, Jonathan S. Geography and the debate over medicare reform. *Health Affairs*, 2002.

Zuckerman, Stephen; Waidmann, Timothy; Berenson, Robert, and Jack Hadley, Ph.D. Clarifying sources of geographicdifferences in medicare spending. *New England Journal of Medicine*, 363(1):54–62, 2010.

Figure 1: Distribution of 2-day cath rates by HRR
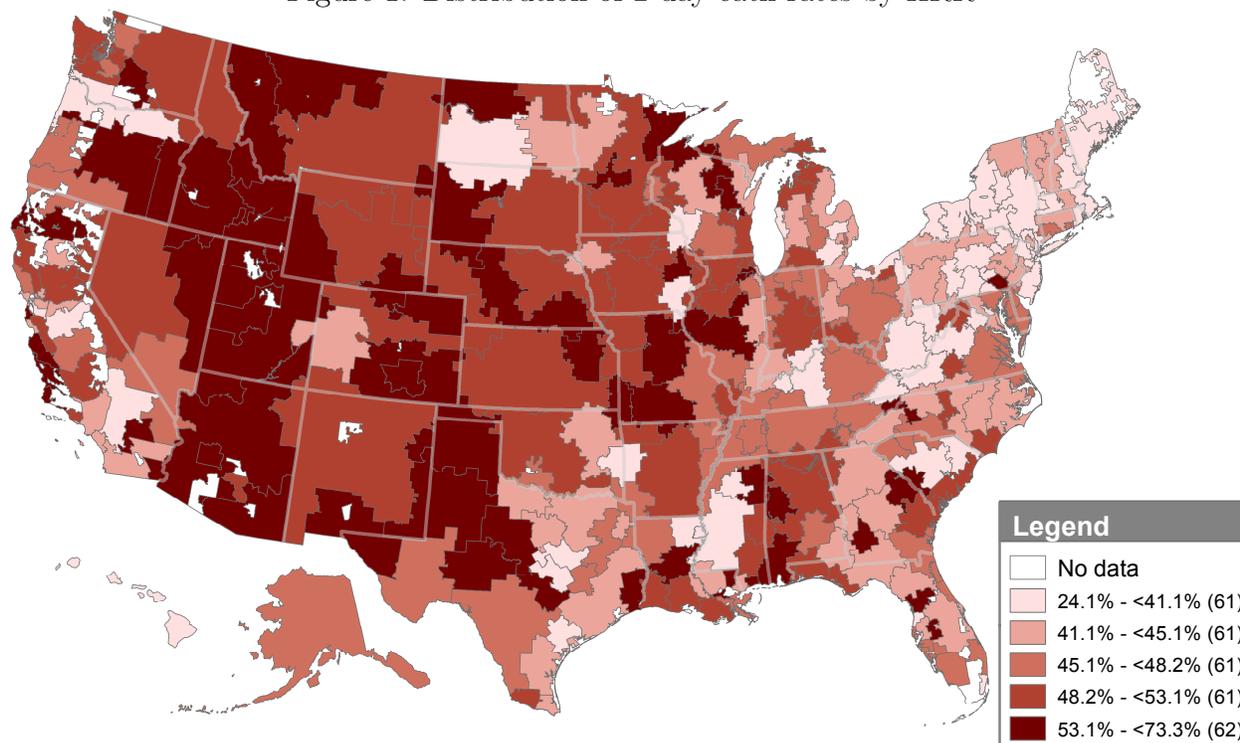
Notes: Map shows the geographic distribution of 2-day cardiac catheterization rates among Medicare heart attack (AMI) patients across the 306 Hospital Referral Regions (HRRs). Cath rates are calculated over pooled years 1998-2012 and weighted by the number of AMI patients treated in each region during this period. Rates are risk-adjusted for patient age, race, sex, and first heart attack.

Figure 2: Distribution of changes in 2-day cath environment across move
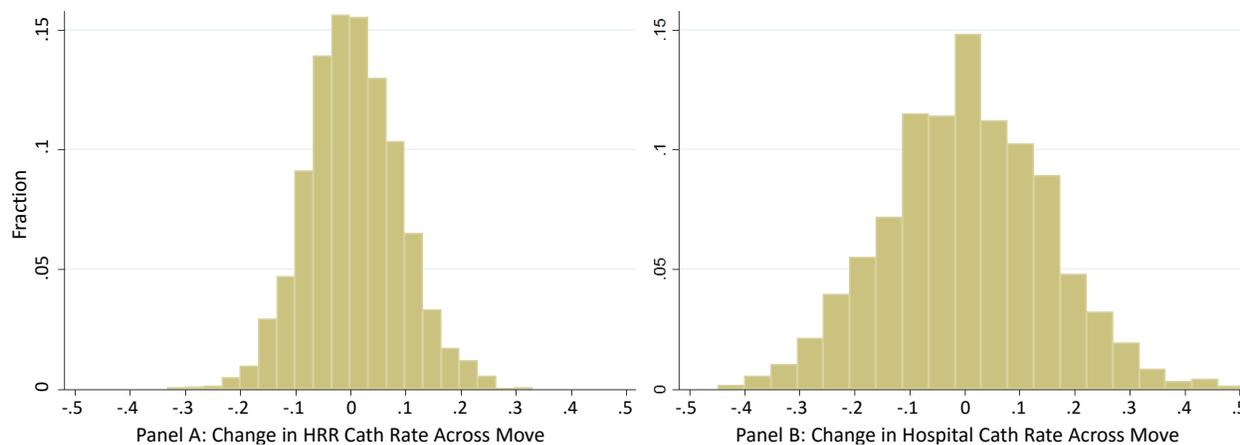
Notes: Figure shows the distribution of changes in cath environment among cardiologists who move HRRs. In Panel A, the change for cardiolgist j is defined as $\Delta_j = (destination\ HRR\ cath\ intensity)_j - (origin\ HRR\ cath\ intensity)_j$, defined as a physician-leave-out mean that omits physician j's own patients from the risk-adjusted HRR cath rates shown in Figure 1. Panel B shows the analogous distribution of changes in hospital cath intensity across the move based on the hospitals where the physician's patients were admitted. In both panels, the distribution is weighted by the number of sample heart attack patients treated by cardiologist movers.

Figure 3: Event study—change in HRR environment
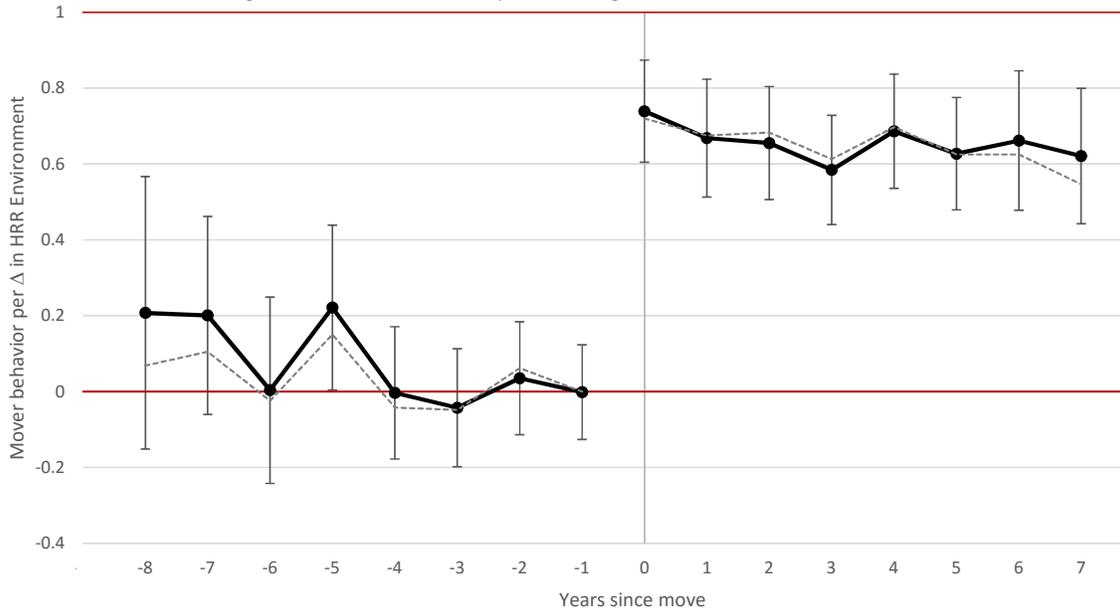
Notes: Graph plots (solid black) estimates of physician practice style t years since move as a function of the change in HRR cath environment experienced across the move (see Figure 2, Panel A). These estimates come from a regression that includes fixed effects for origin HRR, calendar year of patient admission, years since physician move, and patient age, race, sex, and first heart attack. Results controling for physician fixed effects instead of origin HRR are plotted by the dashed gray line. Bands indicate 95% confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels.

