You Had Me At Hello: The Effects of Disruptions to the Patient-Physician Relationship.

Stephen D Schwab

Organizations exist largely to coordinate tasks across individuals with different knowledge (Grant 1996). This knowledge encompasses differences in individual skills, experience and education, as well as "the knowledge of the particular circumstances of time and place," as Hayek (1945) aptly put it. Economists have long recognized, however, that transferring knowledge is costly and that the division of labor increases these coordination costs. In this paper I offer a quantification of these costs in a setting in which significant effort has been placed on reducing coordination costs and which is characterized by considerable uncertainty - the US Health Care System.

The US Health Care system is a prime example, both of the fragmented nature of knowledge and the associated coordination costs. Care is often fragmented across multiple providers (Agha, Frandsen, and Rebitzer 2019) and failures of coordination have been linked to $27 billion to $78 billion in waste each year (Shrank, Rogstad, and Parekh 2019). Fragmented care, though, is not unambiguously negative. First, benefits from the division of labor (hereafter referred to as specialization) may dominate the resulting coordination costs. Patients may suffer from multiple conditions and require different physician specialists, or may transition between sites of care such as from an acute care hospital to a skilled nursing facility. Considerable managerial and policy effort has been exerted into limiting coordination costs while accruing the benefits from specialization. Communication and delivery innovations such as electronic health records (EHR), service line management, and accountable care organizations (ACO's) have all been designed to reduce the costs of transferring knowledge within and between health care organizations.

Fragmented knowledge, however, is not solely caused by specialization. Patients may also decide to see a different physician within the same specialty or site of care. This discontinuity in care (also referred to as "churn" in the medical literature) occurs for a number of reasons including insurance changes, dissatisfaction with quality of care, and patient or physician moves. The potential to change providers offers several benefits. For instance, competition among providers likely increases quality (Allard, Thomas Léger, and Rochaix 2009) and patients and providers may sort into better matches.

While both specialization and discontinuity of care result in knowledge fragmentation, they are conceptually different. While the former potentially decreases the *average* knowledge about a patient due to dispersion across specialists on the care team, discontinuity decreases the *absolute* level of knowledge about the patient within the care team. While significant research has focused on fragmentation of care, discontinuity of care has been largely overlooked in the economics literature.[1] In this paper, I estimate the causal effects of discontinuity in the patient-provider relationship on healthcare utilization. I focus my analysis on the relationship between a patient and his primary care provider (PCP).

A specific challenge of estimating the effects of discontinuity in this relationship is that there are multiple confounders due to the dissociated nature of the US health care system. Electronic medical records often do not transfer between providers, and a patient's primary care provider (PCP) may not even know if her patient is admitted to the hospital (Arora et al. 2010). Any causal analysis of disruptions in the patient-provider relationship needs to isolate changes in the relationship from the effects of changing practices and/or insurance. A secondary challenge is that turnover is rarely exogenous. Providers that leave their practice or retire may differ in unobservable ways from those that remain. I address both of these issues through a unique data source - the Military Health System (MHS). The MHS provides a platform for patient care as well as to maintain the medical skills of military health care providers (Hutter et al. 2019). As military operational missions arise, these providers are

[1] one notable exception is Meltzer (2001) who documented adverse effects from changing hospitalists in an inpatient setting.

pulled from their practices and deployed outside the United States. I use these wartime deployments as a source of exogenous variation in the patient-provider relationship in order to estimate the effects of discontinuity in care. Patients affected by the discontinuity remain enrolled in the same practice, with the same EHR, and same insurance design. As a secondary benefit, using data from the MHS lets me consider the effects on a working age population, a group that has been typically understudied in the health economics literature due to a lack of publicly available data. Despite some idiosyncrasies, the military health system is comparable to the civilian health care system (Bond and Schwab 2019) and has been previously used in the economics literature to study the impact of defensive medicine (Frakes and Gruber 2019).

I restrict my analysis to primary care provider deployments for several reasons associated with the nature of knowledge. First, as health care costs have risen health insurers have increasingly turned to the PCP to limit utilization. Many insurance plans use a 'gate-keeping' model which requires a PCP referral before paying for specialty care. Because the PCP is the first touch point, it is likely that she has the largest stock of knowledge about a patient along with considerable control over the patient's overall specialty care utilization. Second, PCP's generally serve as the center of the team coordinating care among each of a patient's providers (Agha et al. 2018). To put this in context, consider an anecdote chronicled by a general internist (Press 2014). A patient booked an appointment with his primary care PCP due to pain and fever. After tests revealed a tumor, the patient saw 11 clinicians in addition to his PCP over the course of 80 days. The PCP in this tale communicated repeatedly with each of these specialists, with the patient and with the patient's spouse. While the patient's care was fragmented across 12 providers, he received well coordinated care due to continuity with his PCP, who was able to maintain a full awareness of the patient's situation.

Overall, I find that a discontinuity in primary care leads to a 5% increase in physician-payments ($50) in the following year, and that this increase is nearly entirely driven by a 2.1 percentage point (6%) increase in the probability of a specialty care appointment. Because providers often return to their practice after a deployment, not every deployment leads to a documented change in primary care provider, though.[2] I use an instrumental variables approach and find a 30% cost increase ($308) and a 11 percentage point (31%) increase in the probability of a specialty care appointment for those that change PCP's after their provider deploys.

Perhaps equally interesting is what I don't find - that patients reduce their use of primary care after a discontinuity. This may be because the military operates under a restrictive HMO gate-keeping model, or perhaps because the practice and insurer assists the patient in finding a new provider. Previous work on continuity of care has found that patients take an extended period of time to find a new primary care provider after a discontinuity, and that many of these patients substitute through the use of specialty care or the emergency department (Sabety 2020). That I don't observe this substitution indicates that the increased use of specialists is not just a substitution effect, but an overall increase in total utilization. I don't find, however, that this increased use of specialists staves off any negative events. Neither emergency department utilization nor inpatient admissions change after a discontinuity.

In a secondary analysis, I consider what occurs within specialty encounters. I focus on existing patient - specialty care relationships and find that each encounter costs about 8% more after a discontinuity relative to visits before a discontinuity. This is driven by an increase in charges for decision-making, consistent with a model of increased coordination costs. I compare this with referrals that occur after a discontinuity, I find the opposite, these visits cost about 3% less per visit but also require 13% more visits - consistent with a Roy model where these referrals may be less appropriate and less well-targeted.

This research primarily contributes to the literature on variation in health care utilization. While much of the recent literature has focused on differentiating supply side from demand side drivers of variation (Finkelstein, Gentzkow, and Williams 2016 Molitor 2018), Agha et al (2019) focus on fragmentation and show that regional differences in fragmentation causally

[2] some patients remained enrolled on a deployed provider's, and see someone else in the practice until that provider returns

effect total utilization. I complement this work by showing that discontinuity in care has similar affects and is a potential area for managerial and public policy intervention.

