

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

By STEPHEN D SCHWAB *

JOB MARKET PAPER

Draft: September 2, 2022

When a member of a work-team leaves, some information 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 12% increase in costs driven by an increase in the use and intensity of specialty care. Overall, I find significant disruptions in care even in a context in which significant investments have been made in knowledge-management systems. This has implications for how organizations allocate tasks and manage turnover.

* Schwab: Baylor University Department of Economics, stephen.schwab@baylor.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, 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 information is costly and that the division of labor increases these coordination costs. This is further complicated by the fact that information varies in its transferability. While general information can be easily distributed across an organization, more specific information exists only within individual agents. If these agents leave or are replaced, this latter information is lost to the organization. The extent to which information loss affects individual worker productivity is unknown. In this paper I empirically estimate the impact of information loss on organizational productivity. Using quasi-random turnover of primary-care providers (PCP) within the Military Health System, I estimate the causal impact of loss of information 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. 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 case-specific information as products and services are tailored to the individual patient.¹ After the loss of a PCP, some information about the patient's particular case may be lost, potentially decreasing the probability that the new provider can diagnose a condition and increasing the probability of a referral to specialty care. Because specialty care is more expensive than primary care, this is likely to increase total costs while making the marginal referral to a specialist less appropriate.

¹Similar challenges exist in the legal industry and enterprise sales where products and services are tailored to the client.

A particular challenge of estimating the effects of information loss 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 both of these issues through a unique data source - the Military Health System (MHS). The MHS provides a platform for patient care as well as to maintain the medical skills of 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 information loss on patients. An additional challenge is isolating the impact of information loss 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 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 coordinating a patient's care. To put this in context, consider an anecdote chronicled by a general internist (Press (2014))². 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.

²Agha et al (2019) also relay this anecdote

Overall, I find that a primary care provider deployment leads to a 12% (\$186) increase in a patient’s total cost of care in the following year, and that this increase is mostly driven by a 1.9 percentage point (6%) increase in the probability of a referral to a more specialized clinician. Consistent with a Roy model where each new referral is less well-targeted than the previous one, I find that patients are 20% (0.004 percentage points) more likely to only see the specialist once after a discontinuity. I contrast this with a sub-sample of patients who have been receiving care from multiple specialists for at least one year. Visits for these patients become approximately 5% more expensive after the loss of their primary care provider driven by an increased number of procedures for patients that require more coordination.

I provide support that information loss is the mechanism through a triple differences approach. I compare the main analysis with patients whose PCP deployed, but who had never had an appointment with that particular provider either because they did not use much primary care or happened to see another provider. I find the main effects are completely driven by patients with a relationship with their PCP. Finally, I use provider deployments as an instrumental variable in order to estimate the local average treatment effect (LATE) of a provider discontinuity on patient utilization. While deployments offer an average treatment effect, there is considerable heterogeneity in the effect across patients. I find that there is little effect on healthy patients but that there are much larger effects on those who require coordination of care. For instance in the year after a discontinuity, the total cost of care increases by 29% for patients with diabetes.

This research primarily contributes to two strands of literature. First, it contributes to the literature regarding the role of information in the firm production function. Transferring information 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). Information, however, varies in its cost to transmit (Jensen and Meckling (1992)). Some ”gen-

eral” 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 information to differences in firm-specific (Huckman and Pisano (2006)) and customer-specific (Clark, Huckman and Staats (2013)) performance. However, there is little empirical economic work estimating the value of specific information or disentangling it from other disruptions associated with turnover.

Estimating the value of specific information is critical as this information can theoretically be captured and communicated if the organization is willing to absorb the cost. This is in sharp contrast to tacit knowledge (Polanyi (1958)) such as individual skill which is infeasible to transfer outside of large-scale training programs. I offer empirical evidence of the marginal cost of losing specific information at a high level of investment (e.g. large-scale electronic health records). 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, especially in project-based teams. 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 work also contributes to a literature detailing the challenges of coordination in health care. Previous studies have shown that discontinuity in care increases health care costs (Sabety, Jena and Barnett (2021)), while others have

shown closer integration helps mitigate these challenges (David, Rawley and Polsky (2013); Agha, Ericson and Zhao (2020)); I complement these works by showing that discontinuity in care increases health care costs even in a fully integrated organization. Furthermore, I show these costs are not just driven by patients choosing more expensive sites of care, but that there are increases in the intensity of individual encounters for some patients. The Military Health System is comparable across many dimensions, including variation in care, to the civilian health care system (Bond and Schwab (2019)) and has been previously used in the economics literature to study the impact of defensive medicine (Frakes and Gruber (2019)). Using data from the MHS provides several advantages. First, previous work on discontinuity has been hampered by confounders associated with the loss of a primary care provider. For instance, patients generally change practices or forego finding a new primary care provider altogether (Sabety (2020)). New practices may have a different mix of payers offering different incentives for care. In the military setting, patients must enroll with a new PCP, all providers face the same incentive scheme, and insurance rules prevent substitution to a specialist without a PCP referral. The MHS also lets me consider the effects on a working age population, a group that has been typically understudied in the health economics literature due to a lack of publicly available data.

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

II. Conceptual Framework

When the cost of acquiring knowledge is expensive, the division of labor lets workers focus on different problem sets. I build on Garicano’s 2000 model of an organization as a partition of workers into L classes where each class has a discrete (potentially overlapping) knowledge set. Workers pay a “helping cost” in assisting other workers with a problem. The organizational challenge is to match

a problem with the worker who has the appropriate knowledge (Garicano (2000)).

I differ from Garicano in two ways. First, I break a problem into two component stages. In the first "diagnosis" stage, a worker must identify what the root cause of the problem is and in the second "treatment" stage they must fix the problem. Second, and much more substantially, I assume that an individual worker's ability to diagnose a problem is related to that particular worker's past experience and cannot be transferred to other workers. In other words, helping costs within a partition are infinite. A worker passes a problem to the next stage if they lack the requisite knowledge to either diagnose or solve a problem. This implies that some problems that are passed up the hierarchy due to inability to diagnose could be fixed at the lower level.

For simplicity I consider a two-tiered hierarchy. The first layer of the organization are the general problem-solvers. These workers decide either solve a problem or pass it up to the second level. The second layer are the specialized problem solvers. These workers either solve the problem if they have the requisite knowledge or attempt to gather new information to eventually solve the problem. In a health care setting this model reflects the relationship between primary care providers and specialists.