Figure 4: Event study—change in hospital environment
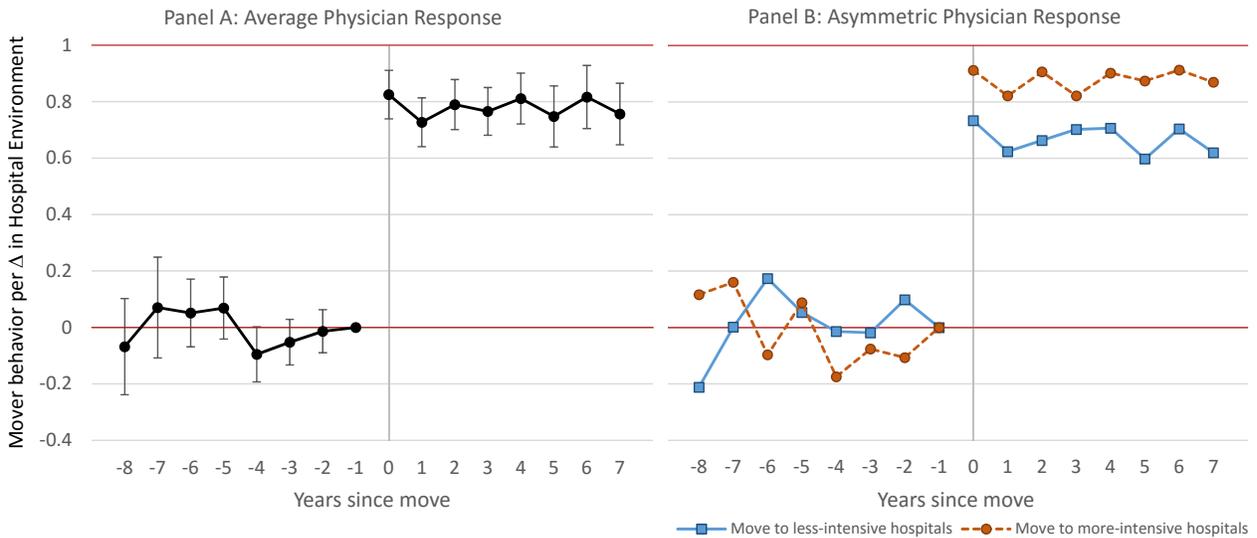
Notes: Graph plots estimates of physician practice style t years since move as a function of the change in hospital cath environment experienced across the move (see Figure 2, Panel B). Panels A and B plot estimates from separate regressions that include the same controls as the physician fixed effects regression behind Figure 3 (dashed gray line). Bands indicate 95% confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels. For Panel B, the regression allows for separate effects by whether the physician moved to more-intensive ($\Delta > 0$) or less-intensive ($\Delta \leq 0$) hospitals. Physician behavior is normalized to zero in the year immediately prior to the move ($t = -1$).

39

Table 1: Sample summary statistics

| | | HRR charactersitics | | | Patient characteristics | | | | | Cardiologist characteristics | |
| | | | 2-day cath rate | | | admitted to | | | | | |
year	N	p25	mean	p75	N	cath hospital	age	male	white	N	# movers
1998	306	0.275	0.340	0.398	43,929	0.840	75.8	0.523	0.895	11,617	55
1999	306	0.287	0.347	0.398	46,427	0.841	76.2	0.515	0.888	12,259	154
2000	306	0.308	0.364	0.417	48,730	0.841	76.3	0.516	0.891	12,750	221
2001	306	0.334	0.389	0.443	50,260	0.852	76.3	0.515	0.885	13,098	239
2002	306	0.364	0.417	0.472	51,705	0.867	76.2	0.518	0.883	13,694	261
2003	306	0.391	0.440	0.493	52,689	0.874	76.4	0.515	0.880	14,033	262
2004	306	0.413	0.467	0.519	50,870	0.889	76.3	0.520	0.879	14,337	290
2005	306	0.446	0.488	0.542	48,226	0.897	76.3	0.520	0.874	14,456	264
2006	306	0.463	0.510	0.563	44,712	0.908	76.3	0.519	0.875	14,596	323
2007	306	0.463	0.510	0.558	42,902	0.915	76.4	0.522	0.875	14,405	287
2008	306	0.472	0.510	0.557	41,405	0.916	76.4	0.522	0.873	13,809	201
2009	306	0.485	0.535	0.581	38,799	0.926	76.0	0.530	0.868	13,112	164
2010	306	0.506	0.552	0.595	38,198	0.927	76.1	0.526	0.863	12,560	164
2011	306	0.521	0.572	0.619	36,481	0.934	75.8	0.538	0.855	11,895	128
2012	306	0.533	0.576	0.619	34,064	0.935	75.9	0.532	0.856	11,197	76
1998-2012	306	0.431	0.479	0.530	669,397	0.888	76.2	0.521	0.877	19,945	3,089

Notes: The first set of columns describes the distribution of cath rates across the 306 HRRs for each year separately, as well as for the pooled sample 1998-2012 (final row). These rates are based on patients treated by the non-mover cardiologist sample and are risk-adjusted for patient age, race, sex, and first heart attack. The second set of columns describe the characteristics of patients with new heart attack episodes who form the primary analysis sample used to measure cardiologist practice styles. Reported characteristics include whether the patient was first admitted to a hospital with a cardiac cath lab in operation that year, average age, and fraction male and white. The final two columns report the number of cardiologists observed that year treating at least one heart attack patient within two days of hospital admission, as well as the fraction of physicians who moved that year. Of the 19,945 unique cardiologists in the sample, 15.5% (3,089) moved between 1998-2012.

Table 2: Migration patterns

| Origin Census Region | Destination Census Region | | | | |
	Northeast	Midwest	South	West	Total
Northeast	298	86	189	58	631
Midwest	78	484	272	124	958
South	118	170	725	147	1,160
West	31	58	67	184	340
Total	525	798	1,253	513	3,089

	N	(% of movers)
Same State	663	(21.5%)
Same Census Division	1,218	(39.4%)
Same Census Region	1,691	(54.7%)

Notes: The top panel shows the number of sample cardiologists moving across HRRs, by Census Region of the origin and destination HRRs. The bottom panel shows the number of migrants for whom the move across HRRs remains within various geographic regions.

Table 3: Comparison of migrant cardiologists to non-migrants

Census Region	Geography† % out	Geography† % in	Years since cardiology fellowship‡ migrants p25	migrants median	migrants p75	non-migs p25	non-migs median	non-migs p75	Female migrants	Female non-migs	US Born migrants	US Born non-migs
Northeast	10.7%	8.9%	4	8	16	8	15	21	11.5%	6.9%	65.3%	72.0%
Midwest	18.3%	15.2%	4	8	16	8	14	21	7.5%	5.7%	53.2%	64.1%
South	14.5%	15.6%	4	8	14	7	13	20	7.4%	4.8%	57.0%	67.5%
West	9.0%	13.5%	4	8	16	8	15	22	7.1%	7.1%	59.8%	64.2%
Total	13.5%	13.5%	4	8	15	8	14	21	8.1%	5.8%	57.7%	67.4%

Notes: †The geography comparison describes the fraction of emigrants (% out) and immigrants (% in) to total cardiologists in each region, weighted by total patients each physician treated in that region from 1998-2012. These totals are slightly less than the fraction of sample cardiologists moving between 1998-2012 (15.5%; see Table 1) since some non-migrants practice in multiple hospitals across regional boundaries.

‡Years-since-fellowship for migrants is defined as the time between year of move and cardiology fellowship completion. For non-migrants, years-since-fellowship is defined as the time between a given patient's admission date and the cardiologist's fellowship completion date. Statistics are calculated over physician-patient pairs, and region is that of hospital admission.

Table 4: Difference-in-differences estimates

Dependent variable: (*cath*)$_i$ ∈ {0,1}, indicating cath within 2 days

	Cardiologist movers only Δ in HRR environment Full sample (1)	Full sample (2)	Single first specialist (3)	One admit specialist (4)	Cath lab hospitals (5)	Δ in hospital environment Full sample (6)	One admit specialist (7)	Cath lab hospitals (8)	All cardiologists Δ in HRR environment Full sample (9)	Full sample (10)
Δ	0.037 (0.057)	--	--	--	--	--	--	--	--	--
Δ*(after)	0.628*** (0.055)	0.652*** (0.059)	0.712*** (0.073)	0.626*** (0.089)	0.643*** (0.056)	0.796*** (0.031)	0.770*** (0.050)	0.754*** (0.034)	0.591*** (0.062)	0.652*** (0.059)
mover	--	--	--	--	--	--	--	--	0.002 (0.006)	--
HRR1 FEs	X								X	
HRR2 FEs									X	
Physician FEs		X	X	X	X	X	X	X		X
Observations	124,650	124,650	59,337	41,209	111,429	124,650	41,209	111,429	932,543	932,543

Notes: Table presents difference-in-differences estimates of the change in a physician's practice style across a move as a function of the change Δ in cath environment. Each column presents results from a separate regression. Columns (1-8) are estimated over physician-patient observations limited to migrant physicians only. Columns (9-10) include non-movers for whom Origin HRR = Destination HRR, and thus $\Delta = 0$. The change Δ in cath environment is as defined in Figure 2 except for columns (5,8) where HRR and hospital cath rates are calculated over the subset of patients treated at hospitals with cath lab facilities that year. All regressions include fixed effects for calendar year of patient admission, years since physician move, and patient age, race, sex, and first heart attack. Two-way clustered standard errors at the physician and HRR levels shown in parentheses. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.

Table 5: Cross-sectional estimates

	Full sample	Full sample	One admit specialist
	(1)	(2)	(3)
Panel A: Dependent variable = **1**(*2-day cardiac catheterization*) ∈ {0,1}			
Origin - Current HRR cath rate	0.265***	0.165***	0.112**
	(0.053)	(0.042)	(0.054)
Panel B: Dependent variable = Predicted Pr(*2-day cardiac catheterization*) ∈ (0,1)			
Origin - Current HRR cath rate	0.025**	0.018*	-0.0002
	(0.011)	(0.009)	(0.0138)
HRR FEs	X		
Hospital FEs		X	X
Observations	882,912	882,912	275,496

Notes: Panel A presents estimates of how patient treatment within a region or hospital depends on the treating cardiologist's prior environment, defined as the origin HRR cath rate for movers and the current HRR cath rate for non-movers. Panel B tests for potential sorting, using an empirical measure $\hat{Pr}(cath)$ of patient appropriateness for 2-day cath as the outcome. Each table column corresponds to a separate specification or sample; Panels A and B differ only in the dependent variable. All regressions include calendar year fixed effects and a mover indicator. Two-way clustered standard errors at the cardiologist's origin and current HRRs shown in parentheses. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.