This work also contributes to a literature on specialization and coordination costs. An important consideration is that the cost of transferring knowledge is not uniform. Jensen and Meckling (1992) describe knowledge as existing on a continuum of how costly it is to transfer. "General" knowledge is easily written down - for instance, a patient's blood pressure or the results of a lab test. "Specific" knowledge, in contrast, is the knowledge that is much more costly to transfer and could include the physician's understanding of a patient's personality or idiosyncratic tendencies. While this theoretical construct has driven a large management literature (Grant 1996; Huckman and Pisano 2006; Clark and Huckman 2012),there is little empirical economic work that is able to disentangle the concepts. My setting offers an opportunity to use discontinuity in care as a proxy for the loss of specific knowledge while maintaining generalizable knowledge. By empirically examining how an increase in coordination costs change providers' actions, I am able to provide an empirical test for Becker and Murphy (1992) who theorized that coordination costs would limit the extent of specialization.

Understanding the effects of discontinuity in care has important implications for both policy and practice. Relatively little policy focus has been placed on this interpersonal relationship (Guthrie et al. 2008) despite observed coordination failures. In fact policies have been largely associated with increasing discontinuity in the physician-patient relationship. For instance, the rise of managed care has been associated with annual contracts that may lead to forced discontinuities Flocke, Stange, and Zyzanski 1997. While the Centers for Medicare and Medicaid Services (CMS) has focused on care coordination as a potential cost-saving tool (McClellan et al. 2010), there has been significant patient churn in accountable care organizations (Hsu et al. 2017) and Medicaid policy often require patients to frequently change policies (Cutler, Wikler, and Basch 2012). Additionally the individual market under the Affordable Care Act (ACA) potentially leads to annual changes in primary care physicians. Narrow and changing networks often prevent an enrollee from developing a personal relationship with her physician (Buettgens, Nichols, and Dorn 2012). The US system of employer based health insurance also contributes to the frequency of discontinuities. There is an approximately 21% average annual private insurance cancellation rate, about one third of which is due to changes in employer group offerings (Cebul et al. 2008).

This paper proceeds as follows. In the next section I provide a conceptual framework for the analysis . In section 3 I detail the data and empirical specification. In section 4 I discuss the results and in section 5 I offer conclusions.

Conceptual Framework

The conceptual framework is based on the premise that patient conditions exist across a domain of uncertainty. For instance a broken bone may be easily observed and diagnosed from an x-ray, while shoulder pain could indicate a minor muscle strain or a much more serious heart attack. Primary care providers faced with this uncertainty must balance their limited time with the patient between gathering knowledge that can reduce the uncertainty about a patient's underlying condition and therapeutic work meant to ameliorate the problem.

The provider begins the appointment by seeking information. If she can achieve a sufficient level of certainty regarding the problem, she will transition to treatment. However, if after gathering information there is too much residual uncertainty, either about the condition or the appropriate treatment, she will refer to another provider with more specialized knowledge. Note that a 'sufficient level of certainty' will likely vary both across providers and across potential conditions. An important consideration, though, is that the patient's condition has both exogenous and endogenous components. That is, even with perfect treatment, each underlying condition has it's own probability distribution of outcomes. Knowledge only increases health through its effect on reducing uncertainty and increasing the probability of optimal treatment.

Knowledge, however, is dynamic. The primary care provider can not only access knowledge gathered during the current visit, but can also access a portion of knowledge previously gathered. The more generalizable that knowledge, the more likely the provider will have access to it regardless of a past interpersonal relationship with the patient. Generalizable knowledge is knowledge that is nearly costlessly transferred from one individual to another such as something that is easily automated in the electronic health record. An example would be the results of an x-ray or patient's blood pressure at a point in time. On the other hand, specific knowledge is that which is more costly to transfer such as personal information about the patient's preferences such as their propensity to seek care, pain tolerance, or levels of trust. This type of knowledge is generally only available within an existing relationship. This implies that the knowledge available to a provider is a function of both the total number of visits the patient has had within the health system and proportion of those visits that were with the current provider. After a discontinuity, this latter knowledge goes to zero, decreasing the probability that the provider can reach their sufficient level of certainty and increasing the probability of a specialty referral. Because specialty care is more expensive than primary care, this is likely to increase total costs while making the marginal referral to a specialist less appropriate.

Empirical Setting

The data for this study come from the Military Health System (MHS) Data Repository[3]. The MHS is an integrated health care system that provides care for active duty military, military retirees[4], and their family members. It is a dual system that combines care delivered in military-run hospitals (direct care) and care delivered by a local network of providers (purchased care). All care is paid for by the Tricare insurance benefit. Overall TRICARE

[3]The MHS is distinct from the Veteran's Administration health system which has different eligibility criteria.

[4]military retirees are those who have left military service after serving long enough for their pension to vest - typically 20 years

covers about 9 million individuals, but does not cover care delivered in a war zone.[5]

Tricare beneficiaries other than active duty service-members[6], predominantly choose between two health plans - Tricare Prime, a staff-model Health Maintenance Organization equivalent, and Tricare Select, a Preferred Provider Organization equivalent. Enrollees in Tricare Prime have near zero out of pocket costs while Tricare Select enrollees face a deductible and coinsurance. I focus my analysis on adult dependents of active duty service-members enrolled in Tricare Prime who receive primary care in a direct care clinic. Approximately 89% of military dependents are enrolled in Tricare Prime. Of note, beneficiaries that live within reasonable driving distance (e.g. 30 minutes) of a direct-care military clinic must enroll for primary care in that clinic unless it has reached capacity or otherwise provided an exemption. I focus on military dependents because I am able to observe their assigned PCP regardless of whether they seek care, tend to live near military posts, and do not have some of the idiosyncrasies of military service. I further limit the sample to those that are continuously enrolled in Tricare Prime and assigned within the continental United States. This excludes families that are stationed overseas as well as those that have a gap in coverage, for instance if the service member left the military and then returned to active duty.

Tricare Prime uses a gatekeeper model for primary care in which a Primary Care Provider (PCP) is responsible for coordinating a patient's care. Each PCP has a set panel of patients for whom that provider is responsible. When one of these providers permanently leaves a hospital or clinic, her patients are automatically reassigned to a new provider and a letter is mailed to the patient informing him of the change. PCP's include active duty, permanent civilian government employees, and contract civilian employees generally on renewable one-year contracts. Like civilian health care practices, both physicians and advanced practice providers (i.e. nurse practitioners and physician assistants) serve as PCP's within the direct care system. Also like civilian health care practices, patients don't necessarily see their PCP

[5]For a comprehensive review of the Military Health System, see *Evaluation of the TRICARE Program: Fiscal Year 2017 Report to Congress* 2017.

[6]active duty service members are legally required to enroll in tricare prime

during a primary care visit. Overall, about half of primary care visits are with a patient's assigned PCP.

