The Value of Specific Knowledge: Evidence from Disruptions to the Patient- Physician Relationship

Stephen D Schwab *

April 7, 2023

Abstract

When a member of a work-team leaves, some knowledge is lost to the organization. Exploiting quasi-random turnover among military physicians due to deployments, I estimate the effects of turnover on patients and other providers in the same care-team. I find that a discontinuity in primary-care leads to a 3-5% increase in costs driven primarily by an increase in the use and intensity of specialty care, especially in areas with high levels of uncertainty. This has implications for how organizations allocate tasks and manage turnover.

*Schwab: University of Texas at San Antonio stephen.schwab@utsa.edu. I am thankful to Lawton Burns, Guy David, Abby Alpert, Rahul Kapoor, Ingrid Nembhardt, Ashley Swanson, David Meltzer, Stuart Craig, Amelia Bond, and Karen Zhang and seminar participants at ASHEcon, OT-HC, Wharton, UNC, Vanderbilt, Baylor, Temple, UTSA, Texas State, and USUHS for helpful thoughts and suggestions. I acknowledge data support from the Army Analytics Group Person-Event Data Environment.

Organizational economics posits that organizations exist largely to coordinate tasks across individuals with different knowledge. This knowledge encompasses differences in individual skills, experiences and education, as well as idiosyncratic information regarding “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. This is further complicated by the fact that knowledge varies in its transferability. While general knowledge can be easily distributed across an organization, more specific knowledge exists only within individual agents. How this knowledge gets acquired and transferred is a function of communication and trust between individuals. If agents leave or are replaced, some knowledge is lost to the organization. In this paper I use quasi-random turnover of primary-care providers (PCP) within the Military Health System to estimate the causal impact of turnover on patients and other clinicians within the same care team.

I analyze this in the medical context as health care is a prime example both of the fragmented nature of knowledge and the associated coordination costs in a context in which trust and communication are key inputs. Care is often fragmented across multiple clinicians (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)). Health care is also one of several industries that are particularly reliant on trust and communication as products and services are tailored to the individual patient. After the loss of a PCP, some knowledge about the patient's particular case may be lost and/or a patient may not communicate as much with their new PCP due to lack of trust. Both mechanisms mean that the PCP has less specific knowledge regrading the patient, potentially decreasing the probability that the new provider can diagnose a condition and increasing the probability of a referral to specialty care. Alternatively, patients may "demand" more referrals due to the loss of trust in their provider. Because specialty care is more expensive than primary care, this is likely to increase total costs.

A particular challenge of estimating the effects of turnover is that employment changes are rarely exogenous to the organization. Clinicians and other workers who move or retire may differ in unobservable ways from those that remain. Furthermore, because knowledge is non-rival, workers may take information (and patients) with them to a competitor organization. I address 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 active duty military health care providers (Hutter et al. (2019)). As military operational missions arise, these providers are pulled from their practices and deployed outside the United States. I use wartime deployments as a source of exogenous variation in PCP turnover in order to estimate the effects of turnover on patients. An additional challenge is isolating the impact of an individual relationship (and the trust and knowledge embedded within it) from other effects of turnover. For instance, the loss of a physician from a medical practice may also result in the patient moving to a different practice. In the military setting, patients affected by the discontinuity generally remain enrolled in the same practice with the same electronic health record, and the same insurance design.

I begin my analysis by investigating overall care utilization after the loss of a provider. In this setting, the primary care provider is the center of the clinical team and is overall responsible for knowing the patient and coordinating their care. To put this in context, consider two andecdotes. The first one was chronicled by a general internist (Press (2014))[1]. A patient booked an appointment with his 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 likely due to continuity with his PCP, who was able to maintain a full awareness of the patient's situation.

The second anecdote was relayed by a physician on twitter[2] under a pseudonym. He had a long-time patient come in who appeared to be suffering from depression and regressing in other healthy behaviors such as quitting smoking. The physician knew the patient was fond of his dog and in normal rapport asked the patient about the dog. The doctor learned through this communication that dog had recently passed away. Rather than prescribing medication or referring to specialty care, the physician "prescribed" adopting another dog. This resolved the patient's symptoms.

Overall, I find that a primary care provider deployment leads to a 3-5% long-term increase in a patient's total utilization of physician services. I find, though, that this average includes substantial heterogeneity. Many healthy patients aren't affected at all, while patients who need their care coordinated the most have increases of 20% or more. Much of the increase is driven by a 1.7 percentage point (5%) increase in the probability of a referral to a more specialized clinician, especially in services where there is more uncertainty, and this effect is driven by new referrals rather than patients returning to specialists where there is an existing relationship. This drives just under half of the increase in utilization, with the remainder split between primary care and the emergency department. I also find that specialist visits become more costly after the loss of a PCP, especially for patients with multiple comorbidities.

This study adds to the body of research on turnover by isolating the effects of the knowledge and

[1] Agha et al (2019) also relay this anecdote

[2] As of this writing it appears the author has deleted his twitter account and the thread is no longer available

trust embodied in the doctor-patient relationship from the other correlated effects of physician exits. By considering the effect of these disruptions on other members of the team, I am also able to show that removing the primary coordinator in a team reduces the productivity of more specialized workers. This provides empirical support to Becker and Murphy's (1992) theoretical model of coordination costs and suggests that additional investment in generalized workers may also lead to increased productivity from more specialized workers. My setting offers an opportunity to isolate the effects of turnover at the individual level while separating production loss from accounting costs. Understanding this cost contributes to our understanding of the impact of organizational turnover, especially in a knowledge-based organization.

This paper proceeds as follows. In the next section I provide background on physician exits, and the role of knowledge, and trust. In section 3 I detail the data and empirical specification. In section 4 I discuss the results. Section 5 concludes.

1 Background

Physician Exits

This research primarily contributes to the growing literature on physician turnover and *discontinuity* in care. Previous studies have shown that discontinuity in care increases health care costs (Sabety, Jena and Barnett (2021); Kwok (2019); Simonsen et al. (2021); Bischof and Kaiser (2021)). Both Sabety (2021) and Staiger (2021) found that patients reduce their use of primary care in the year after a physician exit, resulting in an increase in much more costly inpatient admissions. Kwok (2019) considered whether finding a better patient-physician match could be beneficial and found that switching costs when patients change primary care providers dominates the spending effects driven by differences in physician practice styles. In contrast, Simonsen et al. (2021) used Danish practice closures and found that physician changes led to new treatment for some conditions that were potentially missed by the previous provider. Bischof and Kaiser (2021) used Swiss data to also analyze practice closures and found that patients were more likely to substitute to specialist physicians than find a new Primary Care Provider.

I add to this literature by isolating the effect of the relationship itself, and the knowledge, trust, and communication embedded within it, from other, correlated effects of discontinuity in care. In the Medicare setting used by Sabety (2021) and Kwok (2019), the Medicaid setting studied by Staiger (2021) as well as the Danish and Swiss data used by Simonsen et al. (2021) and Bischof and Kaiser (2021), physician exits may be particularly pernicious as patients will have to find a new practice. This may be particularly difficult for the Medicaid enrollees studied by Staiger. In the military setting, though, practice changes are uncorrelated with physician exits. This provides an opportunity to isolate the effect of the relationship itself from other factors associated with discontinuity in care. An additional contribution is that the military may more closely resemble the US employer-based insurance population than other government payers given the HMO model which mandates gate-keeping (i.e. no specialist visits without a primary care referral) and the younger (and healthier) population.

Information and Trust

Considerable work on continuity of care has focused on the role of trust in the Doctor-Patient relationship (Mainous et al. (2001); Saultz (2003); Fritz and Holton (2019)). The academic literature on trust has many definitions, but I follow Mayer, Davis and Schoorman (1995) and define trust as "the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party." In the health care context, the patient is particularly vulnerable to the physician's decisions (which are often legally codified) and, due to significant information asymmetry, has little opportunity to monitor or control the physician. Ford (2004) analyzes the connection between trust and knowledge management and asserts that interpersonal trust is a key input into knowledge acquisition and generation. Conceptually, this is because there is risk in sharing private information (Ford (2004)). This is quantitatively supported in a study by Muthusamy and White (2005) and qualitatively supported by Connelly and Kelloway (2003). Ford also points out that causality can go in either direction as sharing information may also lead to increasing levels of trust. This has implications for the doctor-patient relationship as more contact may lead to more trust, which in turn may lead to more knowledge-sharing. In later work Ford and Staples (2010) differentiate full and partial knowledge sharing with partial knowledge sharing involving some type of protective behavior, including protecting oneself but also protecting the recipient's time or feelings. This could impact patient care if the unshared portion of knowledge is critical to diagnosis or coordination of care. In the healthcare context, Agha et al. (2018) found that interpersonal relationships between primary care providers and specialists, which conceptually involve both trust and knowledge transfer, lead to lower healthcare utilization without sacrificing quality.

Complementing the work on trust, is a literature on knowledge sharing more broadly. Transferring knowledge across an organization can be costly (Garicano (2000)), even more so when the optimal solution to a problem is unknown (or unknowable) by any member of the organization (Chan (2021)). Knowledge, however, varies in its cost to transmit (Jensen and Meckling (1992)). Some“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 much more costly to transfer and could include the physician’s understanding of a patient’s personality such as whether the patient is likely to adhere to a daily medicine and the most effective forms of communication with that patient. The management literature has linked variation in specific knowledge to differences in firm-specific (Huckman and Pisano (2006)) and customer-specific (Clark, Huckman and Staats (2013)) performance. Frakes and Wasserman (2021) provider further evidence on knowledge transfers “spillovers” within teams. I add to this literature by estimating the value of specific knowledge in a relationship where trust is a key input and disentangling it from other disruptions associated with turnover.