A. The Primary Care Provider

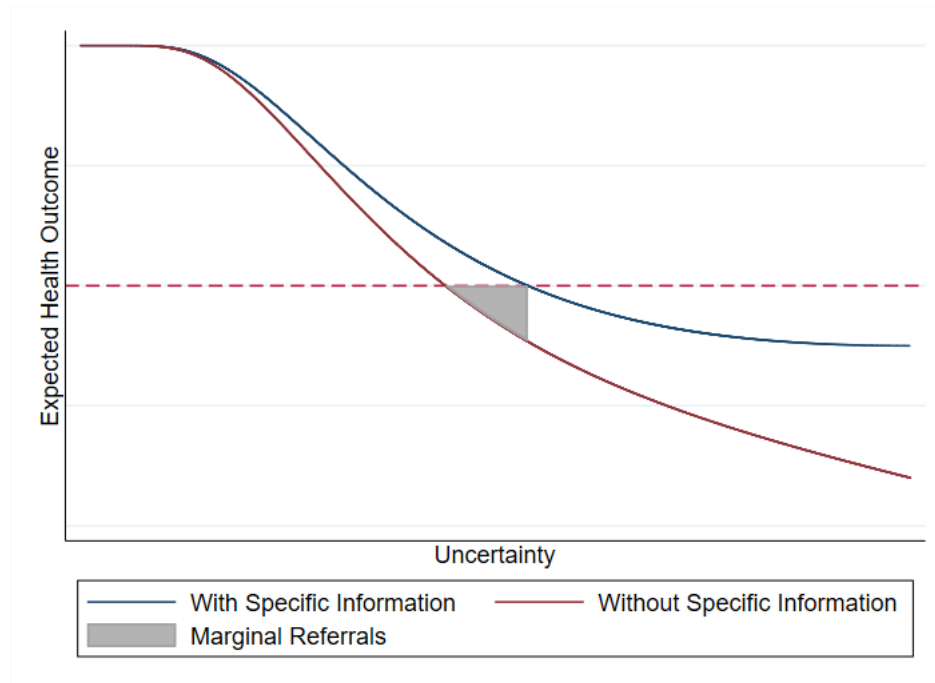
At the onset of illness or injury, the patient sees his PCP. The PCP faces the decision of whether to treat the patient, or refer to a specialist and coordinate the patient's care. Note that coordinating care is expensive for the PCP, as this time is not separately compensated. In order to treat the patient, however, the PCP must meet two necessary conditions. First, the PCP must diagnose the patient. Second, he must have the requisite skill to treat the patient's illness. If either of these conditions are not met, the PCP refers the patient to a specialist and transitions to a coordination role. While the former is dependent on the level of available information, the latter is independent of any case-specific information.

The PCP begins the appointment trying to learn about the patient's condition. Knowledge, however, is dynamic. The primary care provider can not only access information gathered during the current visit, but can also access at least a portion of information previously gathered. The more generalizable that information, the more likely the provider will have access to it regardless of a past interpersonal relationship with the patient. An example would be the results of an x-ray or patient's blood pressure at a point in time as these are both likely documented in the patient's medical record. On the other hand, specific information encompasses "the idiosyncrasies of time and place" (Hayek (1945)) and includes that which is more costly to transfer such as personal information about the patient's preferences such as their propensity to seek care and pain tolerance, but also includes previously-learned knowledge such as how much to trust each other and how to effectively communicate with each other. This type of information is generally only available within an existing relationship. This implies that the knowledge available to a provider is a function of both the visits the patient has had within the health system and the quantity of those visits that were with the current provider. After a discontinuity, this proportion goes to zero, decreasing the probability that the provider can diagnose the illness and increasing the probability of a specialty referral. Because specialty care is more expensive than primary care, this is likely to increase total costs while also making the marginal referral to a specialist less appropriate.

Figure 1 depicts these effects. The Y access denotes the expected health outcomes of a distribution of patients for a given level of information. The X access denotes levels of uncertainty in the appropriate treatment for the specific patient's condition. As uncertainty increases, information has a greater impact on the expected health outcome. The horizontal dotted line reflects a provider-specific threshold to refer a patient to specialty care. A PCP that does not believe she can achieve this a health outcome above this threshold will refer the patient to a specialist. The marginal referral triangle are patients who would have been

treated by the PCP with more information, but are instead referred to a specialist.

While I present this conceptualization as provider-focused, health care decisions are often framed in a 'shared decision-making' framework (Elwyn et al. (2012)) where the provider provides information and the patient elects his preferred care decision. While I focus on the PCP as the decision-maker for simplicity, the shared decision-making framework is a simple modification that yields the same predictions. For instance, the decision to refer would be the minimum expected health outcome of the patient's beliefs and the provider's beliefs.



Note: Author's depiction of the relationship between treatment uncertainty and information. The dotted line reflects a primary care provider's (PCP's) threshold for referring a patient to specialty care. Patients whose expected outcome is above this line are treated by the PCP. Patients whose expected outcome is below this line are referred to a specialist. The grey box denotes patients that are treated by the PCP when information is available, but referred to a specialist without this information.

Figure 1. : Information as an Input in Production of Health

B. The Specialist Clinicians

Once a PCP refers the patient, one or more specialist clinicians join the care team. The specialist, like the PCP, must first diagnose and then treat the patient. However, her choices are constrained to more or less intensive care. I assume the specialist will only provide the intensity of care needed to cure the patient. At the onset of the visit, the specialist seeks information regarding the patient's illness. In addition to the general information available to the whole care team, each clinician has access to the specific information gathered through their individual encounters, along with any other information transferred through coordination between the specialty and primary care providers. After a discontinuity, these discussions offer less information. This implies opposing effects. Specialists will treat appropriate referrals more intensively after a discontinuity in primary care. However, because the marginal referral is less appropriate, new referrals are likely to be treated less intensively on average.

III. 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 retirees³, 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 individuals, but does not cover care delivered in a war zone.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.

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

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

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.

A. 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 instance⁵: 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.