The source of discontinuities in this study is through military PCP deployments. Military providers are generally not assigned to operational (combat) units so that they can practice medicine in hospitals when not needed in combat. This serves the dual purpose of maintaining their medical skills while also providing care to Tricare beneficiaries. The military services vary in how they select providers for deployments, however none base it on patient outcomes or other quality metrics. Take the Army for instance[7]. At least annually, the Army reviews operational needs and submits requirements to hospitals for a provider of a specified specialty. Individual hospitals have discretionary power for how they choose which provider will fill these assignments. These providers are then administratively aligned with a specific operational unit while continuing in their practice. Should that unit deploy to combat, the provider leaves the practice and accompanies the unit. An important note is that there is variation in when a provider learns of a pending deployment but that this tends to be bounded by one year prior to a deployment.

Outcome Variables

The primary dependent variable for this analysis is the probability of a specialty care appointment. The use of specialty care indicates that a primary care provider has decided that she has insufficient knowledge to effectively treat a patient and will coordinate rather than direct a patient's care.

I also consider total resource-based relative value units (RBRVU's) consumed by a patient. RBRVU's are a measure of how much resources were expended for a particular visit and are designed to standardize reimbursement for care across specialties. Tricare paid approximately $36 per RVU during the sample period. However caution should be taken in interpreting the RVU coefficient as a total dollar effect. RVU's are only intended to capture physician expenses and do not include facility charges (e.g. the payment to the hospital)

[7]The Army changed this system after the sample period

which are paid based on Medicare's inpatient and outpatient prospective payment systems. These prospective payment systems primarily apply for ED utilization and inpatient admissions.

Finally, I consider other measures of utilization including the probability of a primary care appointment, emergency department visit and inpatient admission.

Treatment Group Identification I assign patients to a treatment group based on whether they have had their PCP deploy. This is complicated by the fact that physicians may begin to transfer their patients prior to the quarter in which I observe them deploy overseas. I take a somewhat conservative approach and assign any patients who were enrolled with a deploying physician one year (four quarters) prior to the deployment to the treatment group. I use one year as this is the earliest a provider is likely to learn of an impending deployment based on military standard operating procedures. While I include the entire panel for all deploying providers as treated, I also include a separate identifier for whether a patient changes PCP's based on the enrollment file. This is because a portion of patients remain enrolled to their PCP during the deployment seeing another provider in the practice temporarily, and then resuming the relationship after the provider returns. Overall 53% of treated patients change primary care providers in the six months prior to an observed departure, or during the first six months of the deployment. I return to this in the instrumental variables approach later in the paper.

Sample Construction

The full sample consists of adult dependents of active duty service-members under age 65 and continuously enrolled in Tricare Prime within the United States between 2008-2017. I chose these dates because national provider identifiers (NPI) were first introduced in late 2006 and somewhat sparsely populated for the first quarter of 2007. I use provider NPI's to link primary care providers to their military records. I omit patients whose provider deploys for less than 200 days as these shorter deployments are less likely to result in a documented change in PCP enrollment.

Table 1 describes the full sample. The unbalanced full sample is not well suited to the two-way fixed effects difference in differences design for several reasons. As Goodman-Bacon (2018) points out, the staggered timing difference in difference estimator is a weighted average of all possible two by two estimators. When treatment effects vary over time, those who are treated early end up receiving higher weights and potentially biasing the results. Second, the traditional approach does not account for individual time trends uncorrelated with calendar time. This is particularly likely in health care where utilization can be cyclical based on incidences of illness or injury.

I address these concerns by creating a balanced analysis sample through a blocking and exact matching strategy. First, I block by the quarter-year time period. For each quarter, I limit the data to patients whose PCP deploys one year later, have been enrolled for at least one year with the same provider, and who remain enrolled in Tricare Prime for at least two more years (i.e. one year after the PCP deploys). I subsequently refer to this point in time as the "index quarter." I then match to the control group based on the length of relationship (in quarter-years) with their assigned PCP and assign an index quarter - relationship length identifier. I do not restrict to a specific number of matches per treated unit. I do, however, exclude any patients that are assigned in the index-quarter to a purchased care (civilian network) provider. I further limit the control group to patients that are never treated in the data. That is the control group at a point in time is *never* a treated unit in the sample. Finally, I restrict control group patients to one matching group. For instance if patient A is in the data as both a two year relationship in the first quarter of 2010 and a three year relationship in the first quarter of 2011, I randomize which one of these groups to keep. While not econometrically required, this simplifies the analysis. I relax this restriction in the appendix and show nearly identical results.

In the final step, I drop patients who have never had an appointment with their assigned PCP by one year after the index quarter, as by definition there is no relationship to discontinue. I return to this group as a falsification test later in the paper.

The final sample consists of 225,373 individuals over 2,028,357 patient-quarters. Table 2 shows descriptive statistics for the analysis sample in the year prior to the index-quarter. I measure health using the Charlson comorbidity index (Quan et al. 2005), a standardized score based on documented comborbidities. The sample is generally young and healthy and skews heavily female. While there are some statistically significant differences, these are all of low magnitude. For instance the treatment group is about 4 months older than the control group and about 4 percentage points more likely to be Caucasian[8]. They tend to use less than a half of an RVU fewer resources, but have the same probability of primary and specialty care in each quarter.

Empirical Specification

Identification Strategy

My primarily analysis uses a difference in differences approach that treats a physician deployment as a discrete event. Because the treatments occur at different time periods for different individuals I use a traditional two-way fixed effects model with a slight modification. I include a fixed effect that interacts time with the specific matched group identifier. In other words, I include a relationship-length - index-quarter - time fixed-effect. I estimate

$$Y_{ijt} = \alpha + \beta_1^* I(t \geq deployment) + \gamma^* X_{it} + \theta_i + \delta_{jt} + \epsilon_{ijt}$$

where y_{ijt} represents an outcome of interest for individual i, part of matched-group j at time t, and β_1 represents the effect of the PCM deployment. X_{it} are a vector of time-varying controls. Note that most standard patient-level controls are collinear with the fixed effects given that I've restricted the analysis to two years of data. I do include controls regarding the patient's assigned pcp in quarter t, including whether the PCP is a physician or advanced

[8] note that I only observe race for about 68% of patients

practice provider, and whether the PCP is active-duty military, along with how long an individual has been enrolled in their specific clinic. θ represents a vector of individual fixed-effects and δ_{jt} represent the quarter-year fixed-effects interacted with the relationship-group and the quarter on which the group was matched as described above. Standard errors are clustered by individual to account for potential serial correlation in the error term within an individual over time. Each regression includes the four quarters prior and the four quarters after a deployment, but omits the deployment-quarter.