2 Military Setting

I conduct this study in the context of the Military Health System (MHS). The MHS is an integrated health care system that provides care for active duty military, military retirees[3], and their family members. It is a dual system that combines care delivered in military-run clinics (direct care) and care delivered by a local network of private sector providers (purchased care). All care is paid for by the TRICARE insurance benefit. Overall TRICARE covers about 9 million beneficiaries..[4] I focus my analysis on adult dependents of active-duty military. These patients tend to live near military clinics, but do not have some of the idiosyncrasies of military service such as mandatory physical training.

Primary care is an ideal setting to study information loss for two reasons associated with the nature of knowledge. First, as health care costs have risen, health insurers have increasingly relied on PCP’s judgement in order to limit utilization, requiring a PCP referral before authorizing payment for specialty care. Tricare Prime follows this “gatekeeper” model for primary care. Second, the PCP is responsible for coordinating a patient’s care meaning the loss of a PCP is likely to impact not just the patient, but also other members of the care team.

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

[4]For a comprehensive review of the Military Health System, see (2017).

PCP Deployments

The source of discontinuities in this study is through military provider deployments. Military providers are generally not assigned to operational (combat) units so that they can practice medicine in clinics and 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 of care metrics. Take the Army for instance[5]: At least annually, the Army reviews operational needs and submits requirements to hospitals for provider of specified specialties including primary care. 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.

PCP Assignment

While the study allows for endogeneity of the PCP match, the military system generally auto-assigns a patient to a primary care provider based on provider availability. Discussions with Primary Care leaders within the MHS implied that even the decision between a physician and a physician extender (i.e. Nurse Practicioner or Physician Assistant) is not taken into account. Appendix figure displays a sample notification of a PCM change due to a provider deployment. However, an important caveat is that a patient can go online or call their TRICARE carrier and request a specific PCP as long as there is room in that provider's panel. In (unshown) diagnostics I find very small, but statistically significant, correlations between PCP practice style and their patients' gender.

The military also uses a civilian network of providers to augment care provided in the Military clinics. Generally speaking, however, patients are not shifted from the direct care to the network absent a move. When a clinic hits capacity it will close to new enrollment and shift these patients to the network rather than disrupting patients that are already enrolled, even if a provider moves or is deployed.

Data

The data for this study primarily come from the Military Health System Data Repository (Defense Health Agency (2007-2017)) and consist of both claims and electronic health records for adult military dependents

[5]The Army changed this system after the sample period, though they still do not deploy providers based on any quality of care metrics

enrolled in Tricare from 2008-2017. For each patient, I observe their assigned clinic and primary care provider, and the date when either of these change along with demographic details including age and gender. I measure health status using the Charlson comorbidity index (Quan et al. (2005)), a standardized score based on documented comborbidities within the medical records. The full sample consists of 1,605,996 individuals with an average enrollment of five years. Due to military moves, changes in PCP are somewhat ubiquitous with the average patient changing providers twice during the sample.

I identify provider's military status and deployments through the Contingency Tracking System (CTS) and military personnel master files (Defense Manpower Data Center (2007-2017*a*);
Defense Manpower Data Center (2007-2017*b*)). The CTS is a database that records when a service-member arrives and finishes an overseas deployment in support of combat operations in either Iraq or Afghanistan or nearby countries supporting those wars. The Master File lists demographic data for all active duty service members.

Using the medical records, I construct a number of utilization and intensity of care measures. Primarily, I focus on relative value units (RVUs). RVUs are a commonly used method of measuring the inputs required to deliver care including physician effort (work RVUs) and administrative overhead (Practice Expense RVUs). RVUs are the basis of physician billing in the United States for TRICARE, Medicare, and most private insurers. Since RVUs represent an aggregate measure of physician utilization this can be thought as a combination of both the extensive margin (how many visits) and the intensive margin (what happens during these visits). As an alternative measure I consider the overall number of visits with physicians or physician extenders (i.e. Physician Assistants/Associates or Nurse Practitioners).

Second, I consider the types of care a patient utilizes. I categorize care as primary care, specialty care, emergency care, or inpatient admissions. While primary care and specialty care utilization denote changes in specialization, the latter two provide insight into whether these changes potentially prevent adverse events. Finally, I consider variation between appointments. I measure this using the total RVUs per visit, as well as the number of specific procedures listed in each visit as coded using the Common Procedural Terminology (CPT) system. CPT codes are primarily used for billing and include both the evaluation and management (E&M) of the patient as well as any specific procedures performed by the provider.

Of note, I use enrollment data to match patients and physicians. In much of the literature, assignment is done based on which provider an individual is observed seeing. While enrollment data makes this assignment cleaner, it means that there are a substantial amount of patients with no utilization in the pre-period. These individuals would have been dropped from any analysis using traditional Medicare claims for instance.

Additionally, some patients have an appointment, but not with their assigned PCP. Overall, about 20% of patients don't use any care in a given year, and 4.8% don't use any care across the sample period. I include these patients in the primary analysis but drop them in further analyses.

Sample Construction

The 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 in the middle of the sample 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 going to the doctor once may be predictive of going again (e.g. for a follow-up visit) and where nearly everyone is affected by a discontinuity-in-care at some point in time.

I address these concerns by using a stacked difference in differences design and later show that my analysis is nearly equivalent to estimators designed by Sun and Abraham (2021) and Callaway and Sant'Anna (2021). Following Deshpande and Li (2019) first I create different datasets for each quarter-year in which a provider deploys and in which I can observe the patient for at least two years prior to the deployment and one year after the deployment. I limit the data to patients who have been enrolled for at least two years prior to a deployment and who remain enrolled in Tricare Prime for at least one year after the deployment. This restriction allows me to identify a provider's patients a year prior to the deployment in order to observe any anticipatory affects. I then create a matched control group of patients with the same length of relationship (in quarter-years) with their assigned PCP in the same quarter-year and assign a cohort identifier. By requiring the same initial relationship length between treated units and their matched controls I am able to control for any lingering effects from a previous discontinuity in care which would jeopardize the parallel trends assumption. While I do not restrict to a specific number of matches per treated unit, I do restrict control group patients to one matching group. For instance if patient A is in the data as a control for 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 further limit the control group to patients who never have a primary care provider deploy in the data. Finally, I drop patients whose provider deploys but is not gone long enough to reasonably cause a discontinuity in care. Analysis shows deployments that are less than 6 months are less likely to lead to a discontinuity in care.

In the main analysis, I consider annual utilization in the 1-3 years before and after the physician deploys. In subsequent analyses I limit the data to the four quarter-years surrounding the deployment and focus in on certain subgroups of patients to better understand heterogeneity in the results. Appendix table 8 lists the sample size reductions from each step.

Empirical Design

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 stacked difference in differences model with a cohort-time fixed effect. I estimate the models using ordinary least squares. Because RVUs are highly skewed, I conduct a log transformation, adding one to each measure in order to deal with any zero-spending, though later I show that my results are robust to a level model.

Identification Strategy

In the main analysis, I estimate the average treatment effect of a discontinuity in care using the following equation:

$$Y_{ipjt} = \alpha + \beta_1^* 1(t > deployment) + \gamma^* X_{pt} + \theta_i + \delta_{jt} + \phi_{ct} + \epsilon_{ipjt} \tag{1}$$

where y_{ipjt} represents an outcome of interest for individual i, assigned to PCP p, part of cohort j at time t. β_1 is the coefficient of interest that indicates the relative change for the treatment group after the provider deploys. X_{pt} is a vector of provider and patient time-varying controls. First, I control for the PCP's specialty. Appendix table 9 displays the breakdown of specialties in the analysis samples by proportion of observations, unique providers, and deployers. Second, I control for whether the patient's assigned PCP is active-duty military since active duty providers may differ from civilian providers in unobservable ways. I return to these potential provider differences in the robustness and alternative explanations section of the paper. Note that most standard patient-level controls are collinear with the fixed effects. At the patient level I control for age, age squared, gender, and Charlson Comorbidity score in regressions where there is variation in these variables (e.g. not collinear with the fixed effects). θ_i represents a vector of individual fixed-effects, δ_{jt} represent a cohort - time interacted fixed effect, and ϕ is a clinic - time fixed effect. Because the cohort is constructed using both the timing of the discontinuity and the existing relationship length, the cohort-time fixed effect controls for any differential trends based on the timing of treatment or previous discontinuities while the clinic-time fixed effect limits the effects of any change that differentially affects a clinic at a period in time. This could occur, for instance, if a provider deploying makes all providers busier within a clinic. Finally, I cluster the standard errors by both the patient and the clinic to account for any serial correlation of the error terms. Note that the timing of the discontinuity is not precise as providers may be required to report to their deploying units 30-90 days prior to the actual deployment. Therefore I focus my analysis on the intent to treat and omit the deployment quarter from regressions.. To the extent that the timing of discontinuities is slightly off, this is likely to bias my results toward zero.