Table 6: Difference-in-differences heterogeneity

Dependent variable: (*cath*)$_i$ ∈ {0,1}, indicating cath within 2 days

	Change in HRR environment			Change in hospital environment		
VAR =	Above-median (>8) years since fellowship	Move to higher-cath HRR ($\Delta > 0$)	Cross-Census Region move	Above-median (>8) years since fellowship	Move to higher-cath hospital ($\Delta > 0$)	Cross-Census Region move
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta*$(after)$*$VAR	-0.078	0.198	0.050	-0.053	0.265***	0.013
	(0.113)	(0.188)	(0.111)	(0.056)	(0.093)	(0.062)
{Δ, after, VAR} one- and two-way effects	X	X	X	X	X	X
Physician FEs	X	X	X	X	X	X
Observations	124,650	124,650	124,650	124,650	124,650	124,650

Notes: Table presents augmented versions of the difference-in-differences estimates (see Table 4) to explore heterogeneity in physician response to a change in cath environment across a move. This is implemented by adding a triple interaction $\Delta*(after)*VAR$, which describes how the difference-in-difference estimate $\Delta*(after)$ changes as VAR increases. All regressions include fixed effects for calendar year of patient admission, years since physician move, and patient age, race, sex, and first heart attack. Two-way clustered standard errors at the physician and HRR levels shown in parentheses. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.

Appendix for Online Publication

Contents

A Sample and variable construction

This section provides additional details on the data, sample selection, and variable measurement in the paper. The primary health care data consist of Medicare administrative and claims records from 1992-2012. These include demographic and enrollment characteristics for 100% of beneficiaries, hospital (MedPAR) records for a 100% sample of the fee-for-service population, and physician (carrier) claims for a 5% sample over years 1992-1997 and for a 20% sample of fee-for-service beneficiaries over the period 1998-2012. The analysis in the paper only relies on claims data from 1998-2012, the period over which physician claims are

available for the 20% sample, but this appendix describes summary statistics for the full period 1992-2012 to provide additional historical context on heart attack (AMI) and cardiac catheterization rates in the Medicare population.

Table C.1 provides a variety of summary statistics related to the sample construction. Column (2) shows the number of unique Medicare beneficiaries eligible for Medicare in each year 1992-2012. The number of enrollees grew from 36.9 million in 1992 to 53.4 million by 2012. Over that same period, the share of beneficiaries enrolled in traditional fee-for-service Medicare dropped from 93.3% in 1992 to 72.4% in 2012 (see column 3). This is important because claims are generally not available for the non-fee-for-service beneficiaries who are enrolled in Medicare Advantage plans which are reimbursed by Medicare on a capitated basis. As a result, the characterization of physician behavior and practice environments in this paper is based on treatment patterns among the Medicare fee-for-service population only.

A.1   AMI cardiac catheterization

The central measure of regional intensity in this paper is the rate of 2-day cardiac catheterization ("cath") among heart attack (AMI) patients in that region. To create this measure, I use 100% MedPAR hospital admission data to identify a "master" sample of new heart attack episodes based on patients admitted to the hospital with a principal diagnosis of AMI (ICD-9-CM codes 410.x) and who had no other AMI hospital admission in the data within the previous year. Column (4) of Table C.1 shows the number of heart attack episodes identified each year, with 5.8 million episodes identified over the period 1992-2012.

For each heart attack patient, I measure whether the patient received a cardiac catheterization within two days of hospital admission. To identify cardiac catheterizations, I use ICD-9 procedure codes and dates submitted on the hospital claim, closely following the well-established methodology used by the Dartmouth Atlas.[13] Specifically, the procedure codes used to indicate cath include 3722-3 (left and combined right and left heart cardiac catheterization); 8855-7 (angiography); 3601-2, 3605, 3609, and 0066 (percutaneous coro-

---

[13]See http://www.dartmouthatlas.org/downloads/methods/research_methods.pdf.

nary intervention); 3606-7 (coronary stent insertion); and 3610-3619 (coronary artery bypass graft). Column (11) of Table C.1 reports the raw 2-day cath rate each year from 1992-2012 over the 100% heart attack sample. Over the 6-year period 1992-1997, 2-day cath rates doubled from 16.3% to 32.7%. While they continued to rise in subsequent years, they did so at a much slower rate, increasing only 71% over the 15-year period 1998-2012.

Cardiac catheterizations are invasive procedures performed in a specialized examination room referred to as a cardiac catheterization (cath) lab. If a heart attack patient is first admitted to a hospital without a cath lab, the patient must be transferred to another hospital in order to receive the procedure. Thus, lack of cath facilities at the hospital of initial admission is likely to be a high barrier to early catheterization. Column (9) of Table C.1 reports the share of all heart attack patients first admitted to a hospital with a cath lab in operation that year. A hospital is defined to have a cath lab in a given year if at least two cardiac catheterization procedures are performed in that hospital that year (based on the 100% MedPAR sample of Medicare patients admitted for any condition, not just AMI). Even in early sample years 1992-1997 when cath rates were low, over 70% of patients were admitted to hospitals with cath facilities. By 2012, nearly 92% of all heart attack patients were admitted to hospitals with cath facilities. As shown in column (7) of Table 1, the fraction of patients admitted to a hospital with a cath lab is even higher among patients seeing a cardiologist within two days of hospital admission. Especially over the period 1998-2012, these results suggest that whether a heart attack patient receives an early cath in most cases is not driven by lack of a cath lab at the hospital of initial admission.

A.1.1 Cardiologist catheterization rates

Beginning in 1992, Medicare required that physicians billing Medicare for services performed must provide their Unique Physician Identification Number (UPIN) on the claim form. Because UPINs uniquely identify individual physicians and remain with the physician throughout their career, it is possible to create histories of the patients a physician treats using the physician billing claims. Beginning in 2007, Medicare transitioned from UPINs to

the National Provider Identifier (NPI) standard. I match NPIs to UPINs using a crosswalk developed by the NBER, and supplemented with Medicare physician claims that contain both UPIN and NPI fields.[14]

To identify histories of cardiologist treatment decisions, I first identify which UPINs correspond to cardiologists. I link the universe of Medicare UPINs to the American Medical Association (AMA) Physician Masterfile and identify cardiologists as those who have completed a 3-year fellowship in cardiovascular disease. The AMA Physician Masterfile includes current and historical data on virtually every Doctor of Medicine (MD) ever trained or licensed to practice in the United States, regardless of physician AMA membership. However, because the merge to the AMA Masterfile was based on UPINs, the set of cardiologists I identify in the sample period 1992-2012 does not include cardiologists first enrolling in Medicare after the transition to the NPI standard in 2007.

As described in Section 2.2.1 of the paper, I measure cardiologist practice styles by linking heart attack patients to the cardiologist(s) who treat them. Because physician claims identify the set of physicians providing services to each patient, I first limit the "master" set of AMI patient episodes to those patients for whom physician claims are available (5% of patients from 1992-1997, and 20% of patients from 1998-2012). I refer to this as the "physician" sample of AMI patients. The number of patients in this sample each year is shown in column (5) of Table C.1. Of the 792,970 AMI patients in the physician sample over the period 1998-2012, 88.2% see at least one cardiologist within two days of hospital admission. These patients are matched to the first cardiologist(s) who treat them to form the basis of the patient-cardiologist observations in the primary regression analysis in the paper. The number of unique cardiologists treating at least one heart attack patient as the first cardiologist in a given year is shown in column (12) of Table C.1, and the average number of AMI patients treated by each cardiologist that year is shown in column (13).

---

[14]The NBER NPI to UPIN crosswalk is available at http://www.nber.org/data/npi-upin-crosswalk.html.

A.1.2 HRR catheterization rates

One of the primary measurement issues in the paper involves constructing 2-day cath rates at the Hospital Referral Region (HRR) level. Because the analysis aims to measure how changes in HRR cath rates across a move drive changes in a migrant's own behavior, it is important to purge the change in region-level cath rates of a mechanical relationship with the migrant's own treatment choices. To do this, I calculate for each cardiologist "leave-out" measures of risk-adjusted regional cath rates that omit the cardiologist's own patients.

Specifically, I first calculate the raw leave-out 2-day cath rate for physician j in HRR h and year t as

$$P(j, h, t) = \frac{1}{|i : i \in H, i \in T, i \notin J|} \sum_{i \in H, i \in T, i \notin J} \left(cath_i \right),$$

where $\left(cath_i \right)$ is an indicator for whether patient i received a cath within two days of hospital admisssion, and where H, T, and J reflect sets of AMI patients treated in HRR h, in year T, and by cardiologist j within two days of hospital admission, respectively.

To risk-adjust the leave-out rates, I calculate the clinical cath appropriateness for each patient by estimating a logistic regression of patient cath within 2 days of a heart attack as a function of calendar year dummies, patient comorbidities (age, race, sex, and first heart attack), and comorbidities interacted with calendar year. Once estimated, the model is used to predict cath receipt for each patient in the sample and these patient-level predictions are averaged at the HRR-year-level to form a new variable $\hat{Pr}(h, t)$ which describes the predicted cath rate in HRR h in calendar year t. The risk-adjusted leave-out rate for physician j in HRR h and year t is then calculated as

$$\tilde{P}(j, h, t) = P(j, h, t) - \hat{Pr}(h, t).$$

$\tilde{P}(j, h, t)$ is a risk-adjusted physician-leave-out mean of the degree to which HRR h is more intensive than the national average in year t, omitting cardiologist j's patients. Using

this measure, I define the physician-leave-out difference in risk-adjusted cath rates between any two HRRs h_1 and h_2 in a given year t as

$$\Delta_j(h_1, h_2, t) = \tilde{P}(j, h_2, t) - \tilde{P}(j, h_1, t). \tag{A.1}$$

A key simplification used throughout the paper is to use time-invariant physician-leave-out differences in cath rates between HRRs over the sample period 1998-2012. To calculate the time invariant differences, I first average $\tilde{P}(j, h, t)$ across years, using as weights the share $w(j, h, t)$ of patients treated each year in that HRR, not counting physician j's patients. This weighted average is a time-invariant physician-leave-out mean of the degree to which a given HRR's cath rate deviated from the national average over the sample period, omitting cardiologist j's patients. The time-invariant physician-leave-out difference between HRRs h_1 and h_2 is then defined as

$$\Delta_j(h_1, h_2) = \sum_t \tilde{P}(j, h_2, t)w(j, h_2, t) - \sum_t \tilde{P}(j, h_1, t)w(j, h_1, t). \tag{A.2}$$

In the paper, I use time-invariant physician-leave-out differences in cath rates $\Delta_j(h_1, h_2)$ over the period 1998-2012 to describe differences in HRR cath environments. Later in this appendix, I explore the relationship between the time-invariant and year-specific cath rates, and also evaluate robustness to using the year-specific differences in cath rates $\Delta_j(h_1, h_2, t)$ to measure physician behavior response to a change in the HRR environment.