Identifying Assumptions

The key identifying assumption for a difference in differences approach is that both the treatment and control groups would have followed parallel trends in the outcome variables if it were not for the discontinuity in care. For this assumption to hold, the timing of the deployment should not be correlated with the patient's health or physician's performance prior to the deployment. While this assumption is inherently untestable, we can provide suggestive evidence by looking at pre-trends. I use an event-study methodology to evaluate whether trends in the treatment and control groups were parallel prior to the physician deployment. The event study takes the form:

$$Y_{ijt} = \beta_{Q=t-t^*} + \theta_i + \delta_{jt} + \epsilon_{ijt}$$

Where Q is the quarter relative to the quarter of the physician deployment t^*. This allows me to consider the incremental changes over time before and after the deployment. I omit the time period $t^* - 4$ as that is the index-quarter on which I matched the groups and is the earliest plausible notification of a deployment. Figures **??** show the event studies for the probability of primary care and specialty care visits on the top, and the log of RVU's and log of total visits on the bottom. Event studies for other utilization outcomes are presented in the Appendix. Trends move in parallel until the quarter prior to the deployment and there is a noticeable increase after the provider deploys. The quarter-level estimates are somewhat noisy but paint a clear picture that both specialty care and RVU's increase after a discontinuity.

Two other assumptions must also hold in order to have unbiased results: Stable unit treatment value assumption (SUTVA), and common support. SUTVA means that there are no spillover effects between the treatment and control units. This assumption would not hold if a deploying provider effects the entire clinic such as through increasing workload for all other providers in the clinic. I address this by running a robustness check that includes clinic-quarter fixed effects. This adjusts for anything that effects the entire clinic population within a quarter.

The common support assumption means that there is covariate overlap between those who are and are not treated by a physician deployment. Unlike the previous two assumptions, common support is observable. Table 3 displays demographics and utilization in the year prior to the match for the treatment and control groups. In order to account for time trends I regress a dummy for the treatment group on each of the dependent variables and include a matched-group-quarter fixed effect. The observed balance is suggestive evidence of the quasi-random nature of physician deployment. Those who experience a discontinuity tend to be slightly older but are otherwise observably similar.

In order to interpret the coefficients as a result of discontinuity in care, though, deployments must satisfy three major assumptions - relevance, exclusion and conditional independence. The relevance assumption is simply that provider deployments are well correlated with discontinuity in care. Figure 2 shows an event-study of the probability of changing primary care providers in the time periods around a deployment. We see a slight uptick in the quarter prior to the deployment. This is not unexpected as providers likely schedule vacation time and spend time training immediately before a deployment.[9]

The exclusion restriction states that the deployment must only affect the outcome through

[9]the deployment file lists the date a provider arrives in an operational mission rather than the date they depart their practice

its impact on continuity of care. This would be violated if, for example, patients begin seeing providers of higher or lower average quality after a deployment. Table 4 shows the results of a series of t-tests between providers that never-deploy and providers that ever deploy in the data. I limit the data to all primary care appointments that include a patient and an index pcp, but do not restrict to a patient-pcp match. In other words, all of a providers encounters with any individual in the analysis sample are included. Each visit is coded using the Common Procedural Terminology (CPT) system, maintained by the American Medical Association. CPT codes are primarily used for billing and include up to three evaluation and management (E&M) codes as well as up to 10 non-E&M procedure codes. E&M codes are meant to denote the discussion and decision-making portion of the visit, while procedure codes are meant to capture anything else. I subset the visits into those that are only E&M related and those that only include a procedure. The remainder of visits include both. I also look for differences in mean CPT codes and RVU's per visit as measures of average resource intensity of a visit. While there are some statistically significant differences between deploying and non-deploying providers along these visit types, they are of very low magnitude. Deployers tend to perform slightly fewer E&M only visits and use about 0.02 additional RVU's or about eighty additional cents per visit.

A second threat to the exclusion restriction is if providers change their practice style prior to the deployment, such as if they are distracted and provide lower quality of care before the deployment. I look for evidence of this by running a series of difference in differences regressions on provider mean appointment values for the same measures described above, using the index quarter as the 'treatment'. The coefficient on the anticipatory period indicator depicts the effect of being within four quarters of deploying. None of the four measures are statistically significant, though I do get a slightly negative coefficient for log of RVU's. Note that I use a poisson regression for CPT codes as it is a discrete count variable bounded between one and thirteen.

Finally, the independence assumption states that there is no common cause of both deployments and discontinuities in care. In other words, there is no confounding variable causing both. The quasi-random nature of physician deployments gives strong suggestive evidence that the independence assumption is met.

Results

Table 6 shows the main difference in differences results for each of the dependent variables. Column 1 shows the effect of a physician deployment on the probability of a specialist visit in the post period. The coefficient indicates a 2.1 percentage point increase or about a 6% increase in probability on a base of 0.37. Column 2 shows the effects on primary care. The standard errors are large enough that I can't rule out no-effect, however positive coefficient is consistent with about a 2% increase in primary care utilization. This is particularly unexpected as previous research has shown that patients who lose their PCP tend to use less primary care (Sabety 2020), perhaps due to inattention or perhaps due to difficulty in finding a new PCP. One potential reason I do not observe this is because Tricare assists patients in finding a new primary care provider. Future research may consider whether a nudge such as recommending a new PCP could mitigate reduced primary care utilization following a discontinuity. Because visits and RVU's are highly skewed, I log transform them, adding one to handle the large number of zero values. Column 3 considers the log of all visits. The coefficient can be interpreted as about a 3% increase in any type of visit following a physician deployment. Finally column 4 indicates about a 5% increase in RVU's. At $36 per RVU this indicates about a $50 increase per patient in physician costs in the year following a discontinuity. In the appendix I show show an alternative poisson model as well as these same regressions for emergency department visits and inpatient admissions. Neither the probability of emergency department nor inpatient admissions are statistically distinguishable from zero despite fairly precise estimation.

While the linear probability model indicates an increase in the probability of a specialist appointment, this could indicate that patients are turning to established relationships after a discontinuity in care or that they are being referred to new physicians. Therefore I consider whether patients have *new* specialty referrals or further use of existing specialty care relationships. I consider a new referral as any first encounter in a clinic's specialty department in the data, based on the military's three digit department coding methodology or the purchased care network specialty codes. Table 7 column 1 shows the results of a regression where a visit to a new specialty is coded as one and all other quarters are coded as 0. The coefficient is slighly positive but the standard errors are too large to distinguish it from zero. Column 2 includes a poisson model on the number of total visits a new specialty referral requires. The 0.12 coefficient indicates that new referrals after a discontinuity require about 13% more visits than new referrals prior to a discontinuity. This could mean that while these patients may not be more (or less) likely to have a new referral, these new referrals after a discontinuity are less well-targeted. I will return to this topic when considering what goes on during a specialty encounter later in the paper.

Instrumental Variables

The difference in differences approach provides us with the treatment effect for patients whose PCP deploys, however not all patients have an observable change in their PCP even after a physician deploys. The MHS as an enterprise does not automatically transfer patients when a provider deploys, though individual providers may initiate a change. Some patients simply do not seek care and are less affected by the provider deployment. Other patients may come in and see whomever is available in the clinic. In this section I use the deployment as an instrumental variable in order to obtain the effects of the deployment specifically on those who change primary care providers. This local average treatment effect (LATE) scales the OLS estimates by the probability of changing providers.