In order to provide evidence that loss of the interpersonal relationship is the mechanism rather than some correlated impact of the deployment itself, I take a triple differences approach leveraging a particular organizational idiosyncrasy. While patients are assigned a specific PCP, they can book an appointment with any PCP within the practice. I leverage patients who, for likely-endogenous reasons such as appointment availability or good health, have never had an appointment with their PCP prior to the deployment. This could be because they didn't use any care or because they happened to see someone else in the practice when they did use care. Overall, about 1/3rd of patients fall into this category. With the assumption that this decision is orthogonal to the physician deployment, I estimate equation 2.

$$\begin{aligned} Y_{ipjt} = \alpha + \beta_1^* 1(t > deployment) + \beta_2 * (t > deployment) * (visits > 0) \\ \beta_3^* 1(t \geq deployment) * (visits > 0) + \gamma^* X_{pt} + \theta_i + \delta_{jt} + \phi_{ct} + \epsilon_{ipjt} \end{aligned} \tag{2}$$

β_1 in this equation is the coefficient on the interaction of the post-period indicator and the treatment group. β_2 is the coefficient on the interaction of the post-period indicator and an indicator for those who have a relationship with their provider (compared to those who never met their provider). Finally β_3 is the coefficient of interest - the triple interaction of the treatment indicator, the post-period indicator, and having a relationship. This can be interpreted as the relative effect of the deployment for those who had some relationship to disrupt compared to those whose provider deployed, but didn't have an existing relationship and all else is as in equation 1. Note that other variables in the canonical triple difference equation, including the treatment group - existing relationship interaction, are collinear with the fixed effects. Of course, patients who don't see their PCP are likely healthier on average as well so the actual estimates should be taken with caution. Importantly though, this identifies the portion of the sample that drive the results providing the opportunity for more detailed analysis of the focal group.

Finally, I consider differences between individual encounters. In these regressions I modify equation 1 by swapping the patient fixed effect for a provider fixed effect. I then cluster standard errors by physician and clinic.

3 Results

I begin the section by explicitly laying out the model's key assumptions and providing evidence that these assumptions are met. In order to interpret the results as the effect of a discontinuity in care, the deployments must be quasi-random and must only affect outcomes through their impact on continuity of care. This would be violated if, for example, patients begin seeing providers of lower average quality after a deployment or if providers change their practice style prior to the deployments. It's also possible providers offload either their unobservably (to the econometrician) healthier or sicker patients prior to the deployment. Therefore I focus my analysis on the"intent to treat" - the actual provider deployment date rather than the change in enrollment date and show there are no pretrends in utilization. To the extent that having an earlier discontinuity in care creates any bias in the estimates though, it is likely to go toward zero.

First, I consider any differences between patients that are affected by the deployments compared to the control group of patients whose PCP never deploys. Table 1 panel A displays the results of a regression of patient demographics on an indicator for undergoing a primary care provider deployment. The regression includes provider military status and specialty controls, clinic-time, and cohort-time fixed effects.

[Table 1 about here.]

The treatment group tends to be about a year and a half older, has fractional more comorbidities, and tilts a bit more female. The age and gender difference may be an artifact of the matching process. Since the longer a patient is enrolled the more likely they are to experience a PCP deployment, conditioning on a never treated group may create a small age disparity. Similarly, female patients remain enrolled in the TRICARE HMO product for longer on average (6 years for females compared to 4.25 years for males). Importantly, the differences in comorbidities disappear when controlling for age and gender (not shown). Panel B shows differences in utilization using the same regression. Only emergency department utilization is significant at the 95% level while specialty care has a small, difference at the 90% level. Panel C repeats these regression but windsorizes the utilization data at the 95th percentile of the distribution for each outcome. The small differences in utilization almost entirely disappear. I use these windsorized values in regressions using visit counts. Note that the individual fixed effect and balanced panel construction used in most regressions remove the small differences in age, gender, and underlying health with the exception of a specific health shock.

Next, I show support for the assumption that deployments are quasi-random. Table 2 shows the results of a series of regressions of practice intensity. Column 2 displays the results of a regression on an indicator for whether a primary care provider deploys within the sample window using the full sample of primary care visits. The regressions include patient level controls (age, age squared, gender, comorbidity score) along with PCP specialty and clinic-quarter fixed effects. Deployers and non-deployers appear to practice in similar fashion, with no statistically significant differences in the log of RVUs per visit, the probability of referring a patient to specialty care or the number of CPT codes per visit. Column 3 restricts the analysis to primary care providers that deploy and those that gain the deploying providers' patients after the deployment. There are no observable differences in this group as well indicating that patients aren't being sorted to higher or lower intensity providers after a deployment. Finally, I check for changes in practice patterns during the year prior to a deployment. Here, I limit the sample to deployers and regress each appointment intensity indicator on an indicator for being within one year of a deployment along with patient controls and physician and clinic-time fixed effects. Column 4 shows the results of these regressions. There's no evidence of a change in practice patterns for any of the measures.

In appendix table 10 I provider further support that provider deployments are random. Each column shows the coefficients and joint F stat on a regression for whether an individual PCP ever deploys conditional on being in the military. Because I observe demographic information for Military Service Members, I am able to look at both patient and provider demographics. I also consider the same three measures of practice intensity. All three regressions have joint F stats less than 1. Provider intensity measures have a microscopic joint F stat of 0.11 indicating that none of these measures are predictive of a PCP deployment.

[Table 2 about here.]

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. I estimate this by exchanging time-dummies for the post-period indicator in equation 1.

[Figure 1 about here.]

Figure 1 displays a series of event studies that include the three years before and after a discontinuity in care with the year prior being the omitted category. Note that the sample restrictions only require two years of pre-data and one year of post data so this is an unbalanced sample. Because attrition may not be random, caution should be taken in interpreting the results in year two and three. Figure 1a shows changes in the natural log of RVUs. The estimates are stable in the years prior to the deployment indicating parallel pre-trends and providing support for the parallel trends assumption. RVUs increase as the PCP departs indicating that there is an effect of the discontinuity on total care utilization. RVUs appear to begin to come back down over time and are not statistically distinguishable from zero in year 3 though the point estimate is still somewhat elevated. Figure 1b shows the natural log of total provider visits over time. There is not a noticeable decline in visits with fairly stable estimates across all three years. Figure 1c & 1d shows the change in specialty care and primary care appointments (windsorized at the 95th percentile). The trend in specialty care is stable with a slight decline in year three, but primary care has the opposite affect with a sharp increase after the provider leaves and increasing over time.

Table 3 presents the difference in differences results for each of the four main utilization variables. Column 1 shows the effect of a physician deployment on the log of RVU's in the post period. The coefficient indicates a 4.4% percent increase in total utilization of physician services. Column 2 displays the effects of the discontinuity on the log of total visits with a slightly smaller 3.9% increase in total visits. Columns 3 and 4 display the results for specialty and primary care visits. Given the (windsorized) sample means these indicate an approximately 3.3% increase in specialty care and 3.7% increase in primary care over the three years after a discontinuity.

[Table 3 about here.]

[Table 4 about here.]

The increased utilization may be efficient if patients are benefiting through better outcomes. For example, previous literature found that after a discontinuity new providers were more likely to diagnose underlying illnesses (Simonsen et al. (2021)). Table 4 presents linear probability models for whether a patient ends up in either the emergency department or is admitted to a hospital as an inpatient as well as the change in the Charlson comorbidity score after a discontinuity. The coefficients on all three are extremely small and statistically indistinguishable from zero. While the administrative data cannot definitively rule out an effect on underlying health, I do not find evidence that the increased use of outpatient care is preventing any negative events.

If the loss of a relationship is the mechanism then the effects should be driven by those with an existing relationship. In table 5, I present the results of estimating equation 2, a triple difference model that compares the effects of a provider deployments for those that have a recent (within the two years before the deployment) visit with their PCP to those who do not. Each row displays one of the three interactions in the model that can be estimated, noting that the fourth interaction is collinear with the fixed effects. The third row shows the coefficient on the triple interaction, which is the coefficient of interest. The estimates are nearly equivalent to the main difference in differences results indicating that those with relationships to disrupt are driving the results. For the remaining analyses I limit the sample to those with relationships and zoom in on the four quarter-years before and after a physician deployment.

[Table 5 about here.]

3.1 Quarter-Year Sample Main Results

I begin this section by repeating the annual analysis at the quarter-year level with the sample of patients that have at least one PCP visit in the two years prior to the discontinuity. An advantage of quarterly data is that I can include clinic-quarter fixed effects which will pick up temporary shocks at the clinic level that last for less than a year - for instance if physicians are busier or access to care has been hampered due to a deployment. Figure 2 displays event studies at the quarter-year level. 2a and 2b display the log of RVU's and total visits respectively. 2c & 2d display linear probability models where the outcome is the probability of a specialty or primary care visit in the quarter-year. The figures show that here are no pretrends in the pre-period and a visual jump on all measures accept primary care immediately after the discontinuity. Noticeably there is not a decline in primary care utilization. This is contrary to previous literature that has shown a drop in primary care after a discontinuity (Van Walraven et al. (2010); Staiger (2022); Sabety, Jena and Barnett (2021)). This could be due to the gatekeeping model where a patient cannot seek specialty care without first seeking primary care.

Table 11 displays the difference in differences results for this sample with the quarter prior to the deployment as the omitted category. The change in RVU's (column 1 - 4.7%) is slightly higher and total visits (columns 2 - 3.3%) are slightly smaller than the annual sample but very similar. Columns 3 & 4 show linear probability models of using any specialty or primary care respectively. The 1.7 percentage point increase on a base of 0.33 indicates that the probability of specialty care utilization goes up by about 5% slightly higher than the 3.3% increase in the annual model. The linear probability model of any primary care use has a smaller increase (2.3% compared to 3.7%) than in the annual sample. This is likely because primary care is affected more on the intensive margin than the extensive margin. That is, individuals don't get sick and require a primary care appointment because their provider deployed, but conditional on being sick may require more appointments. This contrasts with specialty care where the affect may be driven by those who would not be referred to specialty care at all absent the PCP deployment.