A.1.3 Hospital catheterization rates

As an alternative to the HRR definition of a physician's practice environment, I also measure 2-day cath rates at the hospital level. As with the HRR cath rate measures, I calculate hospital cath rates for each physician using a leave-out average that excludes the physician's own patients. Because there are many more hospitals than HRRs, precisely measuring year-specific hospital cath rates using the 20% physician sample of AMI patients is difficult. For this reason, I only consider time-invariant measures of cath intensity at the hospital level,

defined analogously to the time-invariant HRR-level cath intensities described above.

A.2   Cardiologist migration

I define movers to be cardiologists who are observed to move their practice location across Hospital Referral Regions (HRRs). I identify movers, along with their origin and destination HRRs, as follows.

First, I use the "physician" sample of heart attack patients (defined in Section A.1.1) over the period 1998-2012 to identify the first and last dates a cardiologist practices in each HRR (as defined by the date and HRR of hospital admission for each of the cardiologist's patients). I also measure the total number of patients treated by the cardiologist in the HRR. Together, the first/last dates and total number of patients treated characterize the cardiologist's "practice episode" in that HRR. Note that by this definition, it is only possible for a physician to have at most one practice episode per HRR. If a physician moves away from an HRR early in the sample and returns later in the sample, all observations are part of a single practice episode.

Next, I identify in which HRR a physician treats the most patients in the sample, and call this the physician's "primary" practice episode. Similarly, I further define a cardiologist's "secondary" practice episode to be the largest episode (in terms of patients treated) that does not overlap the primary episode, if such an episode exists. Movers are those with both primary and secondary practice episodes, with at least two patients in each episode.

This definition of a move has two additional implications. First, I only identify one move per migrant. Second, the move must involve a clean split in time between the origin and destination HRR, with no overlap in time. If a cardiologist practices in HRR A from dates $d_1 - d_2$ and HRR B from dates $d_3 - d_4$, this would be considered a move as long as $d_2 \leq d_3$. However, if $d_2 > d_3$, which could happen if the cardiologist returns to practice in HRR A after first switching to HRR B, this would not be marked as a move.

Finally, for each mover and corresponding primary and secondary episodes, I mark the earlier practice episode (in terms of practice dates) to be the "origin" HRR and the later

episode to be the "destination" HRR. Some of the specifications in the paper also include non-migrants, which are all cardiologists not identified as movers. For non-migrant cardiologists, I define the origin and destination HRRs to be the same, and equal to the HRR where the cardiologist is treating patients at that point in time. Column (14) of Table C.1 shows the number of cardiologists moving in each year from 1998-2012. Section 2.2.2 of the paper describes additional summary statistics for the migrant sample.

B Robustness

B.1 Time-Varying Cath Rates

B.1.1 Rank-order preservation

In the paper, time-invariant regional cath rates were used to measure the change in intensity a physician experiences across a move. To the extent that the intensity of a region relative to the secular trend remains stable over time, differences in regional cath propensities averaged over a pooled period of time will be the same as the difference in propensity in any given year. However, if regional intensity with respect to the national secular trend changes over time, then measuring regions as having a time-invariant intensity may introduce measurement error into the key independent variable in the analysis $\Delta_j = (destination\ region\ cath\ intensity)_j - (origin\ region\ cath\ intensity)_j$, potentially biasing the estimated environment effects.

HRR cath trends by quartile of average intensity I begin by evaluating the stability of HRR cath intensities relative to the national average over time. As a first approach, I partition the 306 HRRs into quartiles based on each HRR's average risk-adjusted 2-day cath intensity over the period 1998-2012. For each quartile, Figure C.1 plots the average year-specific cath rate across HRRs in that quartile. In 1998, the most intensive quartile of HRRs had a cath rate of 46.4%, compared to 24.9% for the least intensive quartile. This difference of 21.5 percentage points in 1998 had shrunk to 13.1 percentage points in 2012, indicating that absolute differences in cath rates between these groups are not perfectly stable over

time. Importantly, this figure implies that migrants moving, say, from a top-quartile HRR to a bottom-quartile HRR early in the sample period on average experienced somewhat larger changes in regional cath environments across a move than a migrant moving later in the sample.

Rank preservation among top/bottom HRRs   Table C.2 sheds light more directly on the stability of rankings of individual HRRs over time. I first assign each HRR a rank order from 1 (most intensive) to 306 (least intensive) based on each HRR's average 2-day cath intensity over the period 1998-2012. I also rank each HRR based on 1998 cath rates and again on 2012 cath rates. Panels A and B of Table C.2 list the top and bottom 10 HRRs, respectively, based on the average cath rate ranking. While the rank-order of these HRRs was not perfectly preserved from 1998 through 2012, there does appear to be a substantial amount of rank persistence. All regions in Panel A were in the top-25% of regions in both 1998 and 2012, and all but one were in the top-15% both years. Similarly, and regions in Panel B were in the bottom-33% both years, and all but two were in the bottom-10%.

Panels C and D show the top and bottom 10 regions, respectively, based on ranking in 1998. While the most (least) intensive regions sill tended to be more (less) intensive than average in 2012, there is quite a bit of movement in the rank-order for some regions. Perhaps most prominently, McAllen, Texas was the second-most intensive region in 1998, with a 2-day cath rate of nearly 70%, but was the 12th least-intensive region by 2012 with a cath rate of 46.1%. McAllen is an exceptional case in many respects, and was profiled as one of the most expensive health care markets in an influential New Yorker article in 2009 that even highlighted the propensity in McAllen to perform cardiac catheterization (Gawande, 2009). The cath rate in McAllen was already trending downward in McAllen prior to 2009, but this example highlights that it is possible for regions to change relative intensity over time.

Non-parametric rank correlation   While Table C.2 is useful for summarizing the degree of rank preservation for the most and least intensive regions in the sample, it is less useful

for providing a summary measure of rank preservation across all regions. To provide such a summary, I use the year-specific cath rates for each HRR and investigate whether the annual rank-order given by each of these annual intensity measures is preserved over time. If rank-order is preserved, then the intensity measures for any two years should be positively monotonically related.

Table C.3 reports non-parametric Kendall τ_a and Spearman correlation measures for pair-wise comparisons of the annual HRR cath rate measures. When two rankings have a monotonic and positive relationship, both the Kendall and Spearman coefficients are equal to 1, whereas they are both zero when the rankings are independent. The Kendall τ_a also has a convenient interpretation for any value not equal to zero or 1: for any two rank-order measures, the corresponding τ_a coefficient describes how much more likely (in percentage points) the two orderings will agree than disagree for any two randomly selected observations. The estimated correlation coefficients show that rank-order in cath intensity is not fully preserved across years; the fact that the non-parametric correlations drop over time suggests that regions really do change their intensity relative to the secular trend over time, rather than being driven entirely by sampling error (which would result in correlations less than one, but stable over time).

Contemporaneous cath environment trends across move   Finally, I investigate how the contemporaneous difference in cath intensity between a migrant's origin and destination HRRs (as defined in Equation A.1) evolves relative to the average difference in cath intensity between the two HRRs (as defined in Equation A.2). In other words, how well does the time-invariant difference in cath environment capture the year-specific difference in cath environment in the year of the move, as well as the years before and after the move?

I investigate this relationship by estimating a difference-in-differences event study similar to Equation 1, except where the outcome variable is $\Delta_j^t \equiv \Delta_j(origin\ HRR, destination\ HRR, t)$, the contemporaneous difference in cath intensity between migrant j's origin and destination

HRRs t years since the cardiologist moves, omitting physician j's own patients (see Equation A.1). Because the year-specific HRR cath rates are calculated over the patients of non-migrants only, trends in Δ_j^t describe how the migrant's origin and destination environments are differentially evolving across the move, exclusive of the migrant's own choices.

The results of estimating this regression are shown in Figure C.2. In the year of a migrant's move, each unit difference in the time-invariant measure of Δ_j corresponds to just slightly less than 1 unit different in the contemporaneous difference Δ_j^t. This suggests that at least on average, using the time-invariant difference in cath intensity is appropriate for evaluating the change in cath intensity experienced at the time of a physician's move.

Figure C.2 highlights another important consideration. Changes in Δ_j^t around the time of a physician's move are useful for either reinforcing or casting doubt on the validity of the parallel trends assumption underlying the difference-in-differences estimates of physician behavior response which are central to this paper. Specifically, abrupt changes in Δ_j^t at the time of a migrant's move that are driven by abrupt changes in cath rates in the origin environment would suggest that it may not be plausible to assume that the migrant's own behavior absent the move would have continued to follow the abruptly changing trend of cardiologists remaining in the origin region. On the other hand, it would reinforce the plausibility of the parallel trends assumption if Δ_j^t evolved smoothly across migrant moves, which in fact appears to be the case.

B.1.2 Difference-in-differences with time-varying cath environments

In this section, I explore whether the main estimates in the paper that rely on time-invariant differences in HRR cath intensity are sensitive to performing the analysis using year-specific differences in HRR cath intensity.

I begin by re-estimating the event study from Equation 1 in the paper, but replacing the time-invariant difference Δ_j (Equation A.2) in cath intensity between the destination and origin HRRs with the year-specific difference Δ_j^t (Equation A.1). To reduce annual fluctuations in Δ_j^t driven by sampling error, I first smooth the values of Δ_j^t using separate

linear trends for each physician-HRR pair. Smoothing this way requires that changes in Δ_j^t across a physician's move are in fact smooth and approximately linear, an assumption that Figure C.2 supports.

The results of this regression, plotted in solid black in Figure C.3, are very similar to the results obtained using the time-invariant HRR cath rates reported in Figure 3 (also plotted in dashed gray in Figure C.3 for comparison). The main parameters of interest are the β_t coefficients. For a given value of t, β_t describes the difference between treatment styles of physicians t years since move per unit difference in Δ_j^t. The lack of any apparent pre-trend or level difference in physician behavior prior to the move indicates that physicians starting in the same region but moving to different regions practiced similarly before the move, quickly changed their behavior to partially conform to the new practice environment within one year of moving, and experience little to no additional convergence over the next 7 years.

