Appendix B (for Online Publication)

B-1 Pod Trends

In Section 3.2 of the main paper, I introduce two locations, or "pods," within the ED that I use to estimate the overall effect of the self-managed system. Alpha pod always operated as a self-managed system, while Bravo pod switched from a nurse-managed system to a self-manged one in March 2010. I describe in that section that Alpha generally received the more time-intensive and complex patients, because has always been opened 24 hours, but that Bravo began receiving more complex patients over time, in part due to an increase in the overall volume and complexity of patients arriving at the ED.

In this section, I show evidence of these trends. Understanding these trends is relevant for assessing the robustness of causal inference of the overall effect. If over time Bravo received less time-intensive patients according to unobservable characteristics, then I would not be able to distinguish the effect of the self-managed system from patient selection. On the other hand, if Bravo received more unobservably time-intensive patients over time, then I would instead have a conservative estimate for the causal effect of the self-managed system.

The primary approach I take in this appendix is to examine trends in observable patient and physician peer characteristics, even though I control for observable characteristics in estimating the overall effect in Section 4 of the main paper. If increasingly complex and time-intensive patients are assigned to Bravo pod over time according to observable characteristics, then I will have greater confidence that unobservable characteristics are not going in the opposite direction. This approach complements the approach I take in the main paper of using frequent observations in a long time span to show conditionally parallel trends between the pods and robustness to including pod-specific trends, both of which should include unobservable factors at the pod-time level.

B-1.1 Patient Characteristics

Table B-1.1 presents average patient characteristics for patients in Alpha and in Bravo. It shows that older patients and patients with more severe conditions (a lower emergency severity index indicates a more severe condition) were generally sent to Alpha.

Figures B-1.1 to B-1.3 plot average patient age, Emergency Severity Index (ESI), and number of Elixhauser indices, respectively, for patients seen in Alpha and Bravo over time. The ESI is an integer ranging from 1 to 5, which is the product of a triage algorithm based on patient pain level, mental status, vital signs, and medical condition (Tanabe et al, 2004). An ESI of 1 represents the most severe patient, while an ESI of 5 represents the least severe patient. Elixhauser indices capture 30 important medical conditions, based on coded diagnoses in the medical record, including congestive heart failure, diabetes, hypothyroidism, AIDS, metastatic cancer, and drug abuse (Elixhauser et al, 1998). In regressions in the main paper, I include dummies for each one of these conditions. In Figure B-1.3, I simply plot the number of Elixhauser indices. These three figures are consistent with more complex and severe patients being sent to Alpha at baseline and with increasingly complex and severe patients being sent to Bravo over time.

Figure B-1.4 summarizes the combined effect of observable patient characteristics on patient length of stay by plotting predicted log length of stay. I first estimate the following regression, using only visits to Alpha pod in 2005:

$$Y_{it} = \beta \mathbf{X}_{it} + \varepsilon_{ijkt}, \tag{B-1.1}$$

where $\mathbf{X}_{it}$ are rich patient characteristics that can include the following: age, sex, race, language, ESI, Elixhauser indices, and Major Diagnostic Categories (25 mutually exclusive categories generally based on the organ system and determined by the primary diagnostic code) for patient i at visit arrival time t. I then use estimates $\hat{\beta}$ from Equation (B-1.1) to generate expected lengths of stay for all patients. Results in Figure B-1.4 are predictions based on age, sex, and ESI. Other predictions based on more inclusive sets of patient characteristics, potentially endogenous because they are in part based on physician coding, but they show similar relationships. Patients with longer predicted lengths of stay are always sent to Alpha, but this differential reduces over time.

B-1.2 Physician and Peer Characteristics

In also examine the characteristics of physicians working in Alpha versus Bravo over time. Note that an important feature of the empirical setting is that I observe the same physicians and other providers in both pods over time. This allows me condition on workers in estimating the effect of the self-managed system. Nonetheless, in this section, I explore whether there are any systematic trends in assigning physicians and peers to Alpha versus Bravo over time. This complements Section B-2, which shows that the overall effect is robust to including peer characteristics (in addition to physician-nurse-resident identities), and Section B-6, which shows that physicians are as good as randomly assigned to patient and peer types.