Next, I analyze which types of care drive the change in RVUs by decomposing the RVU effect into three buckets - Primary Care RVUs, Specialty Care RVUs, and Emergency Department RVUs. Table 12 displays the change in RVUs in levels to make interpretation easier. I present models using both the observed and the windsorized values for each type of RVU. The windsorized values are somewhat less informative but I include them as suggestive evidence for how much of the change is driven by extreme values. Column 1 shows that RVU's increase by 0.301 per quarter which correlates to a 3.8% increase, slightly lower than the log model and only slightly higher than the windsorized model. Columns 2-4 display the results for each type of care. The largest increase is in specialty care, which drives about 45% of the total increase. Primary Care and the ED drive 27% and 25% respectively. The last 3% is driven by other visits that don't fall into one of these three buckets (for instance physician inpatient rounds). Perhaps the most interesting aspect here is the increase in ED RVUs despite not observing an increase in the probability of an ED visit. However, when looking at the windsorized model the effect disappears indicating that this is driven by a few very high RVU visits to the ED.

While the provider deployment provides an *average treatment effect*, there is likely considerable heterogeneity in the actual effect on patients. Many patients simply come in once a year or come in with acute conditions that have little uncertainty where specific knowledge may not be as impactful. In this section I consider the different effect across three groups of patients - diabetics, those who were in the top quartile in RVUs in the year prior to the discontinuity, and those who are pregnant at some point in the year before or after the discontinuity. Appendix table 14 displays the distribution of these and other subgroups in the paper between the treatment and control groups. Table 6 displays the results of these regressions. Column 1 shows the results for diabetic patients for each of the primary utilization and outcome measures. Diabetic patients have their cost of care increase by 22% with large increases in both specialty care (8.8%) and primary care (6.8%). This group has a much larger coefficient on the probability of ED use than the annual sample, though imprecisely measured. Column two displays the results for those who were in the top-quartile of RVUs in the year prior to the PCP deployment. These patients potentially had a health shock in the previous year and may be more reliant on their primary care provider for either treatment or coordination of care. The first row indicates that there is about an 11% increase in RVUs for these patients - more than double the magnitude of the effect on the main sample and about half of the magnitude of the effect on the diabetic sub-sample. They have a slightly lower though similar increase in the probability of using specialty care (6.2%) as the diabetic sub-sample, but no distinguishable effect on any of the other outcomes. Column three displays the results for patients who were pregnant at some point in the sample period. The effect on this sample is fairly similar to the main sample with about a 5.5% increase in RVUs and 4.5% in the probability of using specialty care indicating that these patients are not driving the results. In total, these results suggest that for patients who require significant coordination of care, the loss of a provider is extremely costly.

[Table 6 about here.]

Next, I consider how the length of the relationship and recency of contact affect the main results. I consider patients who had their most recent visit in the 6 months prior to the discontinuity, 6-12 prior to the discontinuity, and more than one year prior to the discontinuity. In table 13 Panel A I display the results of these regressions. There is a linear relationship between how recent the previous visit with the PCP was and the magnitude of the effects across all four utilization measures. Of course, there is also a linear relationship between overall utilization and recency of visit which supports the results in the previous section that discontinuity in care is most impactful for those who likely require the most care coordination.

In panel B I display the results of a regression on those who were in the top or bottom quartile of relationship length. The top quartile includes relationships that would have been 8- 9 quarters (2 - 2.25 years) at the time of the deployment. The top quartile includes relationships that would have been 8-16 quarters (3-5 years) at the time of the deployment. While there are slight differences between the groups, recency of visit seems to be much more impactful than the total length of the relationship - at least within the relatively narrow margin in this data set.

Specialty Encounters

In this section I consider the impact on specialist visits. For this analysis, I restrict the sample to specialty appointments within military hospitals where I observe more information about the individual specialty provider and where coordination between specialty and primary care physicians is more likely to occur. I exclude from this analysis any visits that do not individually generate RVUs (e.g. post-surgical follow-ups or pregnancies that are paid with a "global payment" designed to cover all pre-natal care as well as the delivery).

I begin the analysis by considering whether primary care providers are referring patients to a new specialty or back to a clinic in which the patient was previously seen. I define a new specialty as one in which the patient was not seen in the year prior to the match (i.e. between 1 2 years before the deployment). For each patient-specialty clinic grouping, I also observe whether the patient is seen in that clinic multiple times or only once. The intuition is that a patient that goes to a new clinic only once may represent a less appropriate referral compared to a patient that requires follow-up care within the specialty clinic. I estimate the change in probability for each of these measures - a new specialty visit and a single-visit referral - using equation 1. Table 15 presents the results of these regressions. The coefficient in column 1 indicates an 11% (0.009 pp) relative increase in the probability of a visit in a new specialty clinic - about double the magnitude of the overall increase in the probability of specialty care utilization. Column 2 indicates a similar 10% (0.004 pp) relative increase in the probability of only having one encounter in a specialty clinic. Together, these findings suggest that the increase in specialty care is not the result of patients reverting to existing specialty relationships, but also does not support the contention that these are inappropriate referrals. While these patients may have been able to be treated within primary care had their provider not deployed, this evidence suggests referring is the appropriate decision given the discontinuity in care.

To better understand what departments drive the increase in new referrals I consider the five types of outpatient specialty care as defined by the first two digits of the Military Health System's cost accounting system and coded in the medical records. These include Obstetrics/Gynecology (OB/GYN), General Medical Services (e.g. dermatology or neurology), Surgical Services, Orthopedics, and Mental Health. I omit one category of specialty care - rehabilitative services - as it is rare in the data. If the referrals are due to the loss of a relationship that involves trust, communication, and specific knowledge, then areas with more uncertainty should see larger increases. While there is no single consensus on the relative uncertainty across specialties, it is reasonable to assume that OB/GYN, general medical services, and mental health care have more uncertainty than orthopedics and general surgery given existing diagnostic technology (e.g. MRIs). I consider three groups in this analysis: The full sample, only patients who are pregnant at some point in the year before or after the discontinuity, and their opposite - only patients who are not pregnant in the sample. Note that I only consider new referrals in this analysis but patients may also increase their use of specialties where they area already being seen.

Appendix table 16 displays the results of these regressions. As expected, the there is a large increase in OB/GYN (21%, 0.003 pp) and general medical services (0.13% 0.003 pp) in the full relationship sample. The rate of new referrals to OB/GYN is similar across both pregnant and not pregnant patients, but pregnant patients have a higher referral rate to general medical services. In contrast, there is no effect on surgical services or orthopedics. Likewise, there is no effect on mental health care.

Finally, I look across all specialty encounters for patients in the main sample for evidence of a change in intensity of care by the specialty provider. Here I include the patients who were previously excluded due to lack of a relationship with their PCP. I do this because even absent a visit, the PCP may have been coordinating a patient's care. The loss of a primary care provider to coordinate care could cause specialists to have to spend more time and effort on each patient. I estimate any changes in the intensity of care in terms of both log RVUs per visit and the number of procedures that are coded on the physician claim. Table 7 panel A presents the results of these regressions applied to the full sample of specialty encounters. I find a marginally significant 1.8% increase in within encounter RVUs but no change in the number of procedures.

I then subset the group along two margins in which theory would predict an effect. First, I consider patients who see a provider with whom they have a long-term relationship. I define long-term relationship as a provider with whom the patient had an encounter in the year leading up to the match-quarter (i.e. more than 1 year prior to the deployment). Second, I consider any specialist visit for patients multiple comorbidites as measured by a Charlson Score ≥ 2. These patients are likely to require the most coordination of care. Table 7 presents the results of these regressions in panels B & C respectively. The existing relationship group has a similar affect as the main sample - an approximately 2.8% increase in RVUs per visit though also significant only at the 90% level and no change in the number of procedures. The multiple comorbidity sample has a much higher, statistically significant 14% increase in RVUs and an imprecisely measured 0.051 (3%) increase in procedures per visit. To put the RVU increase in perspective, this equates to 0.29 RVU's per visit - roughly equivalent to the average effect *per quarter-year* on the full relationship sample.

[Table 7 about here.]

Alternative Explanations & Robustness Checks

While the results so far have shown the effects of a deployment, there could be several reasons why a deployment causes a change in utilization aside from discontinuity in care. First, patients could be moved to higher or lower quality providers. While I previously presented evidence that deployments are quasi-random, I augment this by presenting a model that changes the individual fixed effect for a PCP fixed effect. Appendix table 17 presents the results of these regressions. All four primary measures remain statistically significant.

Second, the type of provider a patient is sent to could drive the results. Patients may be more likely to be matched to a civilian provider, either employed by the military or in the private sector, or a non-physician provider, each of which may differ in their quality from active duty physicians. While I control for both of these in the main estimation, I also run a series of regressions limiting the sample to those with the same provider type in the pre and post periods. Finally, moves are ubiquitous in the military. While moves should not be correlated with deployments, I estimate a model limiting the sample to those who remain enrolled in the same clinic in the year before and after their PCP deploys. Appendix table 18 shows the results of these four robustness checks designed to test alternative explanations. In panel A, I display the results from limiting patients to an active duty military provider in the pre and post-period. In panel B, I present the results from eliminating anyone who switches to a purchased care (private sector) primary care provider. In panel C, I restrict the analysis to those who have physician PCPs in both the pre and post period. As the choice of a physician versus a physician extender may not be random, these are potentially higher intensity patients. In panel D, I limit the sample to those that remain enrolled in the same clinic throughout the sample period. Despite differences in sample size and precision, the four robustness checks offer evidence that the results are stable and are not driven by these alternative explanations.

In a check to my choice of estimation strategy 3 I show event studies for the log of RVU utilization using techniques designed by Sun and Abraham (2021) and Callaway and SantâAnna (2021). Both event studies resemble the quarterly event-study presented in figure 2a indicating that the results are robust to alternative estimation strategies. Due to computing memory constraints the Callaway and SantâAnna is estimated using a 25% random sample (at the individual level to maintain a balanced panel) of controls.