The "step" pattern of physician behavior across a move revealed by the event study suggests that a traditional difference-in-difference (DD) estimate is appropriate for summarizing the change in physician behavior in response to a change in practice environment. The DD estimate comes from replacing the event time dummies in Equation 1 with a single "after" dummy $\mathbf{1}(t \geq 0)$. The DD results are shown in Table C.4, columns (4-8). For comparison, columns (1-3) show DD estimates based on the time-invariant measure of Δ_j. Note that columns (1-2) are the baseline DD estimates from Table 4 based on origin HRR and physician fixed effects, respectively, and repeated here for continuity.

I consider a number of different DD specifications using the year-specific difference Δ_t^j in cath rates between the destination and origin HRRs. First, column (5) repeats the same regression as column (1), except replacing the time-invariant Δ_j with the year-specific Δ_j^t. The resulting estimate changes little across these two specifications.

Columns (5-8) of Table C.4 explore whether the DD estimate appears to meaningfully differ between the first and second half of years in the sample. Since 1998-2012 are the 15 years used to estimate the DD regressions, I partition migrants into those moving prior

to 2005 (the "early sample") and migrants moving in or after 2005 (the "late sample"). I estimate the DD regression over the early sample movers, first using only patients admitted to the hospital prior to 2005 (column 5) and second using only patients admitted to the hospital within 3 years of the cardiologists move year (column 6). Columns (7-8) estimate the DD regression over analogous samples for the late sample movers. The DD estimate remains fairly stable across both the early and late samples, suggesting that environment effects on cardiologist behavior over the period 1998-2012 remained relatively stable.

B.2 Other robustness

B.2.1 Balanced migrant panel

One issue that arises in the difference-in-differences approach estimated using the treatment choices of migrants for up to 8 years before and after a move is that not all migrants are observed for all years in this window. For example, physicians who move in 2000 have at most 2 years of pre-move behavior in the 1998-2012 sample. The trends in pre-move behavior in years 3-8 before the move are therefore only estimated over physicians who move later in the sample. An issue with this imbalance is that trends in measured behavior across a move may partly reflect changes in the composition of which migrants remain in the sample.

I therefore estimate a "balanced panel" DD specification where I estimate the same specification as in column (1) of Table C.4, but with two sample restrictions. First, I limit to the set of 1,358 migrants who are observed treating patients in the sample at least 3 years before and after a move. Second, I further limit the regression sample to patients treated within the same window 3-years before and after the move. The DD estimate from this balanced panel specification, reported in column (3), is again very similar to the other DD estimates in columns (1-2).

B.2.2 Physician volume across moves

As a final robustness check, I explore how the patient volume of migrant physicians changes around the time of the move. If volume changes abruptly across a move, this could

raise concerns that something other than a change in a physician's environment is occurring contemporaneously with the move. While the difference-in-differences framework can account for time-of-move shocks that are common to all physicians who move (as captured by the event time fixed effects), it would be problematic if the size of these shocks were correlated with the change in intensity experienced at the time of the move. In that case, the difference-in-differences estimates would falsely attribute the differential time-of-move shock to the change in regional intensity experienced across the move, even if regional intensity played no role on physician behavior.

I begin by showing summary statistics of the number of sample patients each migrant treats in each of the 8 years before and after a move. I aim to capture how volume may change asymmetrically for physicians who move to more- versus less-intensive regions, and also how volume may change differentially for physicians experiencing a "large" versus "small" change in environment across a move. For ease of analysis, I first normalize the time-invariant cath difference Δ_j between a physician j's origin and destination HRRs into a z-score Δ_z, standardized to have mean zero and variance of 1 across all migrants. This normalization is useful, because $\Delta_z > 0$ essentially captures physicians moving to more-intensive regions (since mean(Δ_j)=0.007 is very close to zero), and $|\Delta_z| > 1$ captures physician's facing a change in cath environment greater than 1 standard deviation.

Figure C.4a plots raw summary statistics of the average number of sample patients each physician treats in the 8 years before and after moving. In the left figure, the statistics are reported separately for physicians moving to more-intensive regions ($\Delta_z > 0$, in orange) and those moving to less-intensive regions ($\Delta_z <= 0$, in blue). There appears to be a slight increase in patient volume following the move, but there does not appear to be any meaningful asymmetry in either levels or trends between the two physician groups. The right column plots the same volume statistics, but separately for physicians experiencing a "large" change in environment ($|\Delta_z| > 1$) versus those facing a small change. Here, patient volume prior to the move appears to be about 15% higher among physicians facing a large change

in environment, with the gap closing and possibly reversing slightly contemporaneous with the move. This differential change in patient volume between the two physician groups may reflect shocks that differentially affect physicians experiencing larger moves, or may simply reflect differential patient volumes common to all cardiologists in the origin and destination HRRs.

Because a cardiologist's patient volume depends not only on shocks to the specific physician, but also depend on characteristics of the HRR, I aim to isolate the physician-specific shocks over time by measuring cardiologist volume relative to the average patient volume across all cardiologists practicing in the same HRR. This is also useful for evaluating whether migrant cardiologists look systematically different in terms of patient volume than non-migrants in the same region.

Figure C.4b plots the same summary statistics as in Figure C.4a, except for relative physician volume. Both the left and right columns tell a similar story: prior to moving, migrants treat a similar number of patients each year compared to non-migrants in the origin HRR. After the move, migrants treat roughly 10% more patients each year than non-migrants in the destination HRR. However, there do not appear to be differential levels or trends in relative patient volume for cardiologists moving to more- versus less-intensive regions, or for cardiologists facing larger versus smaller absolute changes in intensity.

To evaluate more rigorously whether volume changes differentially across a move based on the size of the move, I estimate an event study of the form

$$(relative\ volume)_{jt} = \{origin\ HRR\ FEs\}_j + \sum_{s=-8}^{7} \left[\alpha_t \mathbf{1}(s=t) + \beta_t size(\Delta_z)\mathbf{1}(s=t) \right]$$
$$+ \{calendar\ year\ FEs\}_i + \epsilon_{jt}, \tag{B.1}$$

where an observation is a cardiologist-year over the 8 years before and after a move. The outcome variable $relative\ volume)_{jt}$ is cardiologist j's volume relative to the current HRR average, and is defined each year where the cardiologist treats at least one patient. The key

independent variable of interest is $size(\Delta_z)$, which either takes on the value Δ_z to evaluate the asymmetric effect changes in cath environment on relative volume, or $|\Delta_z|$ to evaluate the volume effect of larger versus smaller absolute changes in the environment.

The results of the volume event study are reported in Figure C.4c. Consistent with the relative volume summary trends shown in Figure C.4b, the event study does not reveal any volume shocks specific to the time of the move that are also correlated with differences in the size of the change in environment. This result further supports the plausibility that changes in physician practice styles across a move are driven by changes in the HRR environment, rather than by idiosyncratic shocks at the time of the move that correlate with the change in environment.

C   Additional Tables and Figures

Figure C.1: Annual 2-day cath rates among heart attack (AMI) patients
by quartiles of average HRR cath intensity

Quartile	# HRRs	HRR Average Cath Rates over 1998-2012		
		Min	Mean	Max
Bottom	76	23.6%	38.0%	42.9%
Second	77	42.9%	45.0%	47.6%
Third	77	47.6%	49.8%	52.4%
Top	76	52.6%	56.2%	75.3%
Total	306	23.6%	45.7%	75.3%
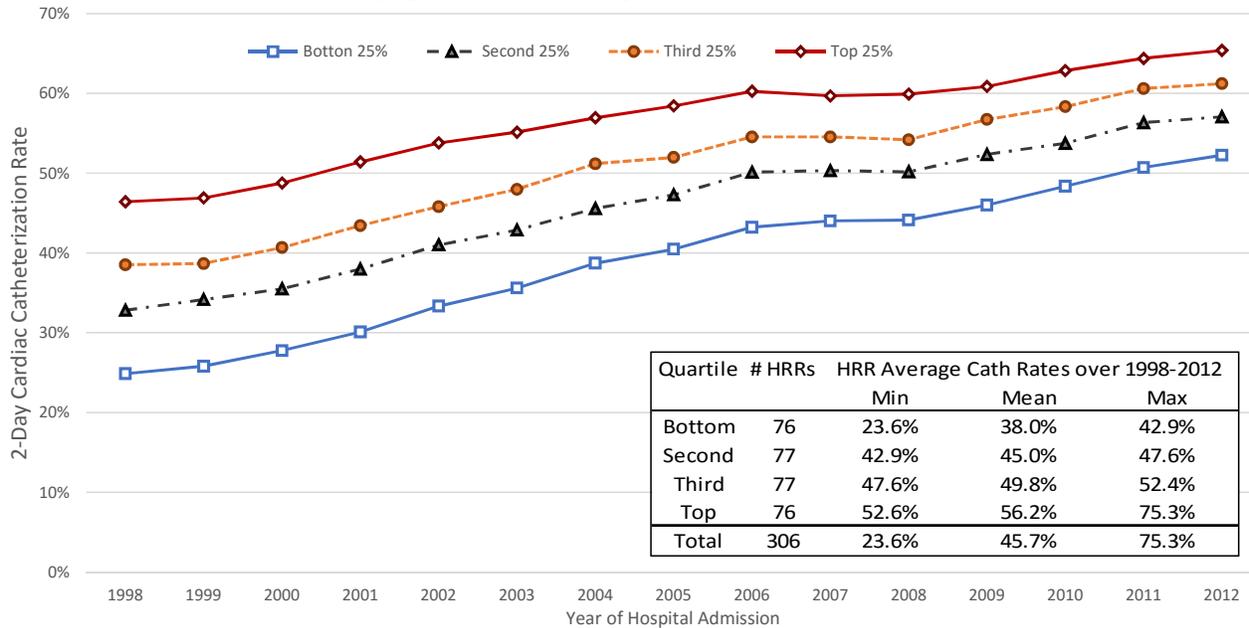
Notes: Figure plots representative cath rates for high and low cath regions by quartile of regional intensity. HRRs are partitioned into quartiles based on average cath rates over all years, such that the composition of HRRs in each quartile remains constant across years. The table shows the distribution of average HRR cath rates over all years, which define the quartiles.