Because patients may change providers shortly before or after a deployment, I use indicators for each time period around a deployment as an instrument for the actual discontinuity in care. I estimate the first stage as:

$$D_{ijt} = \beta_{Z=t-t^*} + \theta_i + \delta_{jt} + \epsilon_{ijt}$$

where D represents an individual in match-group j being in the first year of a relationship with a new provider in time t. Z is a vector of indicators for the 8 consecutive time periods relative to the deployment, beginning with the time period after the index quarter.

The first stage predicted values for $_{ijt}$ are then substituted into the following equation

$$Y_{it} = \beta\hat{D} + \theta_i + \delta_{jt} + \epsilon_{ijt}$$

Standard errors are adjusted for the two-stage least squares approach and clustered by individual. As an alternative specification, I replace the time indicators with a post-deployment indicator as used in the difference in differences estimation.

Table 8 shows these results for each of the main outcomes of interest. Column one is the preferred specification and shows the results of the elapsed time instrument. While primary care estimate is imprecise and approximately the same magnitude as the difference in difference estimate, we see about a much more substantial 30% relative increase (11 pp) in the probability of using any specialty care after a primary care provider deploys. Both overall RVU's and over visits increase substantially as well. The 0.267 coefficient on the log of RVU's indicates a slightly more than 30% increase on a base of 27.92 RVU's per year. Given approximately \$36 per RVU, this equates to about a \$308 increase over the year following a discontinuity. Note that this is a lower bound as RVU's do not capture facility costs paid for through prospective payments systems.

Impact on Specialization

In this section, I consider what happens within specialty care visits after a discontinuity in primary care. As the PCP is the primary coordinator across specialty care, a PCP discontinuity is likely to increase coordination costs. Becker and Murphy (1992) suggested that coordination costs are the limiting factor on specialization. Conceptually, an increase in coordination costs will require the specialist to either extend more effort or accept lower outcomes. I assume that specialists have a minimal acceptable outcome and will exert more effort as required to meet that outcome up to some maximum level.

I test this hypothesis but considering the timing of the first referral to a specialty clinic. I restrict my analysis to specialty care within the MHS direct care system and define a specialty clinic using the MHS 3-digit cost accounting codes. For example, "BAC" is the code for cardiology while "BAO" is rheumatology. I then combine these codes with individual medical facility codes so that cardiology at Brooke Army Medical Center on Fort Sam Houston is a separate specialty clinic from cardiology at Darnall Army Medical Center on Fort Hood.

I focus on four process measures that I can observe in the data. First, I consider the log of RVU's for each visit as a measure of intensity. Second I analyze whether a patient has an Evaluation and Management only appointment. As mentioned previously, E&M codes indicated decision-making. Third, I consider the proportion of procedure visits. These are visits that do not include any decision making. Finally, I consider the number of procedures coded per visit conditional on having any procedures.

Table 9 shows the results of a series regressions for each of these measures for patients that are newly referred after a physician deployment. New visits for the treatment group tend to be about 3% less intensive than those for the control group, driven by fewer procedures per visit - consistent with a roy model where these new referrals are slightly less appropriate after a discontinuity.

Table 10 shows difference in differences estimates for patients who were referred prior to the index quarter. These are patients whose referral could not be caused by the deployment and for whom the pcp was likely coordinating care before deploying. Each regression includes an individual - specialty clinic fixed-effect as well as the group-time fixed-effect described earlier. I cluster standard errors by the individual-specialty clinic. Column one shows the change in log of RVU's. After a PCP deployment, these visits require about 10% more resources per visit. Each visit costs approximately $105 so this is just over $10.50 per visit increase relative to comparable encounters for the control group. Column two and three indicate that patients may be more likely to have a visit that only includes provider decision-making and are less likely to have a visit that does not involve any decision-making. Note that the E&M only visit is not statistically distinguishable from zero but has a positive coefficient consistent with the other estimates. In column 4 I measure the number of coded procedures for each visit conditional on not being an E&M only visit. The increase in procedures is roughly equivalent to the increase in RVU's. Note that, as earlier, this estimate comes from a poisson regression due to the nature of the variable.

Falsification Test and Robustness Checks

I run two falsification tests and a series of robustness checks to provide supporting evidence that discontinuity in care is in fact driving the results rather than some unobserved effect of the deployment or statistical noise. First, I use patients that have never met their primary care provider but whose physician deploys. This could happen because these individuals rarely use primary care or if they see happen to see other providers in the office when they seek care. However, a primary care physician deployment should not effect a patient who has never met his or her provider to begin with. Column 1 of Table 11 shows that there are no effects among this group. Second, I use very short deployments. These are operational assignments where the provider returns in under a month. From the patient's standpoint this is unlikely to create a discontinuity. I show in column 2 that these also have null results.

Table 12 show the results of four robustness checks. In column 1, I add in a clinic-quarter fixed effect to the main specification. This addresses the concern that a loss of a pcp to a deployment limits the overall capacity of the clinic. The magnitude of the effects are somewhat reduced, though not significantly different from the main results. In column two, I eliminate anyone who switches to a purchased care (civilian network) primary care provider. In column 3 I restrict the sample to those that never change clinics. This excludes any normal military moves for the patients. I lose some precision but have similar results. Finally, in column four I restrict the analysis to those who have physicians rather than advanced practice providers in the pre and post period. This is somewhat restrictive as changes between provider types may not be random. Despite differences in sample size and precision, the four robustness checks offer evidence that the results are stable.

Discussion and Conclusions

A story emerges from the data that for most healthy individuals, discontinuity in primary care does not have much effect. However, for individuals that need primary care, continuity becomes important with a substantial 30% increase in costly specialty care utilization despite no reduction in emergency department utilization or inpatient admissions. Within specialty care, individual encounters become more intense, using about 7% more resources on average. These effects are particularly relevant given that the the military uses an HMO staffing model similar to Kaiser-Permanente and Geisinger (*Final Report to the Secretary of Defense: Military Health System Review* 2014). This model has been upheld as the gold standard in integrated care (Curry and Ham 2010). Yet this study indicates that even in this setting there can be a lack of care coordination with provider turnover potentially leading to increased costs.

That I find these effects in a working age, relatively healthy population is particularly important when considering health policy. Recent work has highlighted how common turnover is in health care. The employment-based health insurance and individual marketplaces, with annual contracts, likely increase the amount of turnover. It's unclear if the benefits of competition outweigh these costs. The setting of this research does not allow for that analysis, but future researchers should quantify this tradeoff.

Additionally, many recent policies and organizational innovations have been focused on reducing coordination costs. This is a substantial tradeoff, however, in that policies focused on decreasing loss of information likely also decrease the absolute level of knowledge available within a visit. For instance, Electronic Health Records are designed to expand available knowledge, yet may increase provider burnout and take the provider's attention away from the patient resulting in less accumulation of new, specific knowledge. Another example is the patient centred medical home (PCMH) model. This complex model is designed to promote sharing of information about a patient across a group of providers. However, the patient's relationship with an individual provider may be weakened. This could conceivably mitigate the effects of discontinuity by reducing continuity for every one, something unlikely to lead to better outcomes over all.