Next, I show that the results are not driven by my choice of a 6 month cut-off for a deployment to "count" as a discontinuity in care. While short deployments aren't necessarily a falsification test as they can lead to discontinuities in care - in general they will not and should have smaller or zero results. In contrast longer deployments almost certainly lead to discontinuous care. In appendix table 19 I display the results regressions limited to the top and bottom quartiles of deployment length. The top quartile has higher-magnitude effects across the four utilization measures, while the bottom quartile only has a significant effect on log RVUs.

Finally, I consider whether movement to the purchased care, private sector network could drive the results. Here, I don't have the sample size to subset the data so I equation 2 and estimate a triple differences model that includes the network PCP identifier in a triple interaction term. Appendix table **??** displays the results of these regressions. The triple interaction affect on the log of RVUs indicates an 8.5% increase for those that move to the private sector, though this lacks precision. Thee magnitude on the point estimates for log of total visits and probability of specialty care are also higher than in the main sample, though not statistically distinguishable from zero. This is an interesting subgroup and an open area for future research.

4 Discussion and Conclusions

In this paper I study the effects of provider turnover on patient utilization using quasi-random discontinuity in patient care. I apply a difference in differences model to the Military context where primary care providers are pulled from their practices in the midst of treating and coordinating care for a panel of patients. By considering overall utilization and within specialty visit variation, I am able to provide new information on the value of specific information.

The findings indicate that loss of a relationship causes utilization to increase by 3-5% RVUs per patient and that this effect endures for a long period of time. Additionally, the average hides substantial heterogeneity with patients who require coordinated care increasing utilization by 20% or more while low utilizers aren't impacted at all. I also find that a primary care discontinuity has an effect on other members of the care team. Examining specialty encounters, I find that visits become more intense for patients with multiple comorbidites. These effects add to the literature on physician exits by isolating the effect of the relationship from other factors associated with discontinuity in care.

This study has important implications for health care where many recent policies and organizational innovations have been focused on reducing coordination costs. There is a substantial tradeoff, however, in that policies focused on decreasing loss of information may limit the accumulation of new information. For instance, Electronic Health Records are designed to expand access to generalizable information, yet may take the provider's attention away from the patient potentially inhibiting trust and communication between the patient and provider and resulting in less accumulation of new, specific knowledge. Another example is the patient centered medical home (PCMH) model. This complex model is designed to promote knowledge-sharing 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 overall. It's theoretically ambiguous whether this can improve outcomes.

Beyond the health care context, this work has implications for how organizations organize and allocate tasks. For instance, a firm may weigh the benefits of repeated interactions between team-members and a customer with the potential loss that occurs when a team-member departs. While it's beyond the scope of this paper to offer specific recommendations, firms may consider policies and organizational structures that promote knowledge management and sharing of information. If specific knowledge is contained within individuals, firms may adopt policies that provide more opportunities for 'warm handoffs' in which some knowledge can be shared and trust with new team-member can begin to be built prior to a team-member departing. Firms may also adopt technology that reduces the cost of information transfer, though this is likely not a first-best solution. Finally, effective policies may want to not only address information about the customer, but also information regarding the best way to coordinate within the team.

References

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

Agha, Leila, Keith Marzilli Ericson, Kimberley H Geissler, and James B Rebitzer. 2018. "Team formation and performance: evidence from healthcare referral networks." National Bureau of Economic Research.

Becker, Gary S, and Kevin M Murphy. 1992. "The division of labor, coordination costs, and knowledge." *The Quarterly Journal of Economics*, 107(4): 1137–1160.

Bischof, Tamara, and Boris Kaiser. 2021. "Who cares when you close down? The effects of primary care practice closures on patients." *Health economics*, 30(9): 2004–2025.

Callaway, Brantly, and Pedro HC SantâAnna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Chan, David C. 2021. "Influence and information in team decisions: evidence from medical residency." *American Economic Journal: Economic Policy*, 13(1): 106–37.

Clark, Jonathan R, Robert S Huckman, and Bradley R Staats. 2013. "Learning from customers: Individual and organizational effects in outsourced radiological services." *Organization Science*, 24(5): 1539–1557.

Connelly, Catherine E, and E Kevin Kelloway. 2003. "Predictors of employeesâ perceptions of knowledge sharing cultures." *Leadership & Organization Development Journal*, 24(5): 294–301.

Defense Health Agency. 2007-2017. "Military Health System Data Repository." *https://health.mil/Military-Health-Topics/Technology/Support-Areas/MDR-M2-ICD-Functional-References-and-Specification-Documents.*

Defense Manpower Data Center. 2007-2017*a*. "Contingency Tracking System." *https://dmdcrs.dmdc.osd.mil/dmdcrs.*

Defense Manpower Data Center. 2007-2017*b*. "US Military Master File." *https://dmdcrs.dmdc.osd.mil/dmdcrs.*

Deshpande, Manasi, and Yue Li. 2019. “Who is screened out? Application costs and the targeting of disability programs.” *American Economic Journal: Economic Policy*, 11(4): 213–48.

Evaluation of the TRICARE Program: Fiscal Year 2017 Report to Congress. 2017. “Evaluation of the TRICARE Program: Fiscal Year 2017 Report to Congress.” Department of Defense.

Ford, Dianne P. 2004. “Trust and knowledge management: the seeds of success.” *Handbook on Knowledge Management 1: Knowledge Matters*, 553–575.

Ford, Dianne P, and Sandy Staples. 2010. “Are full and partial knowledge sharing the same?” *Journal of knowledge management*, 14(3): 394–409.

Frakes, Michael D, and Melissa F Wasserman. 2021. “Knowledge spillovers, peer effects, and telecommuting: Evidence from the US Patent Office.” *Journal of Public Economics*, 198: 104425.

Fritz, Zoë, and Richard Holton. 2019. “Too much medicine: not enough trust?” *Journal of medical ethics*, 45(1): 31–35.

Garicano, Luis. 2000. “Hierarchies and the Organization of Knowledge in Production.” *Journal of political economy*, 108(5): 874–904.

Goodman-Bacon, Andrew. 2018. “Difference-in-differences with variation in treatment timing.” National Bureau of Economic Research.

Hayek, Friedrich August. 1945. “The use of knowledge in society.” *The American economic review*, 519–530.

Huckman, Robert S, and Gary P Pisano. 2006. “The firm specificity of individual performance: Evidence from cardiac surgery.” *Management Science*, 52(4): 473–488.

Hutter, Paul J, Joachim Roski, Jonathan Woodson, Allen Middleton, Robert Kneeland, Averston Worthy, Derik Zitelman, Thomas Trinh, Shaun Dela Cruz, and Elaine Cooper. 2019. “Readiness Of Medical Providers In The Military Health System: Overview Of Operational And Policy Considerations.” *Health Affairs*, 38(8): 1274–1280.

Jensen, Michael C, and William H Meckling. 1992. “Specific and general knowledge and organizational structure.”

Kwok, Jennifer. 2019. “How do primary care physicians influence healthcare? Evidence on practice styles and switching costs from Medicare.” *Evidence on Practice Styles and Switching Costs from Medicare (July 22, 2019)*.

Mainous, Arch G, Richard Baker, Margaret M Love, Denis P Gray, and James M Gill. 2001. “Continuity of care and trust in oneâs physician: evidence from primary care in the United States and the United Kingdom.” *Fam Med*, 33(1): 22–27.

Mayer, Roger C, James H Davis, and F David Schoorman. 1995. “An integrative model of organizational trust.” *Academy of management review*, 20(3): 709–734.

Muthusamy, Senthil K, and Margaret A White. 2005. “Learning and knowledge transfer in strategic alliances: A social exchange view.” *Organization studies*, 26(3): 415–441.

Press, Matthew J. 2014. “Instant replayâa quarterback’s view of care coordination.” *New England Journal of Medicine*, 371(6): 489–491.

Quan, Hude, Vijaya Sundararajan, Patricia Halfon, Andrew Fong, Bernard Burnand, Jean-Christophe Luthi, L Duncan Saunders, Cynthia A Beck, Thomas E Feasby, and William A Ghali. 2005. “Coding algorithms for defining comorbidities in ICD-9-CM and ICD-10 administrative data.” *Medical care*, 1130–1139.

Sabety, Adrienne H, Anupam B Jena, and Michael L Barnett. 2021. “Changes in Health Care Use and Outcomes After Turnover in Primary Care.” *JAMA Internal Medicine*, 181(2): 186–194.

Saultz, John W. 2003. “Defining and measuring interpersonal continuity of care.” *The Annals of Family Medicine*, 1(3): 134–143.

Shrank, William H, Teresa L Rogstad, and Natasha Parekh. 2019. “Waste in the US health care system: estimated costs and potential for savings.” *Jama*, 322(15): 1501–1509.

Simonsen, Marianne, Lars Skipper, Niels Skipper, and Peter Rønø Thingholm. 2021. “Discontinuity in care: Practice closures among primary care providers and patient health care utilization.” *Journal of Health Economics*, 80: 102551.

Staiger, Becky. 2022. “Disruptions to the patient-provider relationship and patient utilization and outcomes: Evidence from medicaid managed care.” *Journal of health economics*, 81: 102574.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199.

Van Walraven, Carl, Natalie Oake, Alison Jennings, and Alan J Forster. 2010. "The association between continuity of care and outcomes: a systematic and critical review." *Journal of evaluation in clinical practice*, 16(5): 947–956.

A Appendix

[Figure 2 about here.]

[Table 8 about here.]

[Table 9 about here.]

[Table 10 about here.]

[Table 11 about here.]

[Table 12 about here.]

[Table 13 about here.]

[Table 14 about here.]

[Table 15 about here.]

[Table 16 about here.]

[Table 17 about here.]

[Table 18 about here.]