Figure C.2: Contemporaneous vs. average difference between migrant HRR environments
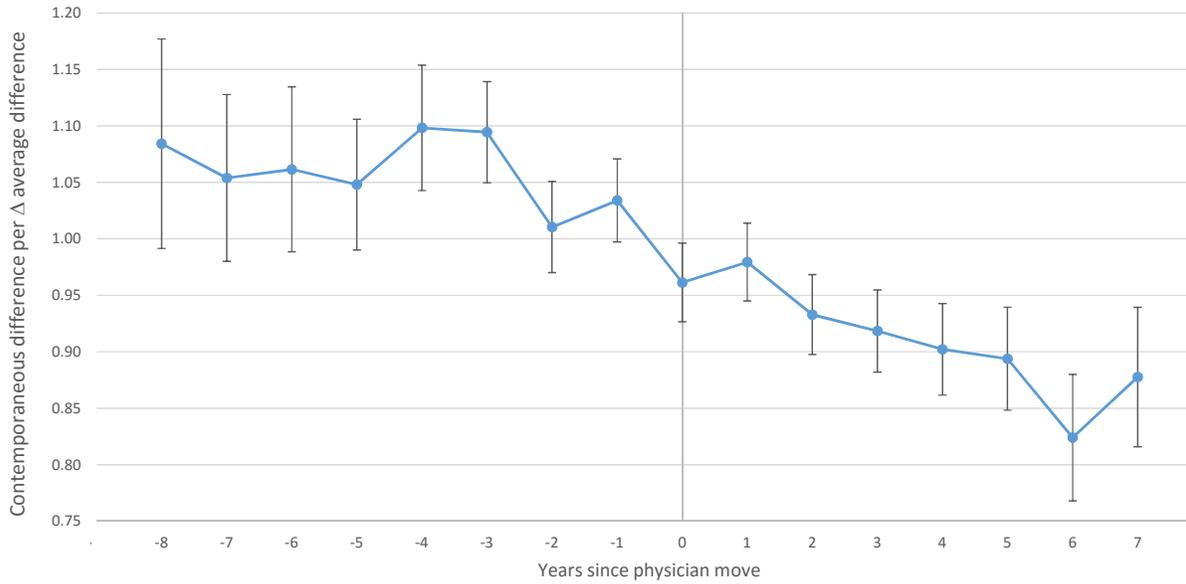
Notes: Graph shows how the year-specific cath difference Δ_j^t between a migrant physician's origin and destination HRRs evolves over time as a function of the time-invariant average cath difference Δ_j between the HRRs. These estimates come from a regression where an observation is a migrant-year and controls include fixed effects for origin HRR, calendar year, and years since physician move. Both Δ_j^t and Δ_j are based on leave-out means that exclude physician j's own patients. Bands indicate 95% confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels.

Figure C.3: Event study–year-specific change in HRR environment
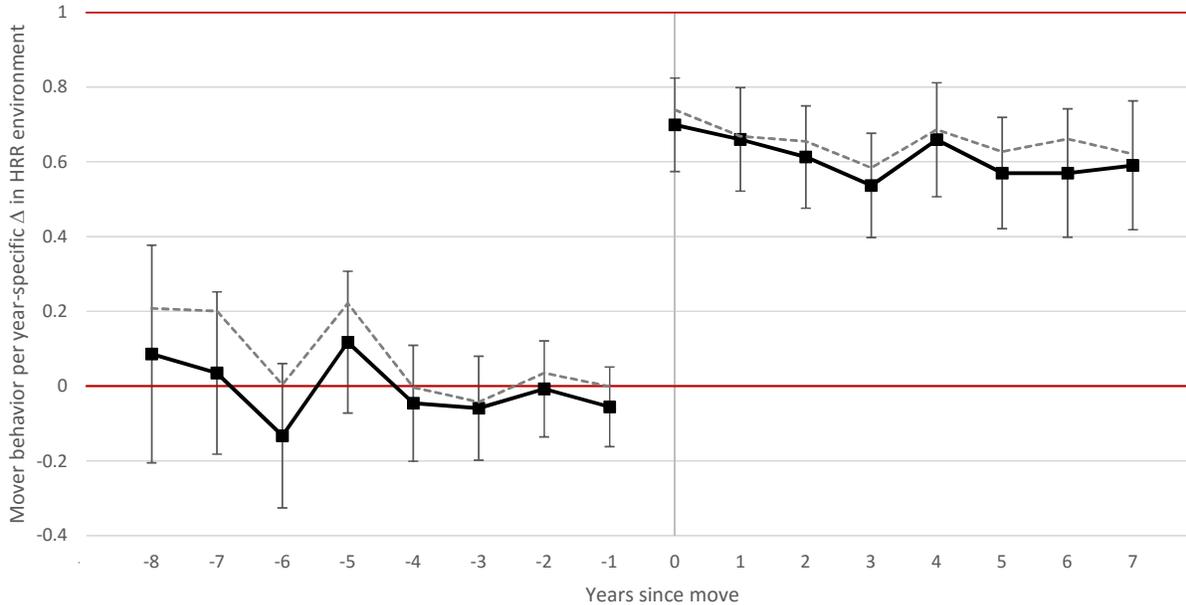
Notes: Graph plots estimates of physician practice style t years since move as a function of the year-specific difference in cath environments Δ_j^t between a migrant physician's destination and origin HRRs. These estimates come from a regression that includes fixed effects for origin HRR, calendar year of patient admission, years since physician move, and patient age, race, sex, and first heart attack. For comparison, the dashed gray line repeats the baseline results based on time-invariant differences in HRR cath rates (Figure 3). Bands indicate 95% confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels.

Figure C.4: Physician volume across a move, by size of move

(a) Physician volume: sample means by size of move
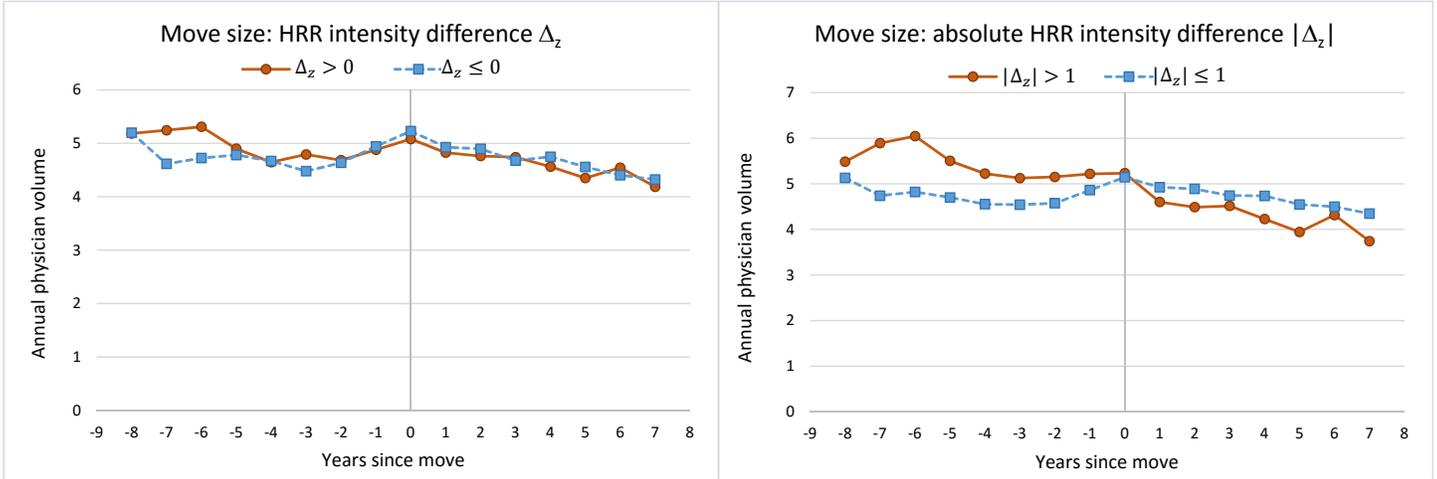
(b) Relative physician volume: sample means by size of move
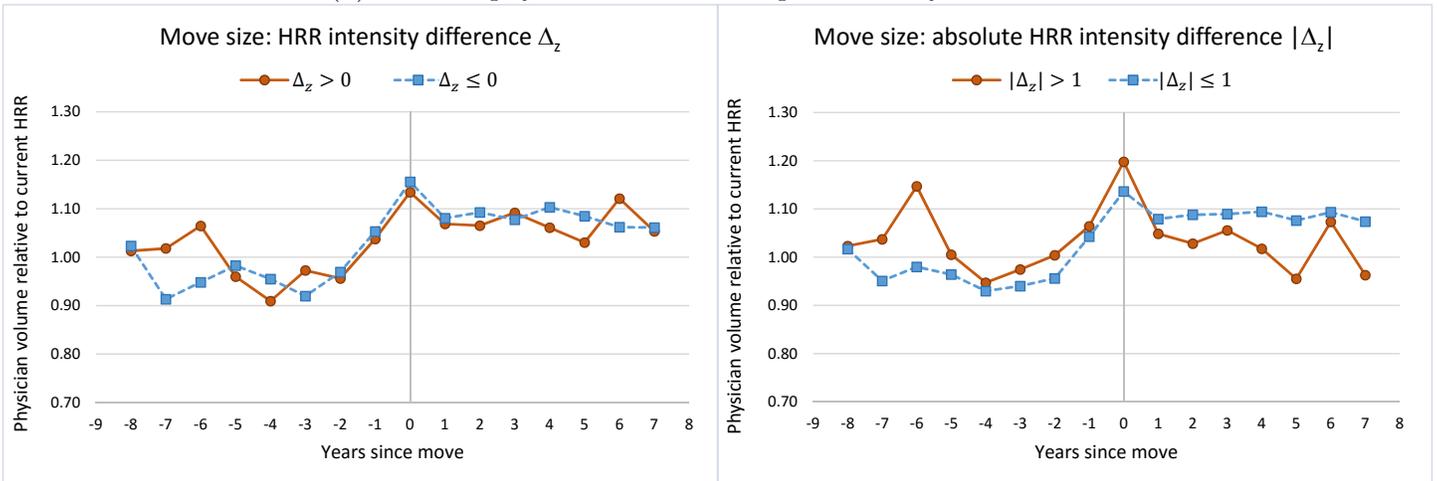
(c) Relative physician volume: event study
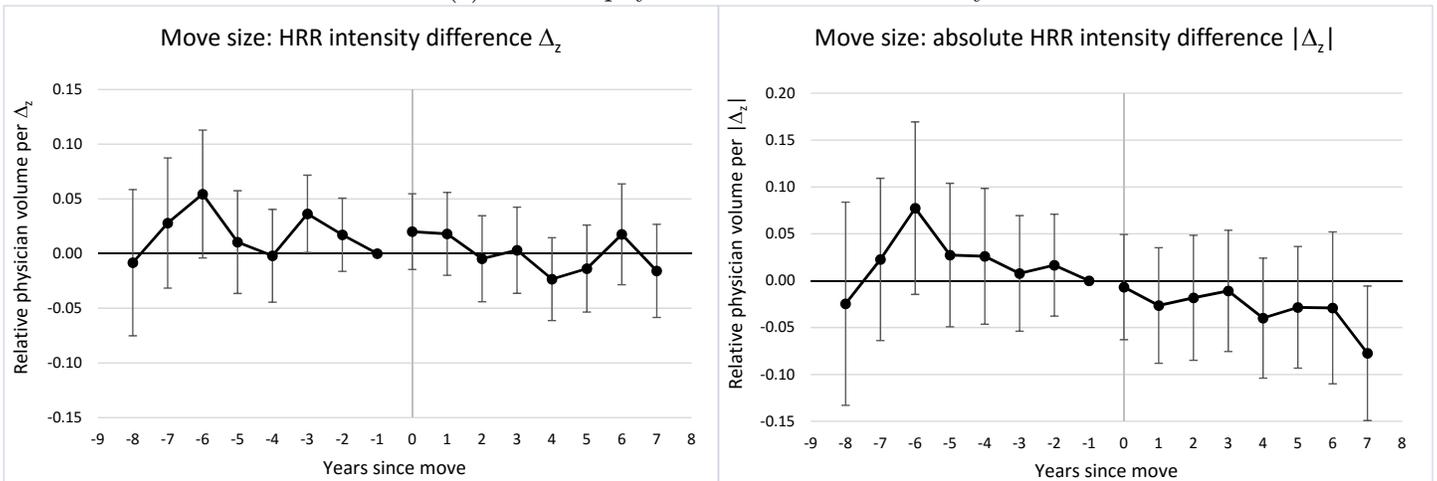
Notes: This figure describes how physician volume, in terms of the number or relative number of AMI sample patients treated each year, varies by time across a move and also by the size of the move. Δ_z is equal to the time-invariant difference in cath between the origin and destination HRRs for each migrant, standardized to have mean zero and a standard deviation of 1 across all migrants. Thus, $\Delta_z > 0$ means an above-average move, while $|\Delta_z| > 1$ indicates a move where the experienced change in cath environment is greater than one standard deviation.