While this work has quantified the impact of discontinuity in care, there is an open question of how to address the problem. As I stated at the outset, organizations exist to coordinate tasks across individuals with different knowledge. However, individual workers have goals, organizations do not (Cyert and March 1963). Creating incentives for these individuals to coordinate care within and across health care organizations has largely been left to government policy rather than managerial initiative. Yet regulated payment systems have created inverse incentives. For instance, hospitals are largely paid on a case-rate basis where the incentive is to minimize costs while physicians are paid on a fee for service basis where the incentive is to maximize revenue. That physicians and hospitals must work together magnifies the principal-agent problem. One way of addressing this problem is for government policy to incentive well-coordinated care while shifting the task of coordination to managers. This makes sense in a model in which managers specialize in coordinating tasks across individuals with different knowledge. The manager, of course, may be the primary-care provider or a non-clinician administrator. Accountable Care Organizations and bundled payment initiatives may offer the most promise for this as they align incentives for well-coordinated care across disparate individuals. Early research on the accountable care organizational model shows that it may actually increase physician turnover and patient churn (Hsu et al. 2017) yet policies focused on continuity of care may be able to address these shortfalls while shifting the problem of coordination to those best able to solve it.

References

Agha, Leila, Brigham Frandsen, and James B Rebitzer (2019). "Fragmented division of labor and healthcare costs: Evidence from moves across regions". In: *Journal of Public Economics* 169, pp. 144–159.

Agha, Leila et al. (2018). *Team formation and performance: evidence from healthcare referral networks*. Tech. rep. National Bureau of Economic Research.

Allard, Marie, Pierre Thomas Léger, and Lise Rochaix (2009). "Provider competition in a dynamic setting". In: *Journal of Economics & Management Strategy* 18.2, pp. 457–486.

Arora, Vineet M et al. (2010). "Problems after discharge and understanding of communication with their primary care physicians among hospitalized seniors: a mixed methods study". In: *Journal of Hospital Medicine* 5.7, pp. 385–391.

Becker, Gary S and Kevin M Murphy (1992). "The division of labor, coordination costs, and knowledge". In: *The Quarterly Journal of Economics* 107.4, pp. 1137–1160.

Bond, Amelia M and Stephen D Schwab (2019). "Utilization Variation In Military Versus Civilian Care: Evidence From TRICARE". In: *Health Affairs* 38.8, pp. 1327–1334.

Buettgens, Matthew, Austin Nichols, and Stan Dorn (2012). "Churning under the ACA and state policy options for mitigation". In: *Prepared for Robert Wood Johnson Foundation, Timely Analysis of Immediate Health Policy Issues, http://www. urban. org/UploadedPDF/412587-Churning-Under-the-ACA-and-State-Policy-Options-for-Mitigation. pdf*.

Cebul, Randall D et al. (2008). "Organizational fragmentation and care quality in the US healthcare system". In: *The Journal of Economic Perspectives* 22.4, pp. 93–113.

Clark, Jonathan R and Robert S Huckman (2012). "Broadening focus: Spillovers, complementarities, and specialization in the hospital industry". In: *Management Science* 58.4, pp. 708–722.

Curry, Natasha and Chris Ham (2010). "Clinical and service integration". In: *The route to improve outcomes. London: The Kings Fund*.

Cutler, David, Elizabeth Wikler, and Peter Basch (2012). "Reducing administrative costs and improving the health care system". In: *New England Journal of Medicine* 367.20, pp. 1875–1878.

Cyert, Richard M, James G March, et al. (1963). "A behavioral theory of the firm". In: *Englewood Cliffs, NJ* 2.

Evaluation of the TRICARE Program: Fiscal Year 2017 Report to Congress (2017). Tech. rep. Department of Defense.

Final Report to the Secretary of Defense: Military Health System Review (2014). Tech. rep.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams (2016). "Sources of geographic variation in health care: Evidence from patient migration". In: *The quarterly journal of economics* 131.4, pp. 1681–1726.

Flocke, Susan A, Kurt C Stange, and Stephen J Zyzanski (1997). "The impact of insurance type and forced discontinuity on the delivery of primary care". In: *Journal of Family Practice* 45.2, pp. 129–136.

Frakes, Michael and Jonathan Gruber (2019). "Defensive medicine: Evidence from military immunity". In: *American Economic Journal: Economic Policy* 11.3, pp. 197–231.

Goodman-Bacon, Andrew (2018). *Difference-in-differences with variation in treatment timing*. Tech. rep. National Bureau of Economic Research.

Grant, Robert M (1996). "Toward a knowledge-based theory of the firm". In: *Strategic management journal* 17.S2, pp. 109–122.

Guthrie, Bruce et al. (2008). "Continuity of care matters". In: *BMJ: British Medical Journal (Online)* 337.

Hayek, Friedrich August (1945). "The use of knowledge in society". In: *The American economic review*, pp. 519–530.

Hsu, John et al. (2017). “Substantial Physician Turnover And Beneficiary âChurnâIn A Large Medicare Pioneer ACO”. In: *Health Affairs* 36.4, pp. 640–648.
Huckman, Robert S and Gary P Pisano (2006). “The firm specificity of individual performance: Evidence from cardiac surgery”. In: *Management Science* 52.4, pp. 473–488.
Hutter, Paul J et al. (2019). “Readiness Of Medical Providers In The Military Health System: Overview Of Operational And Policy Considerations”. In: *Health Affairs* 38.8, pp. 1274–1280.
Jensen, Michael C and William H Meckling (1992). “Specific and general knowledge and organizational structure”. In:
McClellan, Mark et al. (2010). “A national strategy to put accountable care into practice”. In: *Health Affairs* 29.5, pp. 982–990.
Molitor, David (2018). “The evolution of physician practice styles: evidence from cardiologist migration”. In: *American Economic Journal: Economic Policy* 10.1, pp. 326–56.
Press, Matthew J (2014). “Instant replayâa quarterback’s view of care coordination”. In: *New England Journal of Medicine* 371.6, pp. 489–491.
Quan, Hude et al. (2005). “Coding algorithms for defining comorbidities in ICD-9-CM and ICD-10 administrative data”. In: *Medical care*, pp. 1130–1139.
Sabety, Adrienne (2020). *Association of Primary Care Physician Turnover with Utilization and Outcomes Among Medicare Patients: Observational Study*. Tech. rep. Working Paper.
Shrank, William H, Teresa L Rogstad, and Natasha Parekh (2019). “Waste in the US health care system: estimated costs and potential for savings”. In: *Jama* 322.15, pp. 1501–1509.