[Figure 3 about here.]

[Table 19 about here.]

[Table 20 about here.]

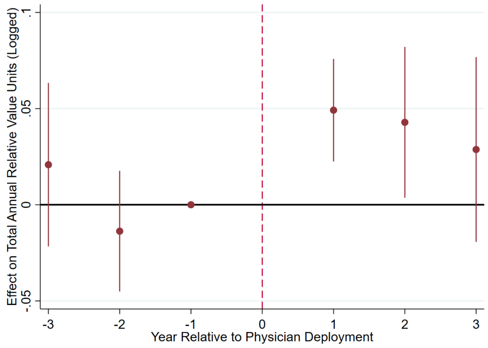

(a) Total RVUs (Logged)

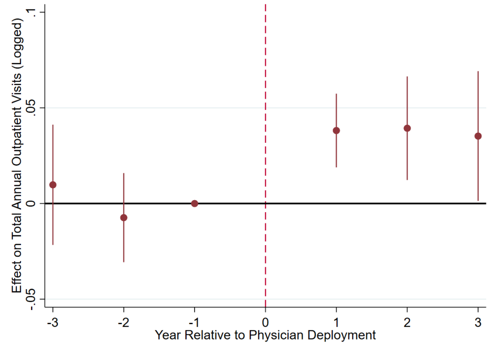

(b) Total Visits (Logged)

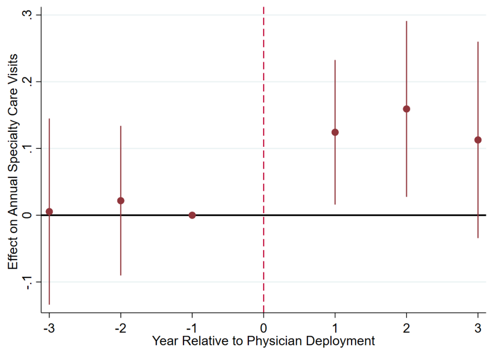

(c) Specialty Care Visits (Windsorized)

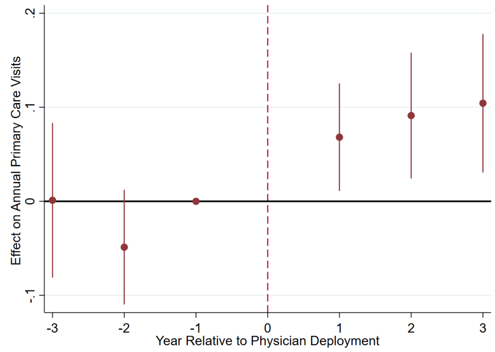

(d) Primary Care Visits (Windsorized)

Note:
Event studies display the coefficient on the interaction of the elapsed-time and treatment group indicator for year relative to the provider deployment. The year prior to the deployment is the omitted category for all figures. 1a shows the change in the log of total RVUs. 1b displays the change in the log of total visits with either a physician or physician extender including nurse practicioners and physician assistants. 1c and 1d show the changes in the total number of specialty and primary care visits respectively using visit counts windsorized at the 95th percentile of the distribution. Regressions include patient controls (age, age squared, gender, and Charlson comorbidity score), PCP controls including specialty and military status and patient, clinic-year, and cohort-year fixed effects. Standard errors are clustered by patient and clinic.

Figure 1: Event Studies of Utilization Measures

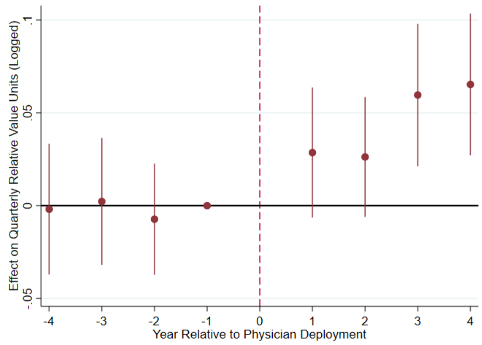

(a) Total RVUs (Logged)

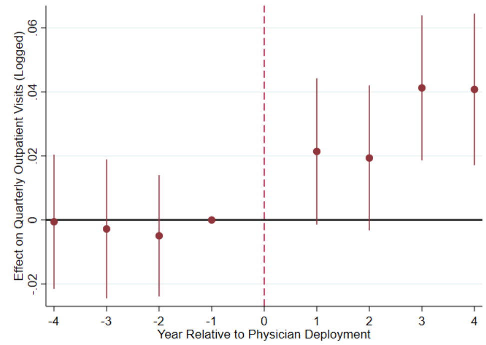

(b) Total Visits (Logged)

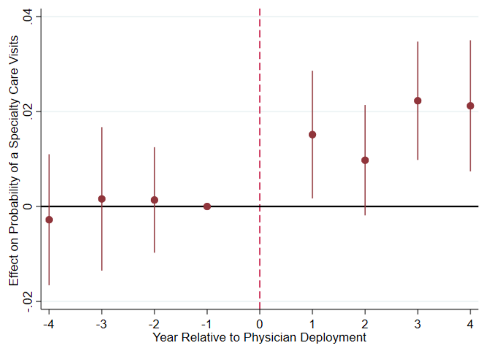

(c) Probability of a Specialty Care Visit

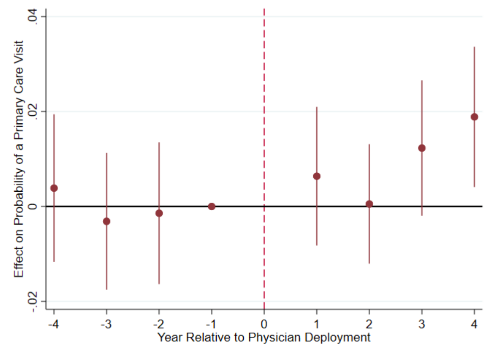

(d) Probability of a Primary Care Visit

Note:
Event studies display the coefficient on the interaction of the elapsed-time and treatment group indicator for each quarter-year relative to the provider deployment. The quarter prior to the deployment is the omitted category for all figures. 2a shows the change in the log of total RVUs. 2b displays the change in the log of total visits with either a physician or physician extender including nurse practitioners and physician assistants. 2c and 2d show the changes in the probability of a specialty and primary care visit respectively. Regressions include PCP controls including specialty and military status and patient, clinic-time, and cohort-time fixed effects. I omit the deployment quarter from these regressions. Standard errors are clustered by patient and clinic.

Figure 2: Event Studies of Utilization Measures

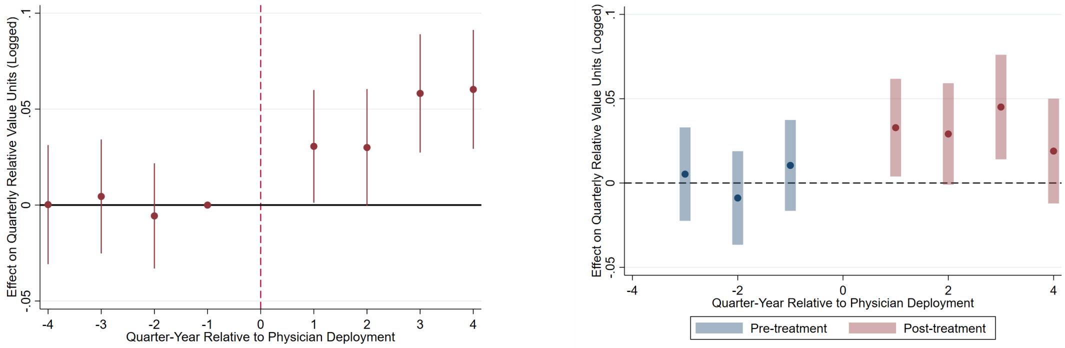

(a) Sun and Abraham Estimator

(b) Callaway and Sant'Anna Estimator

Note:
This figure presents event studies using recent advances in the staggered timing differece in differences literature. 3a presents an event study of Log RVU's on the quarter-year relative to a provider deployment for the relationship sample using the Sun and Abraham estimator (Sun and Abraham (2021)). The quarter prior to the deployment is the omitted category. 3b presents an equivalent event study using the Callaway and Sant'Anna estimator (Callaway and SantâAnna (2021)). Due to computing memory constraints the Callaway and Sant'Anna is estimated using a 25% random sample (at the individual level to maintain a balanced panel) of controls. For this figure the 1st quarter is the omitted category.

Figure 3: Event Studies of Utilization Measures using Alternative Estimation Techniques

Table 1: Patient Group Comparisons

	(1) Sample Mean	(2) Coefficient On Treated	(3) Std Error
Panel A - Demographics			
Age	31.34	1.43***	(0.145)
Gender - Female	0.91	0.017***	(0.005)
Charlson Comborbidity Index	0.20	0.016***	(0.005)
Panel B - Annual Utilization			
Total RVUs	28.10	0.260	(0.365)
Specialist Visits	4.39	0.120*	(0.065)
Primary care Visits	2.97	0.048	(0.038)
Emergency Dept Visits	1.05	-0.047**	(0.023)
Inpatient Admissions	0.15	-0.000	(0.004)
Panel C - Annual Utilization (Windsorized)			
Total RVUs	25.29	0.125	(0.280)
Specialist Visits	3.77	0.059	(0.046)
Primary care Visits	2.70	0.030	(0.032)
Emergency Dept Visits	0.80	-0.021*	(0.012)
Inpatient Admissions	0.13	-0.003	(0.003)
N - Unique Patients			
Control Group	505,204		
Treatment Group	14,992		

Note: Coefficients from a regression of the dependent variable on an indicator for whether the patient is in the treatment group, along with provider controls for specialty and military status. Panels B and C include controls for patient age, age squared, gender, and Charlson comorbidity index. Differences are estimated over the pre-period only.