Table C.1: Sample summary statistics

year	Medicare Beneficiaries		AMI Patient Episodes		AMI Patient Characteristics						Cardiologist Characteristics		
	Total (Millions)	Fraction in FFS	Total in FFS	Physician Sample	Age	White	Male	Admitted to Cath Hospital	Cardiologist within 2 days	2-Day Cath Rate	Total	AMI patients as first card	HRR Moves
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
	Panel A: Years in which physician claims available for 5% of Medicare beneficiaries												
1992	36.9	93.3%	288,663	14,505	75.3	89.3%	53.5%	69.6%	65.6%	16.3%	4,955	1.8	-
1993	37.7	92.6%	282,437	14,173	75.3	89.1%	53.6%	71.1%	69.1%	19.3%	5,293	1.7	-
1994	38.7	91.6%	289,514	14,332	75.3	90.5%	52.9%	73.3%	77.7%	23.0%	6,360	2.0	-
1995	39.3	89.7%	290,493	14,233	75.4	90.2%	52.7%	74.4%	80.5%	27.0%	6,751	2.0	-
1996	39.8	87.3%	291,093	14,432	75.6	89.9%	52.4%	75.6%	82.1%	30.6%	7,170	2.0	-
1997	40.2	84.5%	285,060	14,205	75.8	89.5%	52.2%	77.1%	82.1%	32.7%	7,417	2.0	-
	Panel B: Years in which physician claims available for 20% of Medicare beneficiaries												
1998	40.6	82.3%	285,601	56,174	76.1	89.2%	51.7%	77.9%	84.1%	34.2%	11,790	5.0	55
1999	40.9	81.6%	291,399	58,102	76.4	88.7%	51.3%	78.3%	85.6%	35.0%	12,397	5.1	154
2000	41.4	82.1%	300,255	59,131	76.5	89.0%	51.0%	79.3%	87.3%	36.8%	12,842	5.2	221
2001	41.9	83.9%	306,328	60,139	76.6	88.7%	51.1%	80.4%	87.3%	39.3%	13,160	5.3	239
2002	42.4	85.7%	313,879	61,670	76.5	88.2%	51.2%	82.4%	87.9%	42.2%	13,735	5.4	261
2003	43.0	86.6%	308,448	61,016	76.6	87.9%	51.3%	83.5%	89.9%	44.1%	14,047	5.4	262
2004	43.6	86.4%	294,555	58,287	76.5	87.6%	51.4%	85.4%	90.2%	47.0%	14,339	5.2	290
2005	44.5	84.9%	277,937	54,691	76.6	87.5%	51.7%	86.4%	90.5%	48.6%	14,456	4.9	264
2006	45.3	81.4%	257,325	50,593	76.4	87.4%	52.0%	88.0%	90.4%	51.2%	14,598	4.6	323
2007	46.3	79.1%	246,053	48,551	76.5	87.2%	52.1%	88.9%	90.4%	51.4%	14,472	4.4	287
2008	47.5	76.8%	242,494	47,396	76.5	86.9%	52.0%	88.9%	89.8%	51.3%	13,973	4.3	201
2009	48.7	75.3%	229,787	44,881	76.2	86.3%	52.7%	90.1%	89.4%	53.4%	13,443	4.2	164
2010	49.9	74.6%	229,947	45,317	76.1	86.1%	53.0%	90.7%	88.4%	55.2%	13,030	4.2	164
2011	51.5	74.0%	227,058	44,438	75.9	85.6%	53.6%	91.4%	86.7%	57.4%	12,488	4.2	128
2012	53.4	72.4%	217,829	42,584	76.0	85.6%	53.6%	91.8%	85.0%	58.4%	11,903	4.1	76
1998-2012	45.4	80.1%	4,028,895	792,970	76.4	87.6%	51.9%	85.0%	88.2%	46.2%	19,945	4.8	3,089

Notes: Table shows summary statistics related to the sample and variable construction, as discussed in appendix Section A.

Table C.2: HRR cath rank

HRR Number (1)	HRR City (2)	HRR State (3)	HRR Intensity Rank			HRR Cath Rate			Annual AMI Patients (10)
			1998 (4)	2012 (5)	Average (6)	1998 (7)	2012 (8)	Average (9)	
Panel A: Top 10 HRRs, by 1998-2012 Average Cath Rate									
101	Boulder	CO	4	4	1	62.4%	74.6%	75.3%	100.0
422	Provo	UT	9	1	2	57.4%	82.8%	74.8%	229.7
421	Ogden	UT	6	3	3	59.1%	76.3%	70.1%	211.6
104	Fort Collins	CO	3	75	4	67.1%	63.8%	66.0%	228.3
324	Minot	ND	47	8	5	46.8%	73.2%	65.1%	184.1
69	Palm Springs/Rancho Mira	CA	17	22	6	55.2%	69.8%	64.8%	305.7
142	Albany	GA	22	9	7	52.5%	72.6%	64.1%	239.5
190	Cedar Rapids	IA	14	35	8	56.4%	68.4%	63.6%	305.1
370	Rapid City	SD	27	7	9	51.8%	73.3%	62.8%	234.3
14	Sun City	AZ	1	15	10	71.5%	70.9%	62.7%	333.7
Panel B: Bottom 10 HRRs, by 1998-2012 Average Cath Rate									
289	Newark	NJ	291	297	297	18.6%	45.9%	32.9%	1314.4
367	Florence	SC	205	306	298	31.5%	32.3%	32.0%	670.6
350	Danville	PA	305	286	299	11.7%	49.6%	31.7%	751.2
221	Bangor	ME	302	304	300	14.9%	38.4%	31.4%	925.7
230	Springfield	MA	299	298	301	16.9%	45.2%	31.3%	853.8
443	Charleston	WV	294	293	302	17.8%	48.5%	30.9%	1491.9
347	Altoona	PA	301	252	303	15.0%	53.3%	30.9%	364.9
296	Binghamton	NY	303	305	304	14.1%	34.5%	27.9%	533.0
297	Bronx	NY	297	301	305	17.1%	41.1%	27.2%	716.8
299	Buffalo	NY	304	303	306	12.1%	40.2%	23.6%	1405.6
Panel C: Top 10 HRRs, by 1998 Cath Rate									
14	Sun City	AZ	1	15	10	71.5%	70.9%	62.7%	333.7
402	McAllen	TX	2	295	68	69.7%	46.1%	53.0%	518.3
104	Fort Collins	CO	3	75	4	67.1%	63.8%	66.0%	228.3
101	Boulder	CO	4	4	1	62.4%	74.6%	75.3%	100.0
154	Aurora	IL	5	90	21	59.7%	62.8%	59.5%	157.2
421	Ogden	UT	6	3	3	59.1%	76.3%	70.1%	211.6
456	Wausau	WI	7	14	13	58.7%	70.9%	61.5%	280.0
400	Lubbock	TX	8	64	30	58.0%	64.7%	57.5%	750.7
422	Provo	UT	9	1	2	57.4%	82.8%	74.8%	229.7
152	Idaho Falls	ID	10	5	11	57.0%	74.6%	62.4%	141.1
Panel D: Bottom 10 HRRs, by 1998 Cath Rate									
297	Bronx	NY	297	301	305	17.1%	41.1%	27.2%	716.8
420	Wichita Falls	TX	298	239	200	17.0%	54.5%	44.8%	289.5
230	Springfield	MA	299	298	301	16.9%	45.2%	31.3%	853.8
360	Scranton	PA	300	251	290	16.1%	53.5%	34.9%	517.5
347	Altoona	PA	301	252	303	15.0%	53.3%	30.9%	364.9
221	Bangor	ME	302	304	300	14.9%	38.4%	31.4%	925.7
296	Binghamton	NY	303	305	304	14.1%	34.5%	27.9%	533.0
299	Buffalo	NY	304	303	306	12.1%	40.2%	23.6%	1405.6
350	Danville	PA	305	286	299	11.7%	49.6%	31.7%	751.2
242	Muskegon	MI	306	125	283	11.3%	60.4%	37.1%	308.4