B. 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 military dependents enrolled in Tricare from 2008-

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

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 comorbidities within the medical claims. The full sample consists of 1,553,271 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 Contingency Tracking System (CTS) and military personnel master files provided by the Defense Manpower Data Center (Defense Manpower Data Center (2007-2017a); Defense Manpower Data Center (2007-2017b)). 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 claims data, I construct a number of utilization and intensity of care measures. Primarily, I focus on the cost of care following a discontinuity. Military claims do not include prices so I apply Medicare rules to these claims with two major exceptions. First, I don't geographically adjust so that costs are not driven by regional price differences. Second, I don't include medical equipment or pharmaceuticals as I want to capture provider workload. Equipment can be very expensive but not indicative of the intensity of care. Average spending is about \$1600/year.

Second, I consider the types of care a patient uses each quarter. 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 within between appointments. I measure

this using the total cost an appointment, as well as the specific procedures listed in each visit as coded using the Common Procedural Terminology (CPT) system. CPT codes are primarily used for billing and include up to three evaluation and management (E&M) codes and up to 10 non-E&M procedure codes. E&M codes are meant to denote the discussion and decision-making portion of the visit, while procedure codes are meant to capture anything else including tests and procedures performed. E&M codes have five levels of complexity. I dichotomize the complexity of a visit by coding E&M levels 4 and 5 as "complex."

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. Following Deshpande and Li () 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 whose PCP deploys one year later⁶, have been enrolled for at least one year with the same provider, and who remain enrolled in Tricare Prime for at least two more years (i.e. one year after the PCP deploys). I then match

⁶In order to observe any pre-treatment effects, I only consider patients as treated if they were assigned to the provider one year prior to the deployment date. I chose one year because this is an upper bound on when a PCP may learn about an impending deployment.

to the control group based on the length of relationship (in quarter-years) with their assigned PCP 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.

In the final step, I drop two groups of patients. First, 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. I use deployments that are one to three months as a falsification test later in the paper. I also drop patients who do not have an appointment with their assigned PCP between one year prior to the match quarter and the deployment quarter (two years overall), as by definition there is little specific information to lose ⁷. I return to this group in a triple differences analysis that provides further support for the information loss mechanism. I also relax these restrictions in the instrumental variables approach since 'non-compliers' should not affect the local average treatment effect. The final sample consists of 361,765 individuals over 3,255,885 patient-quarters. Table 1 lists the sample size reductions from each step.

Empirical Design

My primary analysis uses a difference in differences approach that treats a physician deployment as a discrete event. Because the treatments occur at differ-

⁷They may have *some* specific information from coordinating care, but likely much less

Table 1—: Sample Construction

	Unique Patients	Patient-Quarters
Full Data	1,553,271	26,461,009
Matched Sampled	548,454	4,936,086
After Short-Deployment Restriction	536,451	4,828,059
After PCP Relationship Restriction	361,765	3,255,885

Note: This tables lists the sample size after each restriction is imposed. Difference in differences and triple differences samples will have fewer observations due to dropping the treatment quarter in these regressions.

ent time periods for different individuals I use a stacked difference in differences model with a cohort fixed effect. I estimate the models using ordinary least squares. Because spending is highly skewed, I conduct a log transformation, adding one dollar to each quarter in order to deal with any zero-spending. I show the results are robust to alternate transformations and distributional assumptions later in the paper.

Identification Strategy

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

$$(1) \quad Y_{ipjt} = \beta_{Q=t-t^*} + \gamma^* X_{pt} + \theta_i + \delta_{jt} + \epsilon_{ipjt}$$

where y_{ipjt} represents an outcome of interest for individual i , assigned to PCP p , part of cohort j at time t and Q is a series of time dummies for each quarter-year relative to the to the quarter-year of the physician deployment t^* . $\beta_{Q=t-t^*}$ represents the relative difference between the treatment and control groups at each point in time. X_{pt} is a vector of two time-varying controls. First, I control for whether the the patient's PCP in time t is a physician or an advanced practice provider (APP) such as a Nurse Practitioner or Physician Assistant. It's possible that moving from a physician to an APP could lead to more referrals biasing the results since these providers could have less specialized skill. Second, I control for whether the patient's assigned PCP is active duty military since active duty providers may differ from civilian providers in observable ways. I return to these provider differences in the alternative explanations section of the paper. Note that most standard patient-level controls are collinear with the fixed effects given that I've restricted the analysis to two years of data. θ_i represents a vector of individual fixed-effects and δ_{jt} represent a cohort - quarter year interacted fixed effect. Because the cohort is constructed using both the timing of the discontinuity and the existing relationship length, this cohort-time fixed effect controls for any differential trends based on the timing of treatment or previous discontinuities. Finally, I cluster the standard errors by the patient to account for any serial correlation of the error terms.

In the main analysis, I estimate the average treatment effect of a discontinuity in care using equation 2 below.

$$(2) \quad Y_{ipjt} = \alpha + \beta_1^* I(t > deployment) + \gamma^* X_{pt} + \theta_i + \delta_{jt} + \epsilon_{ipjt}$$

where I is an indicator for the treatment group and all else is the same as in equation 1. β_1 is the coefficient of interest that indicates the relative change for the treatment group after the provider deploys. Because I observe the quarter,

but not the exact date, of the deployment, I omit the deployment quarter from these regressions.

In order to provide evidence that loss of information is the mechanism, 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, 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 3.

$$(3) \quad Y_{ipjt} = \alpha + \beta_1^* I(t > \text{deployment}) + \beta_2 * (t > \text{deployment}) * (\text{visits} > 0) \\ \beta_3^* I(t \geq \text{deployment}) * (\text{visits} > 0) + \gamma^* X_{pt} + \theta_i + \delta_{jt} + \epsilon_{ipjt}$$

β_1 in this equation is the effect of the treatment group in the post period. β_2 is the effect of those who have met their provider in the post period compared to those who never met their provider. Finally coefficient β_3 indicates the triple interaction of the treatment indicator, the post-period indicator and having met the provider. This can be interpreted as the relative effect of the deployment for those who had some specific information to lose compared to those whose provider deployed, but didn't have any information to lose.