Table 2: Primary Care Encounter Intensity

	(1) Sample Mean	(2) Coefficient on Deployer	(3) Coefficient on Gaining PCP	(4) Anticipatory Period
Log RVUs	1.53	-0.014	0.017	0.005
		(0.009)	(0.013)	(0.012)
Prob. of a Referral	0.05	-0.001	0.000	0.008
		(0.001)	(0.001)	(0.001)
Num. of Procedures	1.12	0.004	-0.004	-0.010
		(0.005)	(0.008)	(0.006)
Total PCP's	9,679			
Deploying PCP's	586			
Gaining PCP's	2,507			
Observations		12,071,316	6,965,927	788,810

Note: This table displays evidence that deployments are quasi-random. Column 1 displays the sample means from all primary care appointments in the data for deploying and control primary care providers. Column 2 displays the results of a series of regressions of intensity variables on an indicator for whether a primary care provider (PCP) deploys within the sample window using the full sample of primary care visits. The regressions include patient level controls (age, age squared, gender, comorbidity score) along with PCP specialty and clinic-quarter fixed effects. Columns 3 displays results limiting the analysis to losing and gaining PCPs. Column 4 displays the results of a series of regressions limiting the sample to deploying PCP and regressing each appointment intensity measure on an indicator for being within one year of a deployment along with patient controls and physician and clinic-time fixed effects. Standard errors are clustered by PCP and clinic.

Table 3: Main Results - Annual Utilization

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Total Specialist Visits	(4) Total Primary Care Visits
Discontinuous Care	0.044***	0.039***	0.123***	0.101***
	(0.014)	(0.010)	(0.045)	(0.025)
Sample Mean	28.10	8.41	3.77	2.70
Observations	2,290,229	2,290,229	2,290,229	2,290,229

Note: Results of estimating equation 1 for utilization measures on the annual sample. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include patient age, age squared, gender, and Charlson comorbidity index, PCP military status, and PCP specialty. All regressions include individual, clinic-year, and cohort-year fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 4: Main Results - Outcomes

	(1) Probability of Emergency Dept Visit	(2) Probability of Inpatient Admission	(3) Change in Charlson Score
Discontinuous Care	-0.001	-0.001	-0.002
	(0.004)	(0.003)	(0.004)
Sample Mean	0.31	0.04	0.20
Observations	2,290,229	2,290,229	2,290,229

Note: Results of estimating equation 1 for patient outcomes measures on the annual sample. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include patient age, age squared, gender, and Charlson comorbidity index, PCP military status, and PCP specialty. The Charlson index is omitted in column 3. All regressions include individual, clinic-year, and cohort-year fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 5: Results of Triple Difference

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Total Specialist Visits	(4) Total Primary Care Visits
Post-Period*Treatment Group	-0.009	-0.004	0.025	-0.047
	(0.018)	(0.012)	(0.060)	(0.034)
Post-Period*Relationship Group	-0.220***	-0.217***	-0.307***	-0.830***
	(0.007)	(0.005)	(0.020)	(0.016)
Post*Treatment*Relationship	0.064***	0.047***	0.130	0.156***
	(0.023)	(0.017)	(0.087)	(0.041)
Sample Mean	25.29	8.41	3.77	2.70
Observations	2,290,229	2,290,229	2,290,229	2,290,229

Note: Results of estimating equation 2. The first row displays the interaction of the treatment group identifier and the post period identifier. The second row displays the interaction of the post period identifier and an indicator for whether a patient had a visit with their PCP in the 2 years prior to the deployment (or matched time period). Row 3 displays the coefficient of interest - the triple interaction of the post period indicator, the treatment group, and the indicator for a PCP visit within the past two years. Controls include patient age, age squared, gender, and Charlson comorbidity index, PCP military status, and PCP specialty. All regressions include individual, clinic-year, and cohort-year fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 6: Results - Patient Heterogeneity

	(1) Diabetic Patients	(2) Top Quartile Prior Year RVUs	(3) Pregnant Patients
Log RVUs	0.200***	0.102***	0.054*
	(0.058)	(0.023)	(0.030)
Log Total Visits	0.116***	0.075***	0.038*
	(0.040)	(0.016)	(0.019)
Probability of Specialty Care	0.043**	0.038***	0.023**
	(0.021)	(0.008)	(0.011)
Probability of Primary Care	0.042**	0.012	0.008
	(0.018)	(0.008)	(0.010)
Probability of ED Use	0.014	0.010	0.007
	(0.016)	(0.006)	(0.006)
Probability of Inpatient Admission	0.003	-0.000	0.004
	(0.010)	(0.004)	(0.005)
Sub-Sample Means			
RVUs	15.2	18.34	12.72
Total Visits	4.34	5.10	3.34
Probability of Specialty Care	0.49	0.61	0.51
Probability of Primary Care	0.62	0.59	0.46
Probability of ED Use	0.21	0.23	0.17
Probability of Inpatient Admission	0.06	0.08	0.11
Observations	104,088	700,080	556,896

Note: Results of estimating equation 1 on different sub-samples. Each row is a separate regression and displays the coefficient on the interaction of the treatment group identifier and the post period identifier for the listed dependent variable. Column 1 displays these results for patients who have a diagnosis of diabetes. Column 2 restricts the data for patients who were in the top quartile of RVUs in the year prior to the discontinuity. Column 3 restricts the data to patients who are diagnosed as pregnant in the year before or after the discontinuity. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 7: Impact on Specialty Encounters

	(1) Log of RVUs	(2) Number of Procedures
Panel A - Main Sample		
Discontinuity in Primary Care	0.018*	-0.002
	(0.010)	(0.013))
Sample Mean	2.04	1.89
Observations	1,341,619	1,341,619
Panel B - Existing Relationship		
Discontinuity in Primary Care	0.028*	0.009
	(0.016)	(0.019)
Sample Mean	1.95	1.89
Observations	647,576	647,576
Panel C - Multiple Comorbidities		
Discontinuity in Primary Care	0.134**	0.051
	(0.055)	(0.048)
Sample Mean	2.07	1.70
Observations	66,095	66,095

Note: Results of estimating a variation of equation 1 on outpatient specialty visits in Military clinics for the full sample of patients in the year before or after a discontinuity in care. Panel A shows results for all patients in the main sample. Panel B restricts the sample to encounters between a patient and a specialty provider who have at least one encounter together in the year prior to the analysis sample (i.e. 1-2 years before the deployment). Panel C restricts the sample to patients with multiple comorbidities as measured by the Charlson comorbidity index. Controls include patient age, age squared, gender, and Charlson comorbidity index. All regressions include provider, clinic-quarter, and specialty fixed effects. Standard errors are clustered by provider and clinic.

Table 8: Sample Construction

	Unique Patients	Observations	Unit of Observation
Full Data	1,605,996	26,950,914	quarter-year
Matched Sampled	533,127	6,930,651	quarter-year
Annual Sample	520,196	2,290,229	year
Relationship Sample	350,070	2,800,560	quarter-year

Note: This tables lists the sample size after each restriction is imposed. Column 1 includes the number of unique patients and column two lists the unit of observation. The full data includes all patients who met inclusion criteria. The matched sample includes all patients who undergo a provider deployment with sufficient pre and post data, and their matched controls. The annual sample cuts patient-years (relative to the discontinuity) in which the patient is only enrolled for part of the year. The Relationship Sample restricts the analysis to patients who had a visit with their assigned PCP in the two years prior to the deployment.

Table 9: Primary Care Provider Specialties

	Annual Sample	Relationship Sample	Unique Providers	Unique Deployers
Panel A: Direct Care Physicians				
Family Medicine	46.33%	45.68%	4,176	390
Internal Medicine	3.79%	3.68%	1,447	46
Aerospace Medicine	0.76%	0.62%	381	24
Panel B: Other Primary Care Providers				
Nurse Practitioner	21.96%	23.75%	1,227	25
Physician Assistant	17.59%	18.35%	1,478	83
Network Provider	8.25%	6.81%	unknown	0
Other	1.33%	1.11%	970	18

Note: This table shows the distribution of each type of provider specialty. Columns 1 and 2 are proportion of observations in each analysis sample. Columns three and four are at the unique provider level. While private sector network enrollment is observable, network providers are not individually identified in the data.

Table 10: Test of Deploying Provider Selection

	(1) Probability of PCP Deployer	(2) Probability PCP Deployer	(3) Probability PCP Deployer
Panel A - Patient Demographics			
Patient Age	0.001		
	(0.003)		
Patient Gender - Female	-0.000		
	(0.000)		
Comorbidity Score	-0.001		
	(0.001)		
Panel B - Provider Demographics			
PCP - Age		-0.001	
		(0.001)	
PCP Gender - Female		-0.016	
		(0.012)	
PCP Race - White		0.004	
		(0.014)	
Panel C - Provider Intensity			
Log RVUs			-0.001
			(0.002)
Prob. of a Referral			-0.001
			(0.002)
Num. of Procedures			0.003
			(0.003)
Joint F-Stat	0.96	0.82	0.11
Prob > F	0.42	0.48	0.96
Observations	6,143,643	6,142,219	6,143,643

Note: Coefficients from joint test of significance using military providers only. Note there are 6 military providers for whom I am missing the date of birth. In the provider demographic (column 3) regression I drop the 1,424 encounters these providers see. None of these 6 providers are deployers.