The first set of results are based on physician productivity. Physician productivity is first estimated as fixed effects in a regression of length of stay, controlling for all possible interactions of other team members (physician assistant or resident and nurse), coworker, pod location; patient demographics (age, sex, emergency severity index, Elixhauser comorbidities); ED arrival volume; and time dummies (month-year combination, day of the week, and hour of the day). The average difference in productivity between physicians of above- and below-median productivity is 0.28, meaning that physicians with above-median productivity take 28% less time than those with below-median productivity.

Figure B-1.5 shows average physician productivity for each pod, month, and year combination. Figure B-1.6 shows similar productivity averages for peers, when present. Both figures show an increase in the fixed effect over time for both pods, which is an artifact of the fact that physicians observed during earlier times are observed for a longer time period, which results in lower average lengths of stay. Figure B-1.7 shows averages in the difference in productivity between physicians and peers, when peers are present. Note that a difference in average physician fixed effects of 0.01 between pods implies that physician identities explains a 1% difference in length of stay between the two pods. Given results in B-7.1, a difference in average peer fixed effects of 0.01 would lead to an even lower 0.1% difference in length of stay. There are no economically significant differences in physician productivity trends between pods, compared to the overall effect of the self-managed system on length of stay of -11% to -15%.

The second set of results are based on physician tenure, which is calculated as the difference between the patient date of visit and the physician date of hire. Figure B-1.8 shows average physician tenure for each pod, month, and year combination. Figure B-1.9 shows the corresponding tenure for peers, when present. Figure B-1.10 shows the difference between physician and peer tenure, when a peer is present. There does not appear to be any substantial difference in trends between Alpha and Bravo for any of these tenure measures.

B-1.3 Average Log Length of Stay

I finally plot out average, unadjusted log lengths of stay for both pods and each month in Figure B-1.11. These unadjusted numbers show a gradually increasing trend in length of stay in Bravo and also a decreasing trend in Alpha. These trends occur prior to the Bravo regime change in March 2010. Alpha's trend is continuous across the regime change. Although it is difficult to spot a break in Bravo's trend in unadjusted log length of stay, there does appear to be a decrease in average length of stay starting in February 2010. Recall that the regime change was announced in January 2010.

Results in Section B-1.1 suggest that at least part of this is due to more intensive patients (i.e., patients expected to stay longer) assigned to Bravo over time. While these trends are not unconditionally parallel, I show in Section 4 and Figure 4 in the main paper that they are quite parallel when controlling for patient characteristics and provider identities. Moreover, the differential trends in unadjusted log length of stay goes in the opposite direction of the estimated self-managed effect in the regression, which is that switching to a self-managed system *reduced* Bravo's length of stay.

B-2 Robustness of Overall Effect to Peers

In Section 4 of the main paper, I estimate the overall effect of the self-managed system conditioning on increasingly rich sets of patient characteristics. I find that the effect on length of stay increases in magnitude from -11% to -13% as I include more patient controls and increases further to -15% when I allow for pod-specific time trends. This is consistent with the fact, shown in Section B-1.1, that increasingly more severe and complex patients were sent to Bravo, relative to Alpha, over time.

In contrast, Section B-1.2 shows no qualitatively significant difference in trends in peer characteristics between the two pods. In this section, I formally show that controlling for peer characteristics has no economically significant influence on the estimated effect of the self-managed system. Augmenting Equation (4.1), I estimate

$$\begin{aligned} Y_{ijkpt} &= \alpha Self_{pt} + \sum_{s=0}^{1} 1\left(NoPeer_{jt} = s\right)\left(1 + \left(1 - NoPeer_{jt}\right)\mathbf{PeerChar}_{jt}\right) + \\ &\quad \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + \nu_{jk} + \varepsilon_{ijkpt}, \end{aligned}$$

where $NoPeer_{jt}$ is a dummy that equals 1 if there is no peer present, and $\mathbf{PeerChar}_{jt}$ is a vector of characteristics for the peer of physician j, including the cumulative number of days that physician j and his peer have worked together, the tenure of the peer, the difference between the physician's and peer's tenures, the peer's productivity fixed effect, and the difference between the physician's and peer's productivity fixed effects. The coefficient of interest, α, is essentially unchanged at -0.131 with a robust standard error of 0.040, yielding a p-value of 0.002.