Finally, I consider what occurs within individual encounters using the following equation:

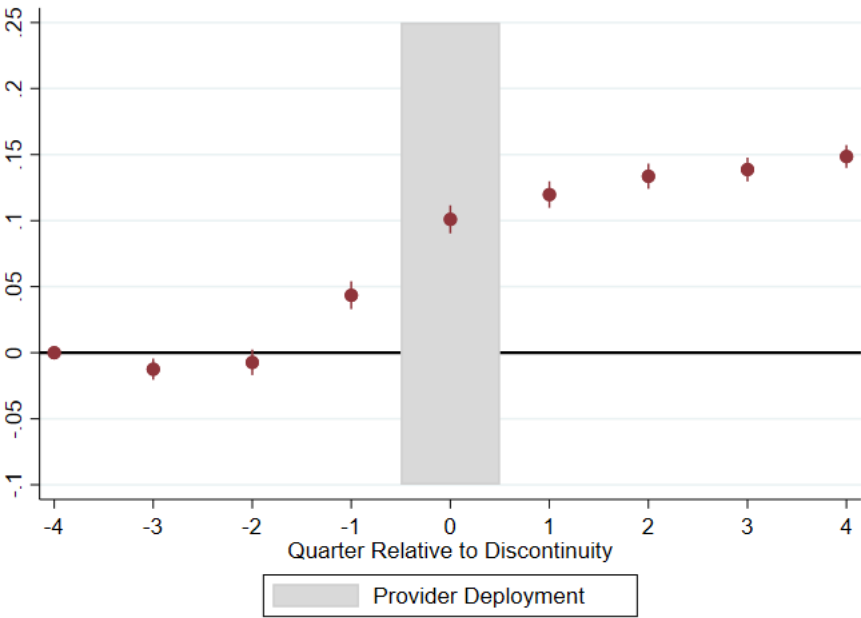
$$(4) \quad Y_{ipjt} = \alpha + \beta_1^* I(t > \text{deployment}) + \gamma^* X_{it} + \theta_p + \delta_{jt} + \epsilon_{ipjt}$$

The index p here refers to the encounter provider rather than the assigned PCP. This is similar to equation 2, however I swap the patient fixed-effect for a provider fixed effect so that we identify differences in how the provider treats these patients compared to the provider’s other patients. This eliminates any differences driven by provider specialty or skill level. Because I no longer include the patient fixed effect I add patient covariates including age, age squared, gender, and charlson comorbidity score and cluster standard errors at the provider level.

Results

I begin this section by presenting evidence that provider deployments do in fact lead to changes in a patient’s primary care provider. I test for this using equation 1 with the outcome being an indicator for having a different primary care provider than in the matching period. Figure 2 shows an event-study of the probability of having changed primary care providers since the matching index quarter in the periods around a deployment. The slight decrease two to three quarters prior to the deployment are a function of the deployment process. Physicians who have been assigned to deploy cannot be moved between hospitals reducing the chance for a discontinuity based on a provider moving. This means a patient assigned to this physician is *slightly* less likely to have a discontinuity in care from the time of assignment until right before the provider deploys. We then observe an uptick in changing primary care providers in the quarter prior to the deployment. This may be due to observing the date that a provider arrives overseas but not the date they left the clinic.

In order to interpret the results as the affect of a discontinuity in care, the deployments must be quasi-random and must only affect outcomes through their impact on discontinuity 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. Table 2 panel A displays the results of a regression of patient demographics in the year prior to the match-



Note: Event study of the combined probability that a patient has changed primary care providers at any point since the match period using equation 1. This is the first stage Each point is the coefficient for the interaction of the elapsed-time quarter and treatment group indicator for quarter-years relative to the provider deployment. The match quarter is the omitted quarter.

Figure 2. : Event Study of Probability of Different Primary Care Provider than Match Period.

quarter on an indicator for undergoing a primary care provider deployment. The regression does not include any controls but does include clinic and cohort-time fixed effects.

The treatment group tends to be about a year and a half older, has fractional more comorbidities, and tilts a bit more female. Panel B shows differences in utilization using the same regression. While there are small differences, none of these are economically meaningful and all are explained by the observable demographic differences. Panel C repeats the utilization regression but adds in controls for age, (linear and quadratic) gender and Charlson score. There are no economic or statistically significant differences between the groups.

Table 3 shows the results of a series of regressions of practice intensity. Column 2

Table 2—: Comparison of Patients in the Analysis Sample Who Do and Do Not Undergo A Deployment Related Discontinuity in Primary Care

	(1) Sample Mean	(2) Coefficient On Treated	(3) Std Error
<i>Panel A - Demographics</i>			
Age	30.25	1.524**	(0.086)
Gender - Female	0.92	0.017***	(0.003)
Charlson Comorbidity Score	0.18	0.021***	(0.005)
<i>Panel B - Annual Utilization</i>			
Total Spend	\$1,586	\$58.59	(\$65.28)
Specialist Visits	5.34	0.282***	(0.10)
Primary care Visit	3.89	0.104**	(0.048)
Emergency Dept Visits	1.31	0.001	(0.037)
Inpatient Admissions	0.17	0.001	(0.005)
<i>Panel C - Annual Utilization with Patient Controls</i>			
Total Spend	\$1,586	-\$8.02	(\$62.79)
Specialist Visits	5.34	0.03	(0.098)
Primary care Visit	3.89	-0.02	(0.046)
Emergency Dept Visits	1.31	0.033	(0.036)
Inpatient Admissions	0.17	0.001	(0.005)
<i>N - Unique Patients</i>			
Control Group	353,178		
Treatment Group	8,587		

Note: Coefficients from a regression of the dependent variable on an indicator for whether the patient is in the treatment group. Panel A&B include a clinic and cohort-time fixed-effect with no controls. Panel C also includes demographic controls including age, age squared, gender, and Charlson Comorbidity Score. Differences are estimated for the year prior to the match-quarter.

displays the results of a regression on an indicator for whether a provider deploys within the sample window. The regressions don't include any controls but do include quarter-year fixed-effects. Deployers and non-deployers appear to practice in similar fashion, with only a small difference in the number of procedures coded on the claim. Both groups spend about the same per primary care appointment and refer to specialists at about the same rate. Column 3 restricts the analysis to primary care providers that deploy and those that gain patients after the provider deployments. I find no significant differences between the providers along any of

the three margins. Finally, I check for changes in practice patterns by modifying equation 4. I regress each appointment intensity indicator on an indicator for being within one year of a deployment and include provider and match cohort fixed effects. Column 4 shows the results of these regressions. There's no evidence of an economically meaningful change in practice patterns for any of the measures, with only the number of coded procedures being statistically significant.