Table 11: Main Results - Relationship Sample

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Discontinuous Care	0.047***	0.033***	0.017***	0.010**
	(0.013)	(0.009)	(0.005)	(0.004)
Sample Mean	7.92	2.39	0.33	0.43
Observations	2,800,560	2,800,560	2,800,560	2,800,560

Note: Results of estimating equation 1 for utilization measures on the relationship sample. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 12: RVU Decomposition

	(1) Total RVUs	(2) Specialty Care RVUs	(3) Primary Care RVUs	(4) Emergency Dept RVUs
Observed	0.301*	0.135	0.08**	0.073*
	(0.175)	(0.154)	(0.035)	(0.0038)
Windsorized	0.285***	0.111*	0.048**	0.013
	(0.101)	(0.058)	(0.016)	(0.010)
Mean Observed	7.92	3.80	1.70	1.45
Mean Windsorized	6.61	2.68	1.41	0.47
Observations	2,800,560	2,800,560	2,800,560	2,800,560

Note: Results of estimating equation 1 on the total relative value units (RVUs) generated in each type of care. Row 1 displays results of the treatment group identifier and the post period identifier using the observed values and row 2 displays the same coefficient after windsorizing the count of RVUs at the 95th percentile of the distribution. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 13: Results - Relationship Heterogeneity

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Panel A - Most Recent PCP Visit				
$\leq$ 6 Months	0.057***	0.047***	0.026***	0.021***
	(0.019)	(0.013)	(0.009)	(0.007)
Sample Mean	9.59	3.00	0.39	0.56
Observations	942,168	942,168	942,168	942,168
6-12 Months	0.049**	0.032**	0.015*	0.009
	(0.022)	(0.014)	(0.008)	(0.007)
Sample Mean	8.30	2.51	0.35	0.46
Observations	783,216	783,216	783,216	783,216
$>$ 12 Months	0.037*	0.024**	0.009	0.007
	(0.019)	(0.012)	(0.007)	(0.005)
Sample Mean	6.18	1.77	0.26	0.31
Observations	1,075,176	1,075,176	1,075,176	1,075,176
Panel B - Length of Relationship				
Bottom Quartile	0.058***	0.048***	0.017**	0.017***
	(0.017)	(0.018)	(0.007)	(0.006)
Sample Mean	8.18	2.47	0.34	0.45
Observations	1,311,272	1,311,272	1,311,272	1,311,272
Top Quartile	0.077***	0.041***	0.023**	0.012
	(0.026)	(0.012)	(0.012)	(0.010)
Sample Mean	7.02	2.10	0.30	0.39
Observations	559,360	559,360	559,360	559,360

Note: Results of estimating equation 1 on different sub-samples related to the strength of the patient-PCP relationship. Each row is a separate regression and displays the coefficient on the interaction of the treatment group identifier and the post period identifier for the listed dependent subgroup. Panel A restricts patients based on length of time between the patient's most recent visit to the PCP and the discontinuity in care. Panel B restricts the data based on the total length of the relationship from start until the PCP deployment. Note some patients changed PCP's prior to the deployment quarter. Quartiles are estimated based on the distribution of length of relationship and can include more than 25% of the data. Bottom quartile is less than or equal to 2.25 years at the time of the deployment, and the top quartile are relationships greater than or equal to 3 years at the time of the deployments. The range of relationships is 2-5 years. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 14: Subgroup Proportions - Relationship Sample

	Control Group	Treatment Group
Diabetic	3.72%	3.69%
Top Quartile Prior Year Spending	24.97%	25.94%
Pregnant	19.87%	20.44%
Top Quartile Relationship	20.07%	16.39%
Bottom Quartile Relationship	46.80%	47.53%
PCP Visit $\leq$ 6 Months before Event	33.58%	35.90%
PCP Visit 6-12 months before Event	27.91%	30.05%
PCP Visit > 1 Year before Event	38.51%	34.05%
Non-Movers	58.47%	57.29%
Direct Care Only	87.51%	91.10%
Physician PCP Only	43.52%	51.66%

Note: This table presents the distribution of each subsample analyzed in the patient heterogeneity or robustness portion of the paper. Percentages are based on the number of patient-quarters in the relationship sample.

Table 15: Specialty Referrals

	(1) New Specialty Encounter	(2) Single Specialty Encounter
Discontinuous Care	0.009***	0.004**
	(0.002)	(0.002)
Sample Mean	0.084	0.040
Observations	2,800,560	2,800,560
2		

Note: Results of estimating equation 1 on the probability of each outcome on the relationship sample. Column 1 displays the probability that a patient is referred to a specialty clinic where the patient was not observed in the year prior to the analysis sample (i.e. 1-2 years prior to the discontinuity). Column 2 displays the probability that a patient was only seen once in a specialty clinic in the year before or after the discontinuity. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 16: New Specialty Decomposition

Specialty	(1) Full Sample	(2) Pregnant Sample	(3) Excluding Pregnancies
OB/GYN	0.003***	0.006**	0.0024**
	(0.001)	(0.003)	(0.001)
Mean	0.015	0.032	0.012
General Medical	0.003***	0.006**	0.0026*
	(0.001)	(0.003)	(0.001)
Mean	0.024	0.032	0.023
Surgical Services	0.001	0.001	-0.001
	(0.001)	(0.001)	(0.002)
Mean	0.014	0.012	0.014
Orthopedics	0.000	-0.000	0.001
	(0.001)	(0.001)	(0.001)
Mean	0.007	0.006	0.007
Mental Health	0.001	0.001	0.001
	(0.001)	(0.002)	(0.001)
Mean	0.010	0.014	0.008

Note: Results of estimating equation 1 on the probability of a new specialty referral, for each type of specialty. Column 1 displays the results for the full relationship sample. Column 2 displays the results limiting the data to patients who are diagnosed as pregnant in the year before or after the discontinuity. Column 3 limits the data to patients who are not observed as pregnant in the year before or after the discontinuity. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 17: Robustness - PCP Fixed Effect

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Discontinuous Care	0.033**	0.023**	0.012***	0.010**
	(0.017)	(0.010)	(0.004)	(0.005)
Sample Mean	7.92	2.39	0.33	0.43
Observations	2,800,560	2,800,560	2,800,560	2,800,560

Note: Results of estimating a variation of equation 1 on the relationship sample. Dis- continuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. All regressions include controls for patient age, age squared, gender, and Charlson comorbidity index and provider, clinic-quarter, and cohort-quarter effects. Network providers are coded with a single dummy ID. Standard errors are clustered by provider and clinic.

Table 18: Robustness Checks - Sample Restrictions

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Panel A - Military PCP Only				
Discontinuous Care	0.037**	0.027***	0.017***	0.009
	(0.014)	(0.010)	(0.006)	(0.005)
Sample Mean	7.81	2.30	0.33	0.41
Observations	774,179	774,179	774,179	774,179
Panel B - Direct Care Only				
Discontinuous Care	0.039***	0.028***	0.016***	0.010**
	(0.013)	(0.009)	(0.005)	(0.004)
Sample Mean	7.67	2.33	0.33	0.43
Observations	2,453,479	2,453,479	2,453,479	2,453,479
Panel C - Physician PCP Only				
Discontinuous Care	0.051***	0.034***	0.014*	0.011*
	(0.018)	(0.012)	(0.007)	(0.006)
Sample Mean	7.82	2.35	0.32	0.42
Observations	1,223,809	1,223,809	1,223,809	1,223,809
Panel D - Non-Movers Only				
Discontinuous Care	0.032**	0.024**	0.015**	0.004
	(0.015)	(0.009)	(0.006)	(0.005)
Sample Mean	7.11	2.16	0.30	0.38
Observations	1,636,520	1,636,520	1,636,520	1,636,520

Note: Results of estimating equation 1 on different subsets of the data. Each panel is a separate regression and displays the coefficient on the interaction of the treatment group identifier and the post period identifier for the dependent variable listed in the column header. Panel A displays these results for patients who have a military provider in the pre and post period. Panel B restricts the data to patients whose assigned PCP is in a military clinic and excludes anyone who has a civilian network PCP. Panel C restricts the sample patients who have a physician PCP in the pre and post period. Panel D restricts the sample to patients who remain enrolled in the same clinic and do not move during the sample period. Controls include, PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 19: Results of Varying Depoloyment Length Restrictions

	(1) Log of RVUs	(2) Log of Outpatient Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Bottom Quartile	0.031**	0.016	0.010	0.005
	(0.015)	(0.011)	(0.007)	(0.007)
Top Quartile	0.071***	0.048***	0.021***	0.018**
	(0.018)	(0.012)	(0.007)	(0.007)

Note: Results of estimating equation 1 using different deployment length restrictions. The Bottom quartile includes deployments between 1-5 months. The top Quartile includes deployments that last at least 8 months. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.

Table 20: Results of Loss of Electronic Medical Record

	(1) Log of RVUs	(2) Log of Total Visits	(3) Probability of Specialist Visit	(4) Probability of Primary Care Visit
Post-Period*Treatment Group	0.032***	0.023***	0.013***	0.008
	(0.010)	(0.007)	(0.004)	(0.004)
Post-Period*Network PCP -0.078***	-0.036***	-0.005	-0.032***	
	(0.018)	(0.012)	(0.004)	(0.009)
Post*Treatment*Network	0.085*	0.045	0.028	-0.000
	(0.048)	(0.031)	(0.023)	(0.023)
Observations	2,800,560	2,800,560	2,800,560	2,800,560

Note: Results of estimating a variation of equation 2 on the relationship sample and substituting an indicator for a network provider for the relationship indicator. The first row displays the interaction of the treatment group identifier and the post period identifier. The second row displays the interaction of the post period identifier and an indicator for whether a patient had a network provider in a quarter-year. Row three displays the coefficient of interest - the triple interaction of the post period indicator, the treatment group, and the indicator for a network provider. Controls include PCP military status, and PCP specialty. All regressions include individual, clinic-quarter, and cohort-quarter fixed effects as described in the text. Standard errors are clustered by patient and clinic.