B-3 Distribution of Physician Productivity

In this appendix, I briefly discuss three methods of estimating the distribution of physician effects on length of stay, taking into consideration the fact that physician effects are estimated with measurement error. Consistent with regressions in the main text, consider the following regression of log length of stay:

$$Y_{ijkpt} = \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + c_j + \nu_k + \varepsilon_{ijkpt}, \tag{B-3.1}$$

for patient i, physician j, nurse-resident k, pod p, and t. As before, I control for patient characteristics $\mathbf{X}_{it}$, a vector of time characteristics (hour of the day, day of the week, and month-year interactions) $\mathbf{T}_t$, and pod fixed effects ζ_p. While I previously partialed out physician-nurse-resident trio identities with a term ν_{jk}, I now am separately interested in the physician effects with the term c_j (and now partial out nurse-resident effects with ν_k).

In the standard fixed effect estimation, c_j is estimated with error:

$$\hat{c}_j = c_j + \xi_j,$$

where $\hat{c}_j$ is a measured physician effect, c_j is the true physician effect, and ξ_j is an error term. The fixed effect estimator is an unbiased estimate of c_j, but because of ξ_j, the standard deviation of the distribution of $\hat{c}_j$ will overestimate the standard deviation of the true distribution of c_j, which is the object of interest. All three methods assume that $c_j \sim N\left(0, \sigma_c^2\right)$ and attempt to recover σ_c^2.

B-3.1 Empirical Bayes Estimator

The preferred method uses an empirical Bayes (EB) procedure, based on Morris (1983). This procedures assumes that $\xi_j \sim N\left(0, \pi_j^2\right)$, or equivalently $\hat{c}_j \left| c_j, \sigma_\xi^2 \right. \sim N\left(c_j, \pi_j^2\right)$. For a prior distribution $c_j \sim N\left(\bar{c}, \sigma_c^2\right)$, the posterior distribution that conditions on estimates $\hat{c}_j$ from fixed effect estimation and σ_ξ^2 is

$$c_j \left| \hat{c}_j, \sigma_c^2, \sigma_\xi^2 \right. \sim N\left(c_j^{EB}, \sigma_\xi^2 \left(1 - B_j\right)\right),$$

where $c_j^{EB} = \left(1 - B_j\right) \hat{c}_j + B_j \bar{c}$ and $B_j = \pi_j^2 / \left(\pi_j^2 + \sigma_c^2\right)$. B_j "shrinks" $\hat{c}_j$ towards the prior mean of $\bar{c}$, the extent of which depending on the degree of measurement error π_j^2.

As described in Morris (1983), I implement the following feasible version of the procedure:

1. Estimate $\hat{c}_j$ in Equation (B-3.1) by fixed effects estimation. The estimated $\hat{c}_j$ will also have a standard error, which is squared to yield a value called $\hat{\pi}_j^2$.

2. Construct a set of weights for each physician, denoting the weight for physician j as W_j. Begin the procedure with $W_j = 1$ for all j.

3. Iterate the following to convergence:

   (a) Estimate $\bar{c} = \sum_j \left(W_j \hat{c}_j\right) / \sum_j W_j$ and

   $$\hat{\sigma}_c^2 = \frac{\max\left\{0, \sum_j W_j \left[\left(\hat{c}_j - \bar{c}\right) - \hat{\pi}_j^2\right]\right\}}{\sum_j W_j}.$$

   (b) Recalculate $W_j = 1/\left(\hat{\pi}_j^2 + \hat{\sigma}_c^2\right)$.

I take the final value of $\hat{\sigma}_c^2$ as the EB-adjusted measure of the variance of the distribution of c_j. The standard deviation of this distribution from EB estimation is $\hat{\sigma}_c^{EB} = 0.091$. This implies that increasing true physician productivity by one standard deviation would result in a 9.1% improvement in length of stay.

