

The Evolution of Physician Practice Styles: Evidence from Cardiologist Migration[†]

By DAVID MOLITOR*

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 starting in the same region and subsequently moving to dissimilar regions practice similarly before the move. After the move, physician behavior in the first year changes by 0.6–0.8 percentage points for each percentage point change in practice environment, with no further changes over time. This suggests environment factors explain between 60–80 percent of regional disparities in physician behavior. (JEL H75, I11, I12, I18)

Health spending per capita varies dramatically across US 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, Florida region compared with \$7,998 in the Minneapolis, Minnesota 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 the classic work of Wennberg and Gittelsohn (1973) 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

* University of Illinois at Urbana-Champaign, 1206 S. Sixth Street, Champaign, IL 61820 and National Bureau of Economic Research (NBER) (email: dmolitor@illinois.edu). This paper is adapted from the first chapter of my dissertation. I am very grateful to my advisors Amitabh Chandra, Amy Finkelstein, and Jonathan Gruber for their guidance and support. I thank David Chan, Joseph Doyle, Mark Duggan, Iuliana Pascu, Michael Powell, Jonathan Skinner, Heidi Williams, and two anonymous referees for helpful comments. I also gratefully acknowledge feedback from seminar participants at Dartmouth College, Georgia State University, George Washington University, Harvard University, Northeastern University, Massachusetts Institute of Technology, RAND, Stanford University, the University of Illinois at Urbana-Champaign, the University of Pennsylvania, the University of Warwick, and Yale University. Research reported in this publication was supported by the National Institute on Aging of the National Institutes of Health under award number T32-AG000186. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health.

[†] Go to <https://doi.org/10.1257/pol.20160319> to visit the article page for additional materials and author disclosure statement or to comment in the online discussion forum.

¹ Price adjustments remove regional differences in Medicare reimbursement rates, such as higher payments to hospitals with medical training programs. Skinner, Gottlieb, and Carmichael (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.

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; and 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, Fisher, and Skinner 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 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. However, 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, and 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, Ewer, and Wagner 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 predate the current environment (Dranove, Ramanarayanan, and Sfekas 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 variation 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 percent) 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 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 0.48 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 across the move by 60–80 percent of the difference ($D - O$), on average. Moreover, this change in behavior occurs within the first year 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 environment 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 percent 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 2011 provides a review). Finkelstein, Gentzkow, and Williams (2016) exploits 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) uses worker movements across branches in a firm to identify the impact of group interactions on shirking behavior; Song et al. (2010) uses patient migration across geographic regions to identify regional diagnosis propensities separately from patient characteristics; and Chetty, Friedman, and Saez (2013) tracks 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, Dubé, and Gentzkow (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 I describes the empirical context and key data elements, and Section II lays out the empirical strategies and results. In Section III, I briefly discuss potential mechanisms, and Section IV concludes.

I. Setting and Data

A. 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 three 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 (e.g., see 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 invasive" 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, and 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

³Source: Centers for Disease Control and Prevention, <http://www.cdc.gov/NCHS/data/nvsr/nvsr58/nvsr5819.pdf>. Death count from 2007.

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.

B. 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 percent sample of hospital admissions records, which are used to identify over four 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 patient hospital records are matched to physician claims to identify the physicians treating the patient, which limits the sample to the 20 percent 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 percent of Medicare beneficiaries those years and because fewer hospitals had cath labs during that period.

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 three-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

⁴ Over the sample period 1998–2012, Medicare Advantage covered 19.9 percent of Medicare enrollees, although the share fluctuated over that period. See the online Appendix for additional sample details.

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 National Bureau of Economic Research, 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.

TABLE 1—SAMPLE SUMMARY STATISTICS

Year	HRR characteristics				Patient characteristics				Cardiologist characteristics		
	N	2-day cath rate			N	Admitted to cath hospital	Age	Male	White	N	Number of movers
		p25	Mean	p75							
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 percent (3,089) moved between 1998–2012.