Tables

Table 1: Full Sample

	mean	sd
gender (female)	0.90	0.30
Caucasian	0.57	0.50
Any Chronic Illness	0.09	0.29
age	30.10	8.07
rvu’s	5.68	12.00
Any specialist visit	0.32	0.47
Any primary care visit	0.34	0.47
Any emergency department visit	0.11	0.32
Any inpatient admission	0.03	0.18
N	16,586,762	

Notes: Description of full sample of patients as described in the text.

Table 2: Analysis Sample

	mean	sd
gender (female)	0.93	0.26
Caucasian	0.56	0.50
Any Chronic Illness	0.13	0.33
age	30.49	7.90
rvu's	7.21	13.47
Any specialist visit	0.37	0.48
Any primary care visit	0.43	0.49
Any emergency department visit	0.13	0.34
Any inpatient admission	0.04	0.19
N	2,929,849	

Notes: Description of analysis sample including all time-periods from one year prior to index-quarter though two years post index-quarter

Table 3: Comparison of Patients in the analysis sample Who Do and Do Not Undergo A Deployment Related Discontinuity in Primary Care

	Control	Treatment	T Stat
gender (female)	0.93	.93	2.87
Caucasian	0.56	0.60	7.28
Any Chronic Illness	0.12	0.13	2.62
age	29.36	29.69	4.84
rvu	7.95	7.51	-3.61
Any specialist visit	0.40	0.40	0.18
Any primary care visit	0.49	0.48	-1.63
Any emergency department visit	0.16	0.16	2.44
Any inpatient admission	0.04	0.04	-0.89
N	888,724	12,768	

Notes: Observations are at the quarter-year level. Data includes the four quarters prior to the index-quarter. Difference in means comes from a regression that includes group-time interacted fixed effects as described in the text. Caucasian regression includes 620,072 observations due to missing data

Table 4: Comparison of Providers Who Ever Deploy

	Non-Deployer	Ever-Deployer	T Stat
Evaluation & Management only	0.82	0.81	-3.18
Procedure Visit	0.05	0.05	-0.28
Coded Procedures (CPT) per visit	1.19	1.19	-2.03
RVU's per Visit	2.57	2.59	4.96
N	88,436	25,402	

Notes: Observations are at the provider-quarter level. Analysis includes mean quarterly estimates for all visits for primary care providers with patients in the analytical sample. Evaluation and Management (E&M) visits are encounters that only include E&M codes. Procedure visits are visits that do not include any E&M code. Coded Procedures calculated for non E&M visits only.

Table 5: Primary Care Physician Anticipatory Effects

	(1) Log RVU	(2) E&M Only	(3) CPT per visit	(4) Log Visits
Anticipatory Period	-0.007	0.009	-0.002	-0.004
	(0.008)	(0.008)	(0.009)	(0.005)
Controls	No	No	No	No
Fixed Effects	Yes	Yes	Yes	Yes
N	55,151	55,151	55,151	55,151

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Observations are at the quarter-year level. The coefficient is the interaction of an indicator for the quarter being greater than or equal to the index-quarter and a deployment indicator. Evaluation and Management (E&M) visits are encounters that only include E&M codes. Coded Procedures calculated for non E&M visits only. All regressions include provider and quarter fixed effects but no other controls.

Table 6: Main Results

	(1) PC Visits	(2) Spec Visits	(3) Log RVU	(4) Log Visits
DiD Coefficient	0.011*	0.020 ***	0.047***	0.035 ***
	(0.006)	(0.007)	(0.016)	(0.012)
Sample Mean	0.41	0.36	6.98	2.58
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
r2	0.27	0.31	0.42	0.45
N	2,025,165	2,025,165	2,025,165	2,025,165

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Results of OLS DiD regressions. Controls include how long the patient has been enrolled in the same clinic, whether these physician is active duty military, and whether the pcp is a physician or advanced practice provider. All regressions include individual and group-quarter fixed effects. Standard Errors are clustered by patient

Table 7: New Specialty Visits

	(1) Probability of a new visit	(2) Total Required Visits
DiD Coefficient	0.007	0.124**
	(0.005)	(0.053)
Sample Mean	0.18	3.67
N	2,025,165	314,253

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Results of DiD regressions. All regressions include individual and group-quarter fixed effects. Column one is a linear probability model indicating whether a patient was seen in a new specialty clinic based on a grouping of clinic and three digit specialty code as described in the text. Column two is a poisson model of the total number of visits required for new visits after a discontinuity in primary care. Only visits within the direct care system with a physician or advanced practice provider are considered. Standard Errors are clustered by patient

Table 8: IV Regression Results

	Elapsed Time	Post Indicator
Primary Care	0.009	0.119*
	(0.043)	(0.067)
Specialty Care	0.109**	0.220***
	(0.046)	(0.075)
Log RVU's	0.267**	0.516***
	(0.110)	(0.183)
Log Visits	0.160*	0.380**
	(0.078)	(0.131)
Controls	Yes	Yes
Fixed Effects	Yes	Yes
Kleibergen-Paap Wald 1st Stage F Stat	64.5	93.7
N	2,253,730	2,247,346

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Results of two stage least square regressions as described in the text. First stage regressions include all controls and fixed effects. Column one uses a series of elapsed time dummies. Column two uses a post indicator that is one starting in the quarter after a provider deploys and omitting the deployment quarter. Dependent Variable is an indicator for being in the first year of a patient-provider relationship.

Table 9: Comparison of Specialty Encounter Intensity for Patients Who Do or Do Not Have a PCP deploy

	(1)	(2)	(3)	(4)
	Log RVU	E&M Only	Procedure Visits	Coded Procedures per Visit
Treated Group	-0.035***	0.012	0.009	-0.032**
	(0.013)	(0.010)	(0.013)	(0.014)
Controls	No	No	No	No
Fixed Effects	Yes	Yes	Yes	Yes
N	197,782	197,782	197,782	145,164

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Observations are at the specialty encounter level for all patients in the analytical sample being seen for specialty care in the military direct care system on an outpatient basis. Evaluation and Management (E&M) visits are encounters that only include E&M codes. Procedure visits are visits that do not include any E&M code. Coded Procedures calculated for non E&M visits only

Table 10: Results of specialty encounter analysis

	(1) Log RVU's	(2) E&M Only Visits	(3) Procedure Visits	(4) Coded Procedures per Visit
DiD Coefficient	0.098***	0.034	-0.084**	0.107**
	(0.037)	(0.037)	(0.041)	(0.043)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
N	261,723	261,723	261,723	198,028

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Results of regressions on intensity of care metrics. All regressions include individual and group-quarter fixed effects. Columns one through three are OLS. Column four is poisson regression. Observations are at the specialty encounter level for all patients in the analytical sample being seen for specialty care in the military direct care system on an outpatient basis. Evaluation and Management (E&M) visits are encounters that only include E&M codes. Procedure visits are visits that do not include any E&M code. Coded Procedures calculated for non E&M visits only. Standard Errors are clustered by patient