B-3.2 Random Effects Estimator

The second method is by random effects estimation, in which I assume that c_j is drawn from a random distribution with variance σ_c^2, and I directly estimate $\hat{\sigma}_c^2$, by maximum likelihood estimation. In order to avoid the incidental parameters problem, I first reduce the dimensionality of the covariates by estimating a OLS-predicted length length of stay based on the covariates. The procedure is thus as follows:

1. Estimate $\hat{Y}_{ikpt}$, predicted log length of stay without conditioning on physician j, by the equation $Y_{ijkpt} = \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + \nu_k + \tilde{\varepsilon}_{ijkpt}$.

2. Estimate by maximum likelihood

$$Y_{ijkpt} = \beta \hat{Y}_{ikpt} + c_j + \varepsilon_{ijkpt}, \tag{B-3.2}$$

   assuming that c_j is randomly distributed as $c_j \sim N\left(0, \sigma_c^2\right)$.

By this procedure, the estimated standard deviation of c_j is $\hat{\sigma}_c^{RE} = 0.065$.

The key random effects assumption is that

$$E\left[c_j \left| \mathbf{X}_{it}, \mathbf{T}_t, \zeta_p, \nu_k \right.\right] = 0, \tag{B-3.3}$$

which means that c_j for physician j is assumed to be orthogonal to patient types, time periods, pods, and nurse-resident teams that the physician is observed to be associated with (e.g., Wooldridge, 2010). In this case, given the above reduction of dimensionality and the assumption that Y_{ijkpt} takes the form of Equation (B-3.2), the assumption in Equation (B-3.3) takes the following weaker form:

$$E\left[c_j \left| \hat{Y}_{ikpt} \right.\right] = 0. \tag{B-3.4}$$

In Section B-6, I show that physicians are as good as randomly assigned to patient types, peer types, and ED conditions such as patient volume to the ED, *conditional* on time periods and pod locations. It is not generally true that physicians are as good as randomly assigned to work at different times and in different pods. Rather, physicians are allowed to state preferences, which are followed to some degree, even though I observe physicians working in both pods and in all days of the week and hours of the day. Thus, any correlation between physician identities, pods, and time categories would cause the random effects estimation to be biased.

B-3.3 Correlated Random Effects Estimator

The third method allows for correlated random effects (e.g., Mundlak, 1978; Altonji and Matzkin, 2005; Wooldridge, 2010), which parametrically model the correlation between c_j and covariates. Specifically, I assume $c_j = u_j + v_j = \lambda \hat{E}\left[\hat{Y}_{ikpt} \middle| j\right] + u_j$, where

$$E\left[v_j \left| \hat{Y}_{ikpt} \right.\right] = 0. \tag{B-3.5}$$

This model considers a "fixed" component to the physician effect, u_j, that is predicted by exposure to covariates, namely average predicted log length of stay, and a random component, v_j, that is assumed orthogonal to $\hat{Y}_{ikpt}$.

I implement this estimation by the following:

1. Estimate $\hat{Y}_{ikpt}$, predicted log length of stay without conditioning on physician j, by the equation $Y_{ijkpt} = \beta\mathbf{X}_{it} + \eta\mathbf{T}_t + \zeta_p + \nu_k + \tilde{\varepsilon}_{ijkpt}$.

2. Calculate empirical expectations $\hat{E}\left[\hat{Y}_{ikpt}\middle| j\right] = \sum_{j(i,t)=j}\left(s_{it}\hat{Y}_{ikpt}\right)/\sum_{j(i,t)=j} s_{it}$, where s_{ikpt} is an indicator function for whether $\hat{Y}_{ikpt}$ exists (i.e., all covariates used in step 1 are non-missing) and $j(i,t)$ is a assignment function indicating the physician associated with patient i arriving at time t.

3. Estimate by maximum likelihood

$$Y_{ijkpt} = \beta\hat{Y}_{ikpt} + \lambda\hat{E}\left[\hat{Y}_{ikpt}\middle| j\right] + v_j + \varepsilon_{ijkpt}, \tag{B-3.6}$$

   assuming that v_j is randomly distributed as $v_j \sim N\left(0, \sigma_c^2\right)$.