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 20 percent random 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 admission. For each patient, I identify the cardiologist(s) who treat the patient first. Because claims only identify the day of service, some patients (34 percent) 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 percent 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).⁷

⁷Physician claims are available for 792,970 (19.7 percent) of the 4.03 million heart attack patients identified by hospital admissions (See Appendix Table C.1). Of these, 669,397 (84.4 percent) have at least one cardiologist claim within 2 days. AMI patients are more likely to have cardiologist claims if admitted to a high-volume hospital: among hospitals with fewer than (at least) 1,000 AMI admissions over the sample—which account for nearly 25 percent of AMI admissions—70.3 percent (89.1 percent) of patients have cardiologist claims. Among these same hospitals, 61.2 percent (9.8 percent) of AMI admissions occur when no cardiologist is the admitting physician for any FFS Medicare patient in the hospital within two days of admission, suggesting that availability of a cardiologist is a key determinant for whether patients have a cardiologist claim. Moreover, patients with no cardiologist claims are over three times more likely (8.4 percent versus 2.7 percent) to die within one day of admission compared to

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, Staiger, and Fisher 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 percent) are identified as movers.⁸

As reported in the last column of Table 1, approximately 1–2 percent of cardiologists observed each year are also identified as moving that year. Table 2 summarizes the migration patterns in the sample: nearly 80 percent of migrants move across states, and over 45 percent 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 percent), and the highest net immigration occurs into the West (4.5 percent). As shown in the last four columns, migrants are also slightly more likely to be female and foreign-born than their nonmigrant counterparts.

Table 3 also compares the time lapse between a migrant's observed move and completion of a cardiology fellowship, relative to the distribution of time since cardiology fellowship completion for nonmigrants in the sample. Half of migrants move within 8 years of finishing cardiology training; comparatively, the median number of years since cardiology training for nonmigrants in the sample is 14 years. The number of years since cardiology training can also be roughly converted to physician age, since cardiologists who go straight through training typically finish at age 32–33 (after

patients with a cardiologist claim, suggesting that sudden death may also partly explain the lack of cardiologist treatment in some cases. On other characteristics, patients with cardiologist claims versus those without are similarly likely to be admitted on a weekend (26.3 percent versus 26.5 percent), are slightly more likely to be male (52.1 percent versus 48.9 percent), and are slightly younger in age (76.2 versus 78.0).

⁸This 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.

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
	Observations	(% of movers)			
Same state	663	(21.5%)			
Same census division	1,218	(39.4%)			
Same census region	1,691	(54.7%)			

Notes: Top panel shows the number of sample cardiologists moving across HRRs, by census region of the origin and destination HRRs. 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 NONMIGRANTS

Census region	Geography ^a		Years since cardiology fellowship ^b						Female		US born	
	% out	% in	Migrants			Nonmigrants			Migrants	Non-migrants	Migrants	Non-migrants
			p25	Median	p75	p25	Median	p75				
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:

^aThe geography comparison describes the fraction of emigrants (percentage out) and immigrants (percentage 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 percent; see Table 1) since some nonmigrants practice in multiple hospitals across regional boundaries.

^bYears-since-fellowship for migrants is defined as the time between year of move and cardiology fellowship completion. For nonmigrants, 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.

college, physicians must complete 4 years of medical school, 3 years internal medicine residency, and a 3 year cardiology fellowship). Thus, the median age in the sample is approximately 41 for migrants (at the time of the move) and 47 for nonmigrants.

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 (two-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 two-day cath rates experienced for each cardiologist 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 a 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, two-day cath rates increased from an average rate of 34.0 percent in 1998 to 57.6 percent 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 propensities 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.