Table 11: Results of Falsification Tests

	(1) No Relationship	(2) Short Deployments
Specialty Care	-0.009	-0.008
	(0.006)	(0.014)
Log of RVU's	-0.003	0.027
	(0.015)	(0.035)
N	1,344,917	2,004,481

* $p<0.1$, ** $p<0.05$, *** $p<0.01$, **** $p<0.001$

Notes: Results of main specification applied to samples in which we would not expect a result. Column one only includes patients that had not met their primary care provider by one year after the index-quarter. Column two shows the results for deployments that are less than 30 days in duration

Table 12: Robustness Checks

	(1) Clinic-Quarter	(2) Military Only	(3) Never Movers	(4) Physician Only
Specialty Care	0.016**	0.022***	0.013	0.018
	(0.007)	(0.007)	(0.008)	(0.011)
Log of RVU's	0.030*	0.037**	0.037*	0.037
	(0.017)	(0.017)	(0.022)	(0.027)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
N	2,024,410	1,752,433	1,258,218	469,822

* p<0.1, ** p<0.05, *** p<0.01 , **** p<0.001

Notes: Results of OLS DiD regressions with conditions. Column 1 includes a clinic-quarter fixed effect. Column 2 omits patients who are ever managed outside the military's direct care system. Column 3 omits patients who change clinic enrollment during the sample period. Column 4 omits patients managed by an advanced practice provider

1 Figures

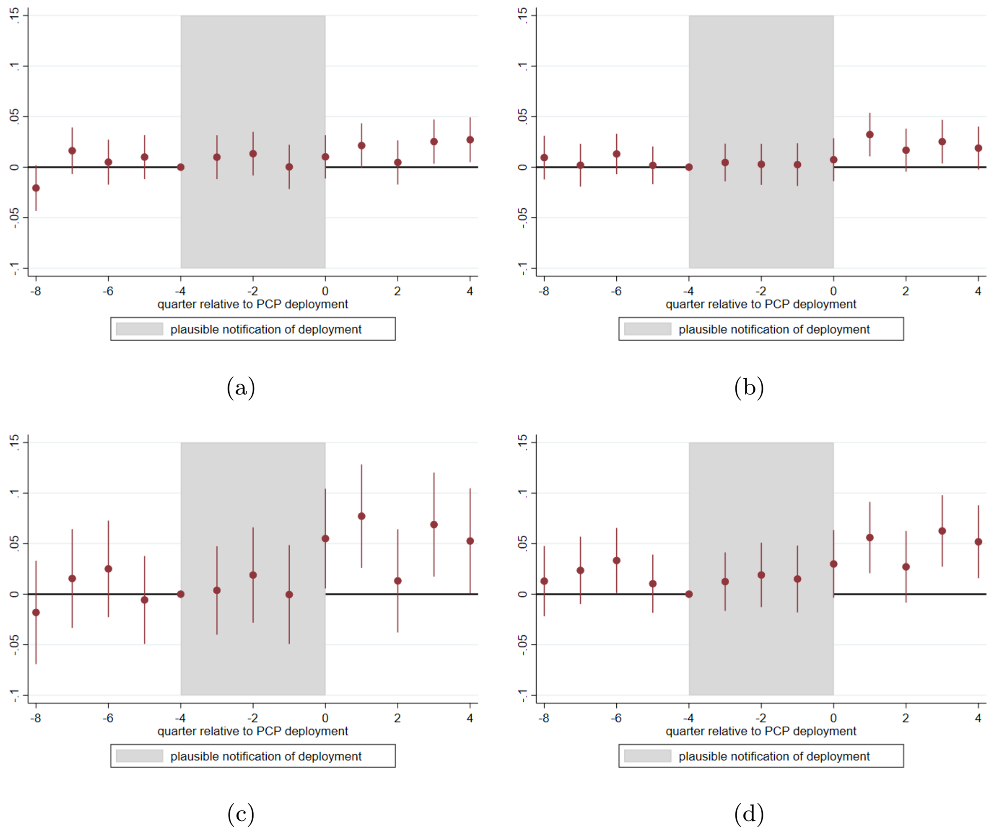

(a) (b)

(c) (d)

Figure 1: Author's portrayal of change in utilization by patient-quarter relative to pcp deployment. Dots indicate point estimates. Lines indicate 95% confidence intervals. T-4 is the omitted period. All regressions as described in the text. Panel (a) depicts the change in the probability of a primary care visit. Panel (b) depicts the probability of a specialty care visit. Panel (c) depicts changes in the log of total RVU's. Panel (d) depicts the change in the log of total visits.

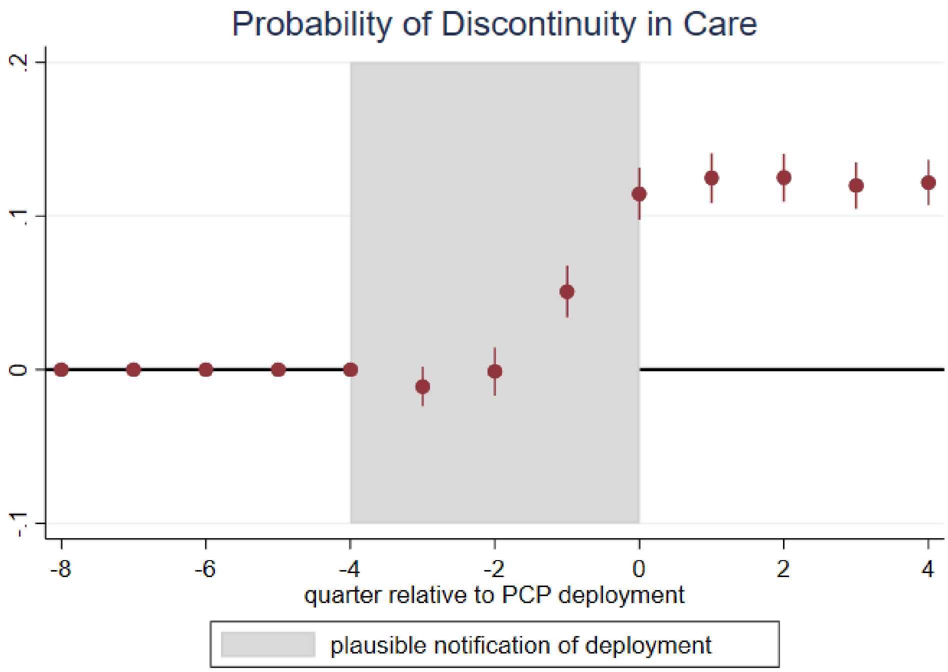

Notes: Graphical portrayal of the probability of change in PCP enrollment on time relative to deployment. X axis is quarter-years relative to provider deployment. Y access is probability that the PCP is different than in the index quarter. Dots are point estimates. Vertical lines are 95% confidence intervals. Grey box is the notification period. Regression includes person and relationship-group-quarter-year fixed effects as descibed in the text. Standard Errors are clustered by individual.

Figure 2: Effect of Discontinuity on Specialty Care Utilization