I estimate the standard deviation of u_j as $\hat{\sigma}_u = 0.056$. The empirical standard deviation of $\hat{v}_j \equiv \hat{\lambda}\hat{E}\left[\hat{Y}_{ikpt}\middle| j\right]$, weighted by $\sum_{j(i,t)=j} s_{it}$, is $\hat{\sigma}_v = 0.036$. Given that u_j and v_j are orthogonal by assumption, the standard deviation of the physician effects $c_j = u_j + v_j$ is $\hat{\sigma}_c^{CRE} = \sqrt{\hat{\sigma}_u^2 + \hat{\sigma}_v^2} = 0.069$.

B-3.4 Summary of Estimators

In summary, I have estimated three different measures of variation in physician-level effects c_j in Equation (B-3.1). The EB estimator suggests a standard deviation of $\hat{\sigma}_c^{EB} = 0.091$ for the distribution of physician effects. Random effects and correlated random effects estimators suggest $\hat{\sigma}_c^{RE} = 0.065$ and $\hat{\sigma}_c^{CRE} = 0.069$, respectively. The random effects estimator suffers from the problem that c_j may be correlated with other covariates that predict length of stay, while the correlated random effects estimator partially addresses this by allowing for parametric correlation between average covariates and a fixed portion u_j of c_j while allowing for a random v_j that is uncorrelated with covariates. Considering the EB estimator as my preferred (and conservative) measure of variation in physician productivity, an overall effect of 11-15% for the self-managed system is equivalent to 1.2-1.7 standard deviations of physician productivity.

B-4 Alternative Methods of Inference

In estimating the overall effect of the self-managed system in Section 4 of the main paper, my baseline specification clusters standard errors by physician, which is equivalent to an experiment sampling at the level of physicians, who are given shifts that translate to pods and organizational systems, before and after the regime change in Bravo. This thought experiment is supported by evidence of conditional quasi-random assignment of physicians to patients and peers, shown in Section B-6.

In Sections B-4.1 and B-4.2, I discuss two alternative methods of inference that allow for pod-level random shocks, given that I have only two pods and cannot cluster by pod for inference. The intuition underlying both of these alternative methods takes advantage of the fact that I observe a long time dimension for both pods. I can therefore compare "random" variation in the difference between the pods over many points in time, which are not associated with any regime change, with the effect associated with Bravo pod's switch from a nurse-managed to self-managed system in March 2010.

B-4.1 Inference with Serially Correlated Pod-level Error Terms

In this approach, I address sampling variation at the pod level across time with a parametric form on the error terms. Specifically, I allow for pod-month random shocks that are serially correlated across months by a first-order autoregressive (AR1) process. That is, I first calculate month-year-pod fixed effects from

$$Y_{ijkpt} = \sum_{m=1}^{M} \sum_{y=1}^{Y} \alpha_{myp} I_{t \in m} I_{t \in y} + \beta \mathbf{X}_{it} + \tilde{\eta} \tilde{\mathbf{T}}_t + \nu_{jk} + \varepsilon_{ijkpt}, \tag{B-4.1}$$

which is the same model used in the paper to generate Figure 4. $I_{t \in m}$ and $I_{t \in y}$ are indicator functions for t belonging in month m and year y, , respectively, and $\tilde{\mathbf{T}}_t$ is a revised vector of time categories only for day of the week and hour of the day. Coefficients estimated for α_{myp} are used as data points.

In the second stage, I estimate a model with observations collapsed to the month-year-pod level:

$$\hat{\alpha}_{myp} = \gamma Self_{myp} + \eta_{my} + \zeta_p + \varepsilon_{myp},$$

where $\hat{\alpha}_{myp}$ are estimated coefficients from Equation (B-4.1), γ is the coefficient of interest on the treatment indicator $Self_{myp}$ of whether pod p is a self-managed system during month m and year y, η_{my} and ζ_p are respective time and pod fixed effects, and ε_{myp} is a serially correlated error term with the AR1 process

$$\varepsilon_{myp} = \rho \varepsilon_{my-1,p} + z_{myp}.$$

I estimate a self-managed effect of $\hat{\gamma} = -0.0981$, with a standard error of 0.0281, which is significant with a p-value of 0.001. This estimate is quite similar to the baseline estimate in the main paper. The estimated correlation in error terms across months is $\hat{\rho} = 0.302$.