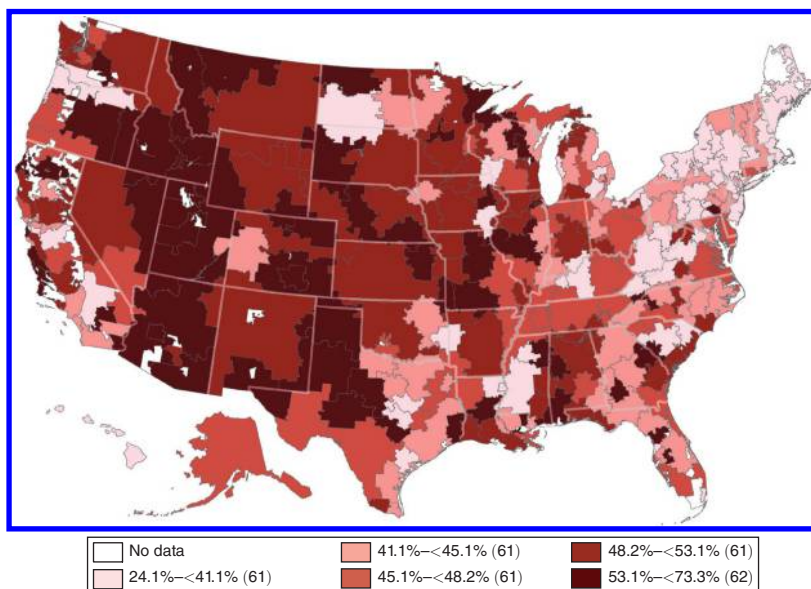


FIGURE 1. DISTRIBUTION OF TWO-DAY CATH RATES BY HRR

Notes: Map shows the geographic distribution of two-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.

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 online Appendix, I provide additional details on the methodology for calculating both HRR and hospital two-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 percent in the lowest quintile of regions to more than 53 percent in the highest quintile. The change in physician j 's practice environment experienced across a move is calculated as $\Delta_j = (\text{destination region cath intensity})_j - (\text{origin region cath intensity})_j$, 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.

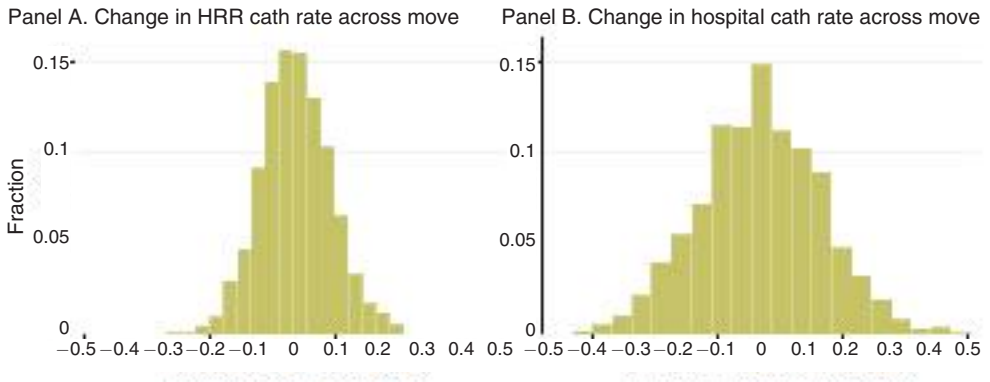


FIGURE 2. DISTRIBUTION OF CHANGES IN TWO-DAY CATH ENVIRONMENT ACROSS MOVE

Notes: Figure 2 shows the distribution of changes in cath environment among cardiologists who move HRRs. In panel A, the change for cardiologist j is defined as $\Delta_j = (\text{destination HRR cath intensity})_j - (\text{origin HRR cath intensity})_j$, 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.

II. Empirical Evidence

A. Difference-in-Differences

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 a 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 $(\text{cath})_{ijt}$ is an indicator for whether heart attack patient i , treated by cardiologist j in event year t , received a cath within two 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 eight years before or after a move, yielding the regression equation

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

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 physician's 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 eight 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 a 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.⁹ 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.

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 percent 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 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,

⁹Since physicians may practice at multiple hospitals within the origin HRR, an "origin" hospital is not well-defined for many physicians.




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 controlling for physician fixed effects instead of origin HRR are plotted by the dashed gray line. Bands indicate 95 percent confidence intervals constructed from two-way clustered standard errors at the physician and HRR levels.