Table 3—: Primary Care Encounter Intensity

	(1) Sample Mean	(2) Coefficient on Deployer	(3) Coefficient on New PCP	4 Anticipatory Period
Cost	\$59.45	-\$0.31 (0.530)	0.98 (0.808)	-0.01 (0.010)
Prob. of a Referral	0.05	-0.000 (0.0001)	0.001 (0.002)	0.02 (0.003)
Num. of Procedures	1.12	0.019** (0.010)	-0.007 (0.014)	-0.021** (0.008)
Total PCP's	6,341			
Deploying PCP's	410			
Gaining PCP's	1,813			
Encounters -				
Total	3,790,252			
Losing & Gaining	1,714,400			

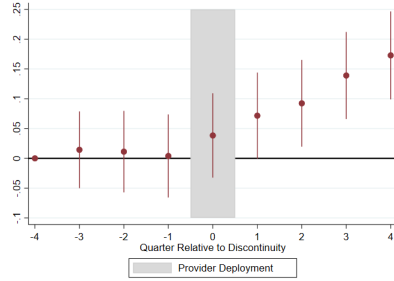
Note: Coefficients from a regression of the dependent variable on an indicator for whether the provider deploys in the analysis sample and a clinic and match cohort fixed-effect as described in the text. Cost per encounter is based on Medicare rates applied to the Military Direct Care System. Probability of a referral is calculated by assigning any specialty visit to the most recent primary care visit if the visit occurred within 1 year. Number of procedures is the total number of Common Procedural Technology (CPT) codes listed on the medical claim. Regression includes any primary care encounter by a primary care provider who has at least one patient from his or her panel in the analysis sample during the match period. For columns 2 and 3 I restrict the analysis to primary care office visits with patients who are not in the treatment group. In column 4 I restrict the regressions to encounters prior to a provider deployment and include the following patient controls: age age squared, gender, and Charlson Comorbidity Score.

Next, I present evidence of parallel pretrends in patient cost and utilization. Figure 3a shows changes in the natural log of total dollars spent per quarter across the year before and after the discontinuity. The estimates are stable in

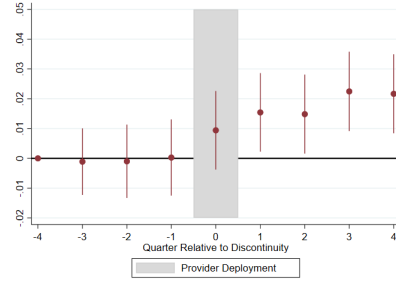
the quarters prior to the deployment indicating parallel pre-trends and providing support for the parallel trends assumption. Cost increases as the PCP departs indicating that there is an effect of the discontinuity on total cost of care. Figure 3b shows the probability of any specialty care appointment. The trend is similar to the spending outcome with a sharp increase after the provider leaves.

The linear probability model event study for primary care shown in 3c is interesting as the effects appear somewhat delayed. This is logical if discontinuities are unlikely to lead to a patient requiring primary care on the extensive margin. A discontinuity in care isn't likely to bring on an illness or injury, but conditional on that illness or injury a patient may be likely to use more or more expensive care, especially if there is a primary care follow-up after a specialty appointment. This is supported by the event study in figure 3d which shows the probability of using multiple primary care appointments in a quarter.

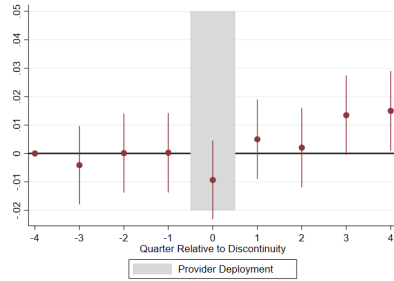
Table 4 presents the difference in differences results for each of the three main utilization variables. Column 1 shows the effect of a physician deployment on the log of spending in the post period. The coefficient indicates about a 12% percent or an approximately \$186 increase over the full year. Column 2 shows the effects of the discontinuity on the probability of using a specialty care appointment. The 1.9 percentage point increase on a base of 0.33 indicates that specialty care utilization goes up by about 6% consistent with the theoretical prediction. Column 3 shows a 1 percentage point increase in primary care. This is contrary to previous literature using Medicare data that has shown a drop in primary care after a discontinuity (Van Walraven et al. (2010); 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. While I estimate changes in spending using a log +1 transformation in the main results, I present an alternative hyperbolic sine transformation in appendix table A1 column 1. The specialty and primary care results are based on linear probability models, however, I allow for different distributional assumptions and present alternative poisson models for these



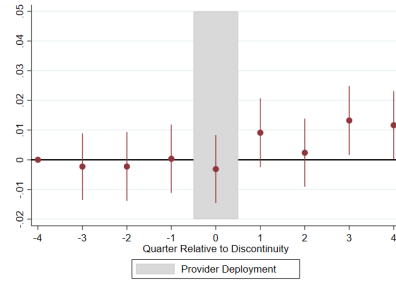
(a) Log of Spending



(b) Probability of a Specialty Care Visit



(c) Probability of a Primary Care Visit



(d) Probability of Multiple Primary Care Visits

Note:

Event studies display the coefficient on the interaction of the elapsed-time quarter and treatment group indicator for each quarter relative to the provider deployment as described in equation 1. The quarter prior to the deployment is the omitted quarter for all figures. 3a shows the change in the log of total spending. 3b displays the change in the probability of at least one specialty visit. 3c and 3d show the changes in the probability of at least one primary visit and at least two primary care visits respectively. Regressions include controls for whether the provider is a physician and whether the provider is active duty along with patient and cohort-time fixed effects. Standard errors are clustered by patient.

Figure 3. : Event Studies of Utilization Measures

outcomes in appendix table A1 columns two and three. Poisson models require some variation in the outcome, meaning that individuals who do not use that type of care over the two year period are dropped from the regressions. However, the magnitude of the estimates are substantially similar for specialty care and much larger for primary care - about a 6% increase in each. This provides further support that the effects on primary care are on the intensive margin rather than the extensive margin.

Table 4—: Main Results - Utilization

	(1) Log of Spending	(2) Probability of Specialty Encounter	(3) Probability of Primary Care Encounter
Discontinuous Care	0.111*** (0.022)	0.019*** (0.004)	0.010** (0.004)
Sample Mean	\$1,586	0.33	0.44
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
N	3,247,298	3,247,298	3,247,298

Note: Results of estimating equation two. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

Table 5 presents linear probability models for whether a patient ends up in either the emergency department or is admitted to a hospital as an inpatient. The coefficients on both are extremely small and statistically indistinguishable from zero. While I cannot definitively rule out a longer-term effect, I do not find evidence that the increased use of outpatient care is preventing any negative events.