B-4.2 Inference with Systematic Placebo Tests (Randomization Inference)

Another alternative method of inference takes sampling as fixed and instead considers randomization at the level of treatment (Rosenbaum, 2002). This is also in the same spirit of systematic placebo tests (e.g., Abadie et al, 2010). Under the sharp null of no effect of the self-managed system, there should be no significant difference between my obtained estimates and those I would obtain if I consider a number of placebo regime changes over each pod and month.

Here, I again use estimated month-year-pod fixed effects $\hat{\alpha}_{myp}$ from Equation (B-4.1). I then perform a regression at placebo regime changes at each month and pod with a bandwidth of three months prior to the placebo regime change date and three months post. That is, using only only observations from within the bandwidth for the two pods, I estimate a difference-in-differences regression

$$\hat{\alpha}_{myp} = \gamma^r \mathit{PlaceboSelf}^r_{myp} + \eta_{my} + \zeta_p + \varepsilon_{myp},$$

where $\mathit{PlaceboSelf}^r_{myp}$ is an indicator for whether pod p at month m and year y is self-managed under the placebo regime change r.

Out of the 130 estimates for γ^r, the true regime change in March 2010 had the second largest coefficient, corresponding to a randomization-inference p-value of about 0.015. Notably, the largest coefficient corresponds to February 2010, the month before the official regime change. The regime change was announced in January 2010; as shown in Section B-8, there appears to be some anticipatory response by physicians in Bravo beginning to see more patients in beds outside of their own (e.g., physicians in Bravo-1 shifts seeing patients in beds belonging to Bravo-2). This ranking of coefficients is robust to using different bandwidths.

B-5 Overall and Pod Patient Volume

Patient volume, or the flow of patients, is a key feature of work in the ED. Although differences in patient volume are expected across different hours of day at the time of scheduling, there remains substantial additional variation in realized patient volume within time categories. Furthermore, while working in a pod, although physicians are aware via the computer interface of the current state of patient volume, there remains substantial uncertainty about the distribution of these patients to pods. I use both of these facts in order to identify the mechanism of foot-dragging as the response to expected future work.

B-5.1 Time Category Variation, Across and Within

I first calculate overall patient volume as the number of patients arriving to the waiting room at each hour in the data sample, including hours with no patients arriving. I then calculate summary statistics for patient volume – extrema and quantiles – across observations within each value of a time category. Figures B-5.1 to B-5.3 show respective variation relevant to hours of the day, days of the week, and month-year interactions. These figures are meant to represent qualitative variation across and within time categories. In essentially all regressions in the paper, I include fixed effects for each value of each time category. In Section B-6.1, I also show evidence of quasi-random exposure to patient volume, conditional on the time categories.

As shown in Figure B-5.1, perhaps the greatest variation across time categories is attributable to hours of the day. As expected, reflected in staffing schedules, more patients arrive during the day than at night. However, there remains significant variation within each hour of the day, with the 95th percentile always more than double the volume of the 5th percentile even at peak hours. Figure B-5.2 shows variation across and within days of the week. Monday is the busiest day, but the variation across days of the week is negligible compared to the variation that remains after accounting for day of the week. Similarly, as shown in Figure B-5.3, variation across month-year observations is negligible compared to within month-year variation. There does appear to be a trend towards greater patient volumes, especially in above-median quantiles, over time.

B-5.2 Low Correlation between Overall and Pod Volume

Figure B-5.4 presents plots that show the relationship between the number of patients arriving at triage in the hour prior to the index patient's pod arrival, which I use in the main manuscript as a measure of expected future work. In order to show that I can separate expected future work from actual (current or future) work, I show that there is wide variation in pod-specific patient volume, even conditional on overall ED volume arriving at triage. Each of the four panels in Figure B-5.4 examines a different hour for pod-specific volume. The median time from triage to pod is about 30 minutes, and the median time from pod to discharge order is about 3.4 hours.

Correlation coefficients are low: The greatest correlation coefficient coefficient is 0.21, between overall volume and pod-specific volume in the same hour previous to the index patient's arrival to the pod. Correlation coefficients between previous-hour overall volume and pod-specific volume for the subsequent hour, the hour ending two hours later, and the hour ending three hours later are 0.16, 0.13, and 0.12, respectively.