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 $t \geq 0$ 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 IIB 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 mismeasure 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).

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 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-differences 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 of Table 4 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

TABLE 4—DIFFERENCE-IN-DIFFERENCES ESTIMATES

Dependent variable: $(cath)_i \in \{0,1\}$, indicating cath within 2 days	Cardiologist movers only								All cardiologists	
	Δ in HRR environment					Δ in hospital environment			Δ in HRR environment	
	Full sample (1)	Full sample (2)	Single first specialist (3)	One admit specialist (4)	Cath lab hospitals (5)	Full sample (6)	One admit specialist (7)	Cath lab hospitals (8)	Full sample (9)	Full sample (10)
Δ	0.037 (0.057)	—	—	—	—	—	—	—	—	—
$\Delta \times$ (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 and 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 are shown in parentheses.

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 one-for-one 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 of Table 4 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 Δ_j^{hosp} 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.

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 percent 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 hospital.¹⁰ 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 Table 4, 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 cath 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. Eighty-four percent 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,

¹⁰The 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 percent 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).

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.

As a final robustness check, I reestimate the HRR environment specification using both migrant and nonmigrant 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 nonmigrants. The second benefit is that it allows a direct comparison of migrant and nonmigrant behavior. In particular, in a specification without physician fixed effects, I will estimate whether migrants are more or less intensive on average than nonmigrants.

To estimate this regression, I define the origin and destination HRRs for nonmigrants 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 nonmigrants.

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 nonmigrants 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).

B. Cross Section

Measurement Error and Estimate Bounds.—While the difference-in-differences results (Section IIA) suggest that physicians are highly responsive to changes in their environment, the results also show less than one-for-one 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 mismeasurement 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 $g(\cdot)$ is the potentially limited set of information available to the econometrician about a physician's change in environment. For example, $g(\cdot)$ 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.¹¹

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 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. Empirically, the higher difference-in-differences estimates when using hospital-level environment changes compared with HRR-level changes (0.75–0.80 versus 0.63–0.71, respectively; see Table 4) could reflect the presence of this type of measurement error.¹²

A second potential source of mismeasurement is classical measurement error, which through attenuation bias would also lead to underestimates 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.¹³

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

¹¹ 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.

¹² The difference in estimates when using HRR- and hospital-level environment changes could also reflect direct impacts the physician has on local practice patterns, a problem which is mitigated when focusing on broader definitions of the region such as the HRR. Section IB discusses this issue in greater detail.

¹³ While we can perform the analysis using time-varying rates directly as in Appendix 2.1, 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.

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 nonconformance, I adopt a “cross-sectional” approach and estimate nonconformance 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 positively sort into practice environments, lack of conformance θ is given by

$$(2) (cath)_{ij} = \theta[(origin\ HRR\ cath\ intensity)_j - (current\ HRR\ cath\ intensity)_j] \\ + \{environment\ FEs\}_i + \{year\ FEs\}_i + \epsilon_{ij}.$$

To estimate this regression, I include patient-physician observations for both post-move migrants and nonmigrants, where the origin HRR is defined to be the current HRR for nonmigrants. 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 nonconformance 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 one.

The results from this regression are shown in panel A of Table 5. Column 1 shows results when the 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 $(origin\ HRR\ cath\ intensity)_j$ 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 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

TABLE 5—CROSS-SECTIONAL ESTIMATES

	Full sample (1)	Full sample (2)	One admit specialist (3)
<i>Panel A. Dependent variable = $\mathbf{1}(2\text{-day cardiac catheterization}) \in \{0, 1\}$</i>			
Origin—current HRR cath rate	0.265 (0.053)	0.165 (0.042)	0.112 (0.054)
<i>Panel B. Dependent variable = Predicted $\Pr(2\text{-day cardiac catheterization}) \in (0, 1)$</i>			
Origin—current HRR cath rate	0.025 (0.011)	0.018 (0.009)	−0.0002 (0.0138)
HRR fixed effects	X		
Hospital fixed effects		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 $\Pr(\text{cath})$ of patient appropriateness for two-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 are shown in parentheses.