While the provider deployment provides an *average treatment effect*, there is

Table 5—: Main Results - Outcomes

	(1) Probability of Emergency Department Encounter	(2) Probability of Inpatient Admission
Discontinuous Care	0.001 (0.003)	0.001 (0.001)
Sample Mean	0.12	0.04
Controls	Yes	Yes
Fixed Effects	Yes	Yes
N	3,247,298	3,247,298

Note: Results of estimating equation 2. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

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. In this section I consider the different effect across three groups of patients - diabetics, those who were above median spenders the previous year, and those who are pregnant. Table 6 displays the results of these regressions. Column 1 shows the results for diabetic patients for each of the three primary outcomes. Diabetic patients have their cost of care increase by 29% with large increases in both specialty and primary care. Column two displays the results for those who were above-median spending in the previous year. These patients likely had a health shock in the previous year and may be more reliant on their primary care for either treatment or coordination. While not as large an increase as diabetics, these patients have their costs go up about 19% and a three percentage point increase in the probability of a specialty care visit. This is slightly over 1.5 times the average effect on both margins. Column three displays the results for patients who were above median spending the year prior. These patients have slightly

Table 6—: Main Results - Patient Heterogeneity

	(1)	(2)	(3)
	Diabetic Patients	Above Median Prior Year Spending	Above Median 2nd Prior Year
Spending	0.255*** (0.104)	0.177*** (0.030)	0.131*** (0.031)
Specialty Care	0.052** (0.021)	0.030*** (0.006)	0.019*** (0.006)
Primary Care	0.041** (0.019)	0.014** (0.006)	0.019*** (0.005)
Mean Annual Spending	\$3,712	\$2,648	2,248
N	121,029	1,623,871	1,624,018

Note: Results of estimating equation 2 on different sub-samples. Each row displays the coefficient on the interaction of the treatment group identifier and the post period identifier for the listed dependent variable. Each column is a different subsample as described in the text. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

higher spending than average (though not significantly different), driven by a larger increase in primary care. In total, these results suggest that for patients who require significant coordination of care, the loss of a provider is extremely costly despite the use of an electronic medical record. This further implies that there is more to the production function than the patient, their medical records, and the provider.

Mechanism

In this section I present suggestive evidence that information loss is the mechanism driving the main results. I begin by presenting the results of a difference in difference in differences model that compares the effects of a provider deployments for those that have a recent (within two years of the deployment) visit with their PCP to those who do not. Table 7 displays the results of these regressions.

The coefficient on the triple interaction in equation 5 is the coefficient of interest. The estimates are slightly lower but substantially similar to then main difference in differences results.

Table 7—: Results of Triple Difference

	(1) Log of of Spending	(2) Probability of Specialty Encounter	(3) Probability of Primary Care Encounter
Post-Deployment	0.023 (0.027)	0.003 (0.005)	-0.002 (0.004)
Triple - Interaction	0.083** (0.035)	0.015*** (0.006)	0.010 (0.006)
Sample Mean	\$346.77	0.29	0.37
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
N	4,291,608	4,291,608	4,291,608

Note: Results of estimating equation three. Post-Deployment coefficient is the interaction of the treatment group identifier and the post period identifier. regardless of whether the patient has a relationship with the provider. Discontinuous relationship is the triple interaction of whether a patient who has a relationship with the provider, whether the provider deploys, and an indicator for the post-deployment period. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

To provide further support for information as the mechanism, I consider the effects on a specific group of patients. The military uses a civilian network of providers to augment military hospitals when there is a capacity constraint. Patients assigned to these primary care providers receive all of their primary care in the private sector which is not on the military's electronic health record. These patients should lose both specific and general information and have significantly higher effects than those who remain within the Military system. I estimate equa-

tion 1 but add in an interaction effect with whether someone is assigned to the civilian network. Table 8 displays the results of these regressions. Total costs increase by an additional 13% for those who are assigned to the network after a discontinuity in care. There are marginally significant additional increases in specialty care but no change in the increase in primary care. While some caution should be taken in interpretation as there could be endogeneity in the assignment to the civilian network, this provides additional suggestive evidence that information is a major input in the production function.

Table 8—: Results of Loss of Electronic Medical Record

	(1) Natural Log of Spending	(2) Probability of Specialty Care Encounter	(3) Probability of Primary Care Encounter
Main Results	0.111*** (0.022)	0.019*** (0.004)	0.010** (0.004)
Post Discontinuity Network Interaction	0.127*** (0.044)	0.014* (0.008)	-0.0036 (0.009)
Network PCP Mean	\$1,429	0.30	0.38
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
N	3,247,298	3,247,298	3,247,298

Note: Results of estimating equation 2 with an additional interaction term for receiving primary care in the civilian network. The first row shows the main results. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. Regression includes indicator for discontinuity in care and an indicator for network PCP in addition to the interaction term. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

Next I consider the impact on specialist encounters. The conceptual framework suggests that information loss could potentially manifest in two different ways. First, primary care provider referrals to specialists will be less appropriate. Second, appropriate specialist visits will be more intensive. For this analysis, I

restrict the sample to specialty appointments within military hospitals where we 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 are not paid for individually (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 also exclude visits that do not require evaluation and management (e.g. an appointment to have an MRI).

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. 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 2. Table 9 presents the results of these regressions. The coefficient in column 1 indicates a 9% (0.009 pp) relative increase in the probability of a visit in a new specialty clinic - about 1.5 times the magnitude of the overall increase in the probability of specialty care utilization. Column 2 indicates a 20% (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 and supports the prediction that much of the increase in specialty care is driven by less appropriate referrals.