B-6 Conditional Quasi-random Exposure to Patients and Peers

Although the effect of the self-managed system is identified by the Bravo regime change, providers working in both pods over time, and conditionally parallel trends in the two pods, it may also be useful to show evidence of conditional quasi-random exposure of physicians to patients and peers. First, this addresses more complicated threats to identification that involve both physician selection and time-varying productivity. Second, it is consistent with the sampling thought experiment in which physicians are assigned random pod experiences conditional on rough time categories. Third, quasi-random exposure to patient volume in particular supports the idea of exogenous shocks to expected future work, which I use to identify foot-dragging.

B-6.1 Similar Exposures across Physician Types

I first investigate whether physicians of different types are exposed to similar patient types and ED volume conditions. I focus on physician differences in terms of preferences and productivity.

I estimate preferences for specific patient types by the probability that a physician will choose a patient type when given the choice. I form productivity measures by fixed effects for physician identities in a regression of log length of stay. Physicians with one standard greater preferences for a patient type are 7.4% more likely to choose that patient type than average. Physicians who are one standard deviation faster than average have 11% shorter lengths of stay.

In Table 1 in the main paper and Table B-6.2, I show that physicians that differ by productivity or preferences, respectively, are exposed to similar average patient types and ED patient volume. In addition to showing that averages are similar, I examine the distribution of patient volume more closely. Figure B-6.1 shows that the distribution of patient volume for high- and low-productivity physicians are indistinguishable.

B-6.2 Joint Insignificance of Physician Identities

I test for the joint significance of physician identities in regressions of patient characteristics arriving at the pod and ED patient volume, while conditioning on rough indicators of time. For patient characteristics, I first summarize patient characteristics of age, sex, emergency severity index (ESI), race, and language into predicted length of stay. I calculate this prediction for each patient and then average these predictions for patients arriving at each pod and hour-date in my data. Using physician schedules, I associate the average predicted length of stay for each pod-hour combination to physicians that are working on that pod and during that hour-date. I then estimate this equation:

$$\overline{E\left[Y\right]}_{pt} = \alpha_j I_{jpt} + \eta \mathbf{D}_{pt} + \varepsilon_{jpt},$$

where $\overline{E\left[Y\right]}_{pt}$ is the average predicted log length of stay for pod p at time (hour-date combination) t, I_{jpt} is an indicator variable equal to 1 if physician j is working in pod p at time t, and $\mathbf{D}_{pt}$ is a vector of interactions between pod p and rough time dummies including month-year dummies, day of the week dummies, and hour of the day dummies. I test for the joint significance of the vector of coefficients $\boldsymbol{\alpha} = (\alpha_j)$. The F-statistic (64, 659042) under the null that $\boldsymbol{\alpha} = \mathbf{0}$ is 1.07 (p-value of 0.32), clustering by date.

For ED volume, I perform a similar exercise. I estimate this equation:

$$EDWork_t = \alpha_j I_{jt} + \eta \mathbf{T}_t + \varepsilon_{jt},$$

where $EDWork_t$ is the number of patients arriving at the ED during time (hour-date combination) t, I_{jt} is an indicator variable equal to 1 if physician j is working in the ED at time t, and $\mathbf{T}_t$ is a vector of indicators for rough time categories of month-year, day of the week, and hour of the day. The F-statistic (64, 2098) under the null that $\boldsymbol{\alpha} = \mathbf{0}$ is 1.11 (p-value of 0.26) when clustering by unique date (each of the 75 physicians in the sample works from 60 to 845 days).

B-6.3 Exogenous Assignment of Physicians to Peers

I also show that physician identities do not explain the preferences or ability of their peers that they happen to be working with on a shift. With respect to physician productivity, I regress the productivity (length of stay) fixed effect of the physician against that of his peer. That is, if $\hat{\alpha}_j$ is the productivity fixed effect for physician j, and $-j\,(t)$ denotes physician j's peer during shift t, I perform this regression:

$$\hat{\alpha}_{-j(t)} = \beta\hat{\alpha}_j + \varepsilon_{jt}.$$

I estimate the coefficient β to be small and insignificant at -0.003 with a standard error of 0.014 (p-value of 0.84). Similarly, the correlation coefficient between physician