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.


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 two days of a heart attack, as

$$(3) \quad \Pr(\text{cath}_{iht}) = G[\theta_h + \theta_h t + X_{it} \Phi].$$

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. The term 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.

Panel A. Average physician response



Panel B. Asymmetric physician response

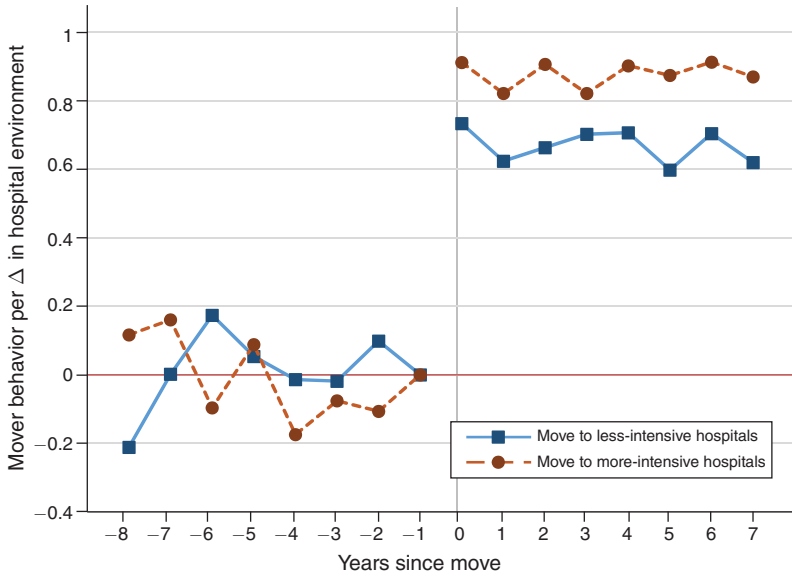


FIGURE 4. EVENT STUDY—CHANGE IN HOSPITAL ENVIRONMENT

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 percent 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 ($t = -1$).

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 percent 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 IIA. 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 IIA.

C. 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.

“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 eight 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 IIA based on whether the physician was more than eight years post-fellowship completion (the median among migrants) at the time of the move. The specific regression I estimate takes the form:

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

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 5, respectively. In both cases, $\hat{\beta}$, is negative but small, and not significant at the 5 percent 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—at least among those who chose to move—remain elastic over time.

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. However, 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.

TABLE 6—DIFFERENCE-IN-DIFFERENCES HETEROGENEITY

Dependent variable: $(cath)_i \in \{0,1\}$, indicating cath within 2 days	Change in HRR environment				Change in hospital environment			
	Above-median (>8) years since fellowship	Move to higher-cath HRR ($\Delta > 0$)	STEMI	Cross-census region move	Above-median (>8) years since fellowship	Move to higher-cath hospital ($\Delta > 0$)	STEMI	Cross-census region move
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \times (\text{after})$	-0.078	0.198	-0.132	0.050	-0.053	0.265	-0.032	0.013
$\times \text{VAR}$	(0.113)	(0.188)	(0.045)	(0.111)	(0.056)	(0.093)	(0.027)	(0.062)
$\{\Delta, \text{after}, \text{VAR}\}$ one- and two- way effects	X	X	X	X	X	X	X	X
Physician FEs	X	X	X	X	X	X	X	X
Observations	124,650	124,650	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 $\Delta \times (\text{after}) \times \text{VAR}$, which describes how the difference-in-difference estimate $\Delta \times (\text{after})$ 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 are shown in parentheses.

Panel B of Figure 4 plots the results of the event study of physician behavior across a 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 1 percentage point more-intensive hospital changes behavior by about 0.21 percentage points more than a physician moving to a 1 percentage point 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 6, 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 percent level). Both of 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.