I then consider each of these measures across the four different specialties with the most new visits in the data - Gynecology, General Surgery, Orthopedics, and Dermatology⁸. If the referrals are due to information loss, then areas with more

⁸I also analyzed Psychiatry and found no effect. I omit it from the table as it's less clear to me

Table 9—: Specialty Referrals

	(1) New Specialty Encounter	(2) Single Specialty Encounter
Discontinuous Care	0.009*** (0.002)	0.004** (0.002)
Sample Mean	0.10	0.02
Controls	Yes	Yes
Fixed Effects	Yes	Yes
N	3,247,298	3,247,298

Note: Results of estimating equation two. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Column one is a linear probability model indicating whether a patient was seen in a specialty clinic (e.g. dermatology, ob-gyn) in which they are not seen in the year prior to the match-quarter. Single encounters refers to the probability that a patient visits a specialty clinic with no follow-up visit in the data as of one-year after the deployment event. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

uncertainty should see larger increases. While there is no single consensus on the relative uncertainty across specialties, it is reasonable to assume that Gynecology and Dermatology have more uncertainty than orthopedics and general surgery given existing diagnostic technology (e.g. MRI's). Table 10 displays the results of these regressions. As expected obstetrics/gynecology has a large increase in both new visits and single visits. This is not driven by any change in the pregnancy-rate (analysis not shown). Surgery has a marginally significant increase in new visits and there is no effect on orthopedics. Perhaps the most interesting is dermatology which has a large increase in new visits. It's plausible that new PCP's notice skin marks that could potentially be cancerous. If these new visits are preventing skin cancer then they could be extremely valuable. Cancer is a rare diagnosis in this sample and econometric analysis is not powered to pick up any change. However, the large increase in single encounters likely implies that there is no follow-up and cancer was probably not found in any biopsies conducted due to the increase in referrals.

whether that is a high or low uncertainty specialty

Table 10—: Specialty Decomposition

Specialty	(1) New Visit	(2) Single Visit
OB/GYN	0.0022** (0.0009)	0.0018** (0.0007)
Mean	0.013	0.007
Surgery	0.0013* (0.0007)	0.0009 (0.0006)
Mean	0.009	0.007
Orthopedics	-0.0003 (0.0008)	-0.0004 (0.0004)
Mean	0.011	0.009
Dermatology	0.0015*** (0.0005)	0.0012*** (0.0004)
Mean	0.011	0.009

Note: Results of estimating equation two across different specialties. Each row shows the coefficient on the interaction of the treatment group identifier and the post period identifier for the dependent variable in each column restricted by specialty department. Column one is a linear probability model indicating whether a patient was seen in a specialty clinic in which they are not seen in the year prior to the match-quarter. Single encounters refers to the probability that a patient visits a specialty clinic with no follow-up visit in the data as of one-year after the deployment event. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.

Finally, I look across all specialty encounters for evidence of a change in intensity of care by the specialty provider. I use equation 4 to estimate any changes in the intensity of care along three measures. Primarily, I consider spending within the visit. Second, I consider whether the provider codes a more complex evaluation and management code (indicated by a 4 or 5 level code). Finally, I consider whether there is a change in the number of procedures conducted in the encounter. Table 11 panel A presents the results of these regressions applied to the full sample of specialty encounters. I find no evidence of an effect of the discontinuity in

primary care on these specialty encounters.

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. Second, I consider any specialist visit for patients who have encounters with multiple specialists in the year leading up to the match quarter. These patients are likely to require more coordination of care. Table 6 presents the results of these regressions in panels B & C respectively. Both groups have a similar increase in spending per visit, although this is imprecisely measured especially for the long-term relationship sample. However, there is a difference in what is driving the increased cost. For long-term relationships, this is driven by an 18% increase in the probability of a level 4 or level 5 evaluation and management code indicating that these providers are exerting more effort in the discussion portion of the visit. The results shown in panel C, however, suggest that for those who require more coordination, the cost is driven by an approximately 5% increase in the number of procedures performed. The difference between these groups makes sense for a couple reasons. One, patients returning to specialists whom they have been seeing for years may have already undergone the procedures that provider commonly performs. These specialist providers may end up providing some of the management the PCP was formerly providing, consistent with previous findings (Sabety (2020)). Two, patients with higher coordination needs may have more uncertainty in their diagnoses resulting in more tests and procedures when information is decreased.

Alternative Explanations

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 information loss. First, the loss of a primary care provider could simply make

Table 11—: Impact on Specialty Encounters

	(1) Log of Spending	(2) Complex Evaluation	(3) Number of Procedures
<i>Panel A - Main Sample</i>			
Discontinuity in Primary Care	0.009 (0.014)	0.003 (0.009)	0.011 (0.019)
Sample Mean	\$81.58	0.277	1.71
N	585,857	585,857	585,857
<i>Panel B - Existing Relationship</i>			
Discontinuity in Primary Care	0.047 (0.050)	0.054* (0.037)	-0.004 (0.066)
Sample Mean	\$78.28	0.298	1.68
N	139,622	139,622	139,622
<i>Panel C - Coordination Required</i>			
Discontinuity in Primary Care	0.054 (0.030)	0.001 (0.019)	0.080* (0.041)
Sample Mean	\$82.78	0.301	1.74
N	179,801	179,801	179,801
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes

Note: Results of estimating equation six. Observations are at the specialty encounter level for all patients in the analytical sample being seen for specialty care in the military direct care system on an outpatient basis. 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 leading up to the match quarter. Panel C restricts the sample to patients who have at least one encounter with two or more specialty providers in the year leading up to the match quarter.

all other providers in the clinic busier and less attentive to a focal patient's care. Second, the type of provider a patient is sent to could drive the results. Patients are more likely to be matched to either a civilian providers or an advanced practice provider, each of which may differ in their quality from active duty physicians. While I control for provider-type 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.

Table A2 show the results of four robustness checks designed to test these alternative explanations. In panel A, I add in a clinic-quarter fixed effect to the main specification. This fixed effect is intended to take out any variation in a clinic's capabilities. For instance if the loss of a PCP to a deployment limits the overall capacity of the clinic. The magnitude of the effects are not significantly different from the main results implying that the results are not driven by some time-varying effect on the clinic. In panel B, I display the results from limiting patients to an active duty military provider in the pre and post-period. The results are a bit attenuated with about a 7% increase in costs, though the sample size is much smaller . In panel C, I present the results from eliminating anyone who switches to a purchased care (private sector) primary care provider. The spending coefficient is about 8%, again slightly lower but in line with the main results. Finally, in column four I restrict the analysis to those who have physicians rather than advanced practice providers in the pre and post period. The results here indicate about a 14% increase in spending, slightly higher than the main results. As the choice of a physician versus an advanced practice provider may not be random, these are potentially higher intensity patients. The change in probability of a specialty encounter is remarkably consistent across the samples - about 1.6 percentage point increase, slightly lower than the main results. 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.