Notes: Table describes the top 10 and bottom 10 HRRs by average cath ranking over the period 1998-2012 (Panels A and B) and by 1998 ranking (Panels C and D). Columns (4-6) list the intensity ranks based on 1998, 2012, and average cath rates, respectively, where a rank of 1 indicates the highest-cath HRR and a rank of 306 indicates the lowest-cath HRR. Columns (7-9) show the cath rates from which the respective ranks in columns (4-6) derive.

Table C.3: HRR cath rank

Panel A: Kendall τ Rank-Order Coefficient

year	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1998	1.00 ()													
1999	0.78 (0.02)	1.00 ()												
2000	0.72 (0.02)	0.77 (0.02)	1.00 ()											
2001	0.68 (0.02)	0.71 (0.02)	0.78 (0.02)	1.00 ()										
2002	0.62 (0.03)	0.65 (0.03)	0.72 (0.02)	0.77 (0.02)	1.00 ()									
2003	0.55 (0.03)	0.58 (0.03)	0.65 (0.02)	0.70 (0.02)	0.77 (0.02)	1.00 ()								
2004	0.50 (0.03)	0.52 (0.03)	0.59 (0.03)	0.63 (0.02)	0.68 (0.02)	0.74 (0.02)	1.00 ()							
2005	0.47 (0.03)	0.49 (0.03)	0.56 (0.03)	0.61 (0.02)	0.66 (0.02)	0.73 (0.02)	0.76 (0.02)	1.00 ()						
2006	0.46 (0.03)	0.48 (0.03)	0.52 (0.03)	0.57 (0.03)	0.62 (0.02)	0.68 (0.02)	0.70 (0.02)	0.76 (0.02)	1.00 ()					
2007	0.42 (0.03)	0.43 (0.03)	0.48 (0.03)	0.53 (0.03)	0.58 (0.03)	0.63 (0.02)	0.63 (0.02)	0.71 (0.02)	0.73 (0.02)	1.00 ()				
2008	0.40 (0.03)	0.40 (0.03)	0.45 (0.03)	0.49 (0.03)	0.54 (0.03)	0.58 (0.02)	0.61 (0.02)	0.66 (0.02)	0.70 (0.02)	0.71 (0.02)	1.00 ()			
2009	0.41 (0.03)	0.43 (0.03)	0.47 (0.03)	0.51 (0.03)	0.55 (0.03)	0.60 (0.02)	0.60 (0.02)	0.66 (0.02)	0.70 (0.02)	0.72 (0.02)	0.72 (0.02)	1.00 ()		
2010	0.39 (0.04)	0.40 (0.03)	0.44 (0.03)	0.48 (0.03)	0.53 (0.03)	0.56 (0.03)	0.54 (0.03)	0.63 (0.02)	0.63 (0.02)	0.69 (0.02)	0.68 (0.02)	0.71 (0.02)	1.00 ()	
2011	0.40 (0.03)	0.40 (0.03)	0.44 (0.03)	0.47 (0.03)	0.52 (0.03)	0.56 (0.03)	0.55 (0.03)	0.60 (0.02)	0.62 (0.02)	0.64 (0.02)	0.64 (0.02)	0.69 (0.02)	0.70 (0.02)	1.00 ()
2012	0.41 (0.03)	0.40 (0.04)	0.43 (0.03)	0.46 (0.03)	0.53 (0.03)	0.55 (0.03)	0.54 (0.03)	0.59 (0.03)	0.61 (0.03)	0.63 (0.02)	0.62 (0.03)	0.65 (0.02)	0.68 (0.02)	0.71 (0.02)

Panel B: Spearman Rank-Order Coefficient

year	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1998	1.00 ()													
1999	0.93 (0.02)	1.00 ()												
2000	0.89 (0.03)	0.92 (0.02)	1.00 ()											
2001	0.85 (0.03)	0.88 (0.03)	0.93 (0.02)	1.00 ()										
2002	0.80 (0.03)	0.83 (0.03)	0.89 (0.03)	0.92 (0.02)	1.00 ()									
2003	0.73 (0.04)	0.75 (0.04)	0.83 (0.03)	0.87 (0.03)	0.92 (0.02)	1.00 ()								
2004	0.68 (0.04)	0.69 (0.04)	0.77 (0.04)	0.82 (0.03)	0.86 (0.03)	0.91 (0.02)	1.00 ()							
2005	0.65 (0.04)	0.67 (0.04)	0.74 (0.04)	0.80 (0.03)	0.84 (0.03)	0.91 (0.02)	0.92 (0.02)	1.00 ()						
2006	0.63 (0.04)	0.65 (0.04)	0.70 (0.04)	0.76 (0.04)	0.81 (0.03)	0.86 (0.03)	0.88 (0.03)	0.92 (0.02)	1.00 ()					
2007	0.59 (0.05)	0.59 (0.05)	0.66 (0.04)	0.71 (0.04)	0.77 (0.04)	0.82 (0.03)	0.82 (0.03)	0.88 (0.03)	0.90 (0.03)	1.00 ()				
2008	0.56 (0.05)	0.56 (0.05)	0.62 (0.04)	0.67 (0.04)	0.73 (0.04)	0.78 (0.04)	0.80 (0.03)	0.84 (0.03)	0.87 (0.03)	0.88 (0.03)	1.00 ()			
2009	0.58 (0.05)	0.59 (0.05)	0.65 (0.04)	0.70 (0.04)	0.73 (0.04)	0.79 (0.04)	0.79 (0.03)	0.84 (0.03)	0.88 (0.03)	0.89 (0.03)	0.89 (0.03)	1.00 ()		
2010	0.55 (0.05)	0.56 (0.05)	0.61 (0.05)	0.66 (0.04)	0.71 (0.04)	0.76 (0.04)	0.73 (0.04)	0.81 (0.03)	0.81 (0.03)	0.87 (0.03)	0.85 (0.03)	0.88 (0.03)	1.00 ()	
2011	0.56 (0.05)	0.56 (0.05)	0.62 (0.05)	0.65 (0.04)	0.70 (0.04)	0.75 (0.04)	0.73 (0.04)	0.79 (0.04)	0.81 (0.03)	0.82 (0.03)	0.82 (0.03)	0.87 (0.03)	0.88 (0.03)	1.00 ()
2012	0.57 (0.05)	0.56 (0.05)	0.60 (0.05)	0.63 (0.04)	0.70 (0.04)	0.73 (0.04)	0.72 (0.04)	0.78 (0.04)	0.79 (0.04)	0.81 (0.03)	0.80 (0.03)	0.84 (0.03)	0.85 (0.03)	0.88 (0.03)

Notes: Table gives the Kendall τ_a (Panel A) and Spearman (Panel B) coefficients corresponding to each pair of annual HRR 2-day cath intensity measures (standard errors in parentheses). When two rankings have a monotonic and positive relationship, both the Kendall and Spearman coefficients are equal to 1, whereas they are both zero when the rankings are independent. The Kendall τ_a also has a convenient interpretation for any value not equal to zero or 1: for any two rank-order measures, the corresponding τ_a coefficient describes how much more likely (in percentage points) the two orderings will agree than disagree for any two randomly selected observations.

Table C.4: Difference-in-differences robustness

	Time-invariant Δ in HRR environment			Year-specific Δ in HRR environment				
					early sample: move yr ≤ 2004		late sample: move yr ≥ 2005	
			bal panel: 3 years before/ after move			admits within 3 years of move		admits within 3 years of move
	all movers	all movers		all movers	admit yr ≤ 2004		admit yr ≥ 2005	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ	0.037	--	--	-0.026	0.097	0.035	0.089	-.005
	(0.057)			(0.047)	(0.064)	(0.058)	(0.103)	(0.097)
Δ*(after)	0.628***	0.652***	0.622***	0.652***	0.696***	0.681***	0.569***	0.680***
	(0.055)	(0.059)	(0.080)	(0.053)	(0.073)	(0.066)	(0.105)	(0.098)
Fixed Effects								
HRR1	X			X	X	X	X	X
Physician		X	X					
Observations	124,650	161,944	38,852	161,944	44,492	50,389	31,588	29,233

Notes: Table presents additional difference-in-differences estimates of the change in a physician's practice style across a move as a function of the change Δ in cath environment. Each column presents results from a separate regression. Columns (1-3) use the time-invariant change in cath rates between the origin and destination HRRs, as defined by Equation A.2. Columns (4-8) use the year-specific difference in cath rates between HRRs, as defined by Equation A.1. The balanced panel specification in column (34) restricts to physicians who treat patients at least 3 years before and after the move, and also restricts to patients treated during that time window. All regressions include fixed effects for years since physician move, as well as for patient age, race, sex, and first heart attack. Two-way clustered standard errors at the physician and HRR levels shown in parentheses. *: $p < 0.10$; **: $p < 0.05$; ***: $p < 0.01$.