Heart Attack Type.—Guidelines for heart attack treatment generally distinguish between two types of heart attacks, characterized by their electrocardiogram tracings. The first type is an ST-elevation myocardial infarction (STEMI), caused by complete blockage of an artery in the heart. The second type is a non-ST-elevation myocardial infarction (NSTEMI), caused by partial blockage. Current clinical evidence and guidelines generally support treating most STEMI patients invasively by diagnostic catheterization with an intent to perform revascularization (O’Gara et al. 2013). In contrast, medical guidelines for NSTEMI patients suggest a more nuanced risk-management strategy to determine patient appropriateness for early

invasive treatment (Amsterdam et al. 2014). Consistent with medical practice broadly following these guidelines, online Appendix Table C.1 (column 14) shows that in 2012, the two-day STEMI cath rate was 26.9 percentage points (52 percent) higher than the NSTEMI cath rate.

The higher degree of medical uncertainty regarding the benefits of early invasive management of NSTEMI versus STEMI patients over the sample period provides an opportunity to test whether physicians are more responsive to changes in their environment based on the degree of medical uncertainty. I code each heart attack patient as NSTEMI if they were diagnosed with a subendocardial infarction (ICD-9-CM codes 410.7), labeling all other heart attacks as STEMI. I estimate equation (4) above, replacing $\mathbf{1}(tsf_j > 8)$ by an indicator for a STEMI heart attack. The results from HRR and hospital changes in cath environment across a move are reported in Table 6, columns 3 and 7, respectively. The point estimate based on changes in HRR cath rates is -0.132 (significant at the 1 percent level), and the estimate based on hospital cath rates is -0.032 (though not statistically significant). Taken together, these estimates suggest that physicians change their treatment of patients more when medical benefits are less certain. Given that STEMI cath rates are higher than NSTEMI cath rates, this result also suggests that hard capacity constraints (such as intraday or intra-week availability in a non-24/7 cath lab setting) that are orthogonal to patient cath appropriateness are not the exclusive factor driving changes in physician behavior in this setting.

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 to 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 nonmigrants 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 percent of migrants who moved across census regions (Table 2). As reported in Table 6, columns 4 and 8, physicians who move across census regions respond about the same to changes in their environment as those who do not.

III. Mechanisms

The results in Section II 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 percent 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 cath, 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.¹⁴ 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. It seems plausible that the first cardiologist treating a heart attack patient plays a key role in deciding on cardiac care options especially in cases that are not clear-cut. The robustness of the key estimates in this paper to situations where there is only one cardiologist suggests the change in treatment patterns across a move does not simply reflect new cardiologists passively deferring to other

¹⁴This 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.

cardiologists also treating the patient. However, this does not rule out a possibly key role played by other team members. 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 from 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 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. An important direction for future work is to understand the role of specialists on teams for influencing patient care and to understand the factors that shape team practice patterns.

IV. Conclusion

Cardiologists vary widely across US 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. However, 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 policymakers 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

- Amsterdam, Ezra A., Nanette K. Wenger, Ralph G. Brindis, Donald E. Casey, Theodore G. Ganiats, David R. Holmes, Allan S. Jaffe, et al. 2014. "2014 AHA/ACC Guideline for the Management of Patients with Non-ST-Elevation Acute Coronary Syndromes: Executive Summary." *Journal of the American College of Cardiology* 64 (24): 2645–87.
- Angrist, Joshua D. 2014. "The perils of peer effects." *Labour Economics* 30: 98–108.
- Baicker, Katherine, and Amitabh Chandra. 2004. "Medicare Spending, the Physician Workforce, and Beneficiaries' Quality of Care." *Health Affairs* W4: W184–97.
- Barnato, Amber E., M. Brooke Herndon, Denise L. Anthony, Patricia M. Gallagher, Jonathan S. Skinner, Julie P. W. Bynum, and Elliott S. Fisher. 2007. "Are Regional Variations in End-of-Life Care Intensity Explained by Patient Preferences? A Study of the US Medicare Population." *Medical Care* 45 (5): 386–93.
- Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review* 102 (6): 2472–2508.
- Brophy, James M., and Peter Bogaty. 2004. "Primary Angioplasty and Thrombolysis Are Both Reasonable Options in Acute Myocardial Infarction." *Annals of Internal Medicine* 141 (4): 292–97.
- Chandra, Amitabh, and Douglas O. Staiger. 2007. "Productivity Spillovers in Health Care: Evidence from the Treatment of Heart Attacks." *Journal of Political Economy* 115 (1): 103–40.
- Chetty, Raj, John N. Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review* 103 (7): 2683–2721.
- Congressional Budget Office. 2008. *Geographic Variation in Health Care Spending*. Congressional Budget Office. Washington, DC, February.
- Doyle, Joseph J. 2016. "Physician Characteristics and Patient Survival: Evidence from Physician Availability." Unpublished.
- Doyle, Joseph J., Steven M. Ewer, and Todd H. Wagner. 2010. "Returns to physician human capital: Evidence from patients randomized to physician teams." *Journal of Health Economics* 29 (6): 866–82.
- Dranove, David, Subramaniam Ramanarayanan, and Andrew Sfekas. 2011. "Does the Market Punish Aggressive Experts? Evidence from Cesarean Sections." *B.E. Journal of Economic Analysis and Policy* 11 (2): Article 5.
- Epstein, Andrew J., and Sean Nicholson. 2009. "The formation and evolution of physician treatment styles: An application to cesarean sections." *Journal of Health Economics* 28 (6): 1126–40.
- Fernández, Raquel. 2011. "Does Culture Matter?" In *Handbook of Social Economics*, Vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 481–510. Amsterdam: North-Holland.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "Sources of Geographic Variation in Health Care: Evidence from Patient Migration." *Quarterly Journal of Economics* 131 (4): 1681–1726.
- Fisher, Elliott S., David E. Wennberg, Thérèse A. Stukel, Daniel J. Gottlieb, F. L. Lucas, and Étoile L. Pinder. 2003a. "The implications of regional variations in Medicare spending. Part 1: The content, quality, and accessibility of care." *Annals of Internal Medicine* 138 (4): 273–87.
- Fisher, Elliott S., David E. Wennberg, Thérèse A. Stukel, Daniel J. Gottlieb, F. L. Lucas, and Étoile L. Pinder. 2003b. "The implications of regional variations in Medicare spending. Part 2: Health outcomes and satisfaction with care." *Annals of Internal Medicine* 138 (4): 288–98.
- Gatsonis, Constantine A., Arnold M. Epstein, Joseph P. Newhouse, Sharon-Lise Normand, and Barbara J. McNeil. 1995. "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): 625–42.
- Glover, J. Alison. 1938. "The Incidence of Tonsillectomy in School Children." *Indian Journal of Pediatrics* 5 (4): 252–58.
- Gottlieb, Daniel J., Weiping Zhou, Yunjie Song, Kathryn Gilman Andrews, Jonathan S. Skinner, and Jason M. Sutherland. 2010. "Prices Don't Drive Regional Medicare Spending Variations." *Health Affairs* 29 (3): 537–43.
- Grytten, Jostein, and Rune Sørensen. 2003. "Practice variation and physician-specific effects." *Journal of Health Economics* 22 (3): 403–18.
- Ichino, Andrea, and Giovanni Maggi. 2000. "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm." *Quarterly Journal of Economics* 115 (3): 1057–90.
- Jollis, James G., Elizabeth R. DeLong, Eric D. Peterson, Lawrence H. Muhlbaier, Donald F. Fortin, Robert M. Califf, and Daniel B. Mark. 1996. "Outcome of Acute Myocardial Infarction

- According to the Specialty of the Admitting Physician." *New England Journal of Medicine* 335 (25): 1880–87.
- Keeley, Ellen C., and Cindy L. Grines.** 2004. "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.
- Kushner, Frederick G., Mary Hand, Sidney C. Smith, Spencer B. King, Jeffrey L. Anderson, Elliott M. Antman, Steven R. Bailey, et al.** 2009. "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.
- Manski, Charles F.** 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60 (3): 531–42.
- Martin, Anne B., Lekha Whittle, Stephen Heffler, Mary Carol Barron, Andrea Sisko, and Benjamin Washington.** 2007. "Health spending by state of residence, 1991–2004." *Health Affairs* 26 (6): w651–63.
- McClellan, Mark, and Joseph P. Newhouse.** 1997. "The marginal cost-effectiveness of medical technology: A panel instrumental-variables approach." *Journal of Econometrics* 77 (1): 39–64.
- Molitor, David.** 2018. "The Evolution of Physician Practice Styles: Evidence from Cardiologist Migration: Dataset." *American Economic Journal: Economic Policy*. <https://doi.org/10.1257/pol.20160319>.
- O'Gara, Patrick T., Frederick G. Kushner, Deborah D. Ascheim, Donald E. Casey, Mina K. Chung, James A. de Lemos, Steven M. Ettinger, et al.** 2013. "2013 ACCF/AHA Guideline for the Management of ST-Elevation Myocardial Infarction." *Circulation* 127 (4): 362–425.
- Phelps, Charles E.** 1992. "Diffusion of Information in Medical Care." *Journal of Economic Perspectives* 6 (3): 23–42.
- Phelps, Charles E.** 2000. "Information Diffusion and Best Practice Adoption." In *Handbook of Health Economics*, Vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 223–64. Amsterdam: North-Holland.
- Phelps, Charles E., and Cathleen Mooney.** 1993. "Variations in Medical Practice Use: Causes and Consequences." In *Competitive Approaches to Health Care Reform*, edited by Richard J. Arnould, Robert F. Rich, and William D. White, 139–75. Washington, DC: Urban Institute Press.
- Philipson, Tomas J., Seth A. Seabury, Lee M. Lockwood, Dana P. Goldman, and Darius N. Lakdawalla.** 2010. "Geographic Variation in Health Care: The Role of Private Markets." *Brookings Papers on Economic Activity* 35 (1): 325–61.
- Scanlon, Patrick J., David P. Faxon, Anne-Marie Audet, Blase Carabello, Gregory J. Dehmer, Kim A. Eagle, Ronald D. Legako, et al.** 1999. "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." *Journal of the American College of Cardiology* 33 (6): 1756–1824.
- Sirovich, Brenda E., Daniel J. Gottlieb, H. Gilbert Welch, and Elliott S. Fisher.** 2006. "Regional Variations in Health Care Intensity and Physician Perceptions of Quality of Care." *Annals of Internal Medicine* 144 (9): 641–49.
- Skinner, Jonathan, and Elliot Fisher.** 1997. "Regional Disparities in Medicare Expenditures: An Opportunity for Reform." *National Tax Journal* 50 (3): 413–25.
- Skinner, Jonathan S., Daniel J. Gottlieb, and Donald Carmichael.** 2011. *A New Series of Medicare Expenditure Measures by Hospital Referral Region: 2003–2008*. Dartmouth Institute for Health Policy and Clinical Practice. Hanover, NH, June.
- Skinner, Jonathan S., Douglas O. Staiger, and Elliott S. Fisher.** 2006. "Is Technological Change in Medicine Always Worth It? The Case of Acute Myocardial Infarction." *Health Affairs* 25 (2): W34–47.
- Song, Yunjie, Jonathan Skinner, Julie Bynum, Jason Sutherland, John E. Wennberg, and Elliott S. Fisher.** 2010. "Regional Variations in Diagnostic Practices." *New England Journal of Medicine* 363 (1): 45–53.
- Wennberg, John E., Elliott S. Fisher, and Jonathan S. Skinner.** 2002. "Geography and the debate over Medicare reform." *Health Affairs* W2: W96–114.
- Wennberg, John, and Alan Gittelsohn.** 1973. "Small Area Variations in Health Care Delivery." *Science* 182 (4117): 1102–08.
- Zuckerman, Stephen, Timothy Waidmann, Robert Berenson, and Jack Hadley.** 2010. "Clarifying Sources of Geographic Differences in Medicare Spending." *New England Journal of Medicine* 363 (1): 54–62.