Finally, I run a falsification test using short deployments to provide supporting

evidence that discontinuity in care is in fact driving the results. These are operational assignments where the provider returns within three months. I expect that there is little information lost when the provider is only gone for a short time period. table A3 presents the results of running equation three on a sample that includes all deployments less than 90 days and excludes any longer deployments. There are no significant effects across the three main utilization measures. This provides additional support that the deployment itself does not cause an effect, but that the effect is generated through the impact of the deployment on discontinuous care and loss of specific information.

IV. Discussion and Conclusions

In this paper I study the effects of losing specific information on productivity 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 information creates about 12% additional cost per patient and that this is driven by a much higher 29% increase in costs for those most affected. These effects are particularly relevant given that the high investment in health information technology by the Military Health System where any provider can see generalizable information regarding the patient. I also find that a primary care discontinuity has an effect on other members of the care team. Examining specialty encounters, I find that new referrals become less appropriate while care of existing patients becomes more costly. This is consistent with the both the literature on specific information (Jensen and Meckling (1992)) as well as work on coordination costs, specifically Becker and Murphy's (1992) model that showed coordination costs limit the extent of the market.

This study has important implications for health care where many recent poli-

cies 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 resulting in less accumulation of new, specific information. Another example is the patient centered medical home (PCMH) model. This complex model is designed to promote sharing of information about a patient across a group of providers. However, the patient's relationship with an individual provider may be weakened. This could conceivably mitigate the effects of discontinuity by reducing continuity overall. It's 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 information is contained within individuals, firms may adopt policies that provide more opportunities for 'warm handoffs' in which some specific information can be shared prior to the team-member departing. Firms may also adopt technology that reduces the cost of information transfer. 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, and Xiaoxi Zhao.** 2020. "The Impact of Organizational Boundaries on Healthcare Coordination and Utilization." 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.
- Bond, Amelia M, and Stephen D Schwab.** 2019. "Utilization Variation In Military Versus Civilian Care: Evidence From TRICARE." *Health Affairs*, 38(8): 1327–1334.
- 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.
- David, Guy, Evan Rawley, and Daniel Polsky.** 2013. "Integration and task allocation: Evidence from patient care." *Journal of economics & management strategy*, 22(3): 617–639.
- 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.** 2007-2017a. “Contingency Tracking System.” <https://dmdcrs.dmdc.osd.mil/dmdcrs>.
- Defense Manpower Data Center.** 2007-2017b. “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.
- Elwyn, Glyn, Dominick Frosch, Richard Thomson, Natalie Joseph-Williams, Amy Lloyd, Paul Kinnersley, Emma Cording, Dave Tomson, Carole Dodd, Stephen Rollnick, et al.** 2012. “Shared decision making: a model for clinical practice.” *Journal of general internal medicine*, 27(10): 1361–1367.
- 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.
- Frakes, Michael, and Jonathan Gruber.** 2019. “Defensive medicine: Evidence from military immunity.” *American Economic Journal: Economic Policy*, 11(3): 197–231.
- 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."
- Polanyi, Michael.** 1958. "Personal Knowledge: Towards a Post-Critical Philosophy."
- Press, Matthew J.** 2014. "Instant replay—a 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.** 2020. "Association of Primary Care Physician Turnover with Utilization and Outcomes Among Medicare Patients: Observational Study." Working Paper.
- 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.

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.

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.

APPENDIX

Table A1—: Alternative Estimation of Main Results

	(1)	(2)	(3)
	Inverse Hyperbolic Sine of Spending	Number of Specialty Encounters	Number of Primary Care Encounters
Discontinuous Care	0.123*** (0.024)	0.051** (0.021)	0.057** (0.014)
Model:	OLS	Poisson	Poisson
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
N	3,247,298	2,484,388	2,875,144

Note: Results of estimating equation 2 using alternative transformation and distributional assumptions. Column 1 transforms total dollars using the inverse hyperbolic sine rather than a log + 1 method. Column 2 estimates the effects of a discontinuity on specialty care using a poisson model. Column 3 does likewise for primary care visits. All regressions include an individual and cohort-time fixed effect as described in the text. Standard errors are clustered by patient.

Table A2—: Robustness Checks

	(1) Log of Spending	(2) Probability of Specialty Encounter
<i>Panel A - Clinic-Quarter Fixed Effect</i>		
Discontinuous Care	0.103*** (0.023)	0.016*** (0.004)
Sample Mean	\$396.41	0.33
N	3,247,298	3,247,298
<i>Panel B - Active Duty Only</i>		
Discontinuous Care	0.065** (0.028)	0.016*** (0.005)
Sample Mean	\$403.08	0.33
N	826,397	826,397
<i>Panel C - Direct Care Only</i>		
Discontinuous Care	0.080*** (0.024)	0.016*** (0.004)
Sample Mean	\$377.56	0.32
N	2,810,132	2,810,132
<i>Panel D - Physician Only</i>		
Discontinuous Care	0.133*** (0.031)	0.017*** (0.006)
Sample Mean	\$399.60	0.32
N	1,427,645	1,427,645
Controls	No	No
Fixed Effects	Yes	Yes

Note: Results from estimating equation two with adjustments and sample restrictions. Panel A adds a clinic-quarter fixed effect to the equation. Panel B omits patients who ever have a non-active duty military primary care provider. Panel C omits patients who are ever managed in the private sector. Panel D omits patients managed by an advanced practice provider.

Table A3—: Results of Falsification Test

	(1) Log of of Spending	(2) Probability of Specialty Encounter	(3) Probability of Primary Care Encounter
Discontinuous Care	0.020 (0.053)	0.004 (0.010)	0.008 (0.009)
Sample Mean	\$396	0.33	0.44
Controls	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
N	3,190,874	3,190,874	3,190,874

Note: Results from estimating equation two to deployments that last between one and three months and omitting longer deployments. Discontinuous Care coefficient is the interaction of the treatment group identifier and the post period identifier. Controls include whether the Primary Care Provider (PCP) in that quarter is active duty military, and whether the PCP is a physician or advanced practice provider. All regressions include individual and cohort-time fixed effects as described in the text. Standard errors are clustered by patient. The deployment quarter is omitted from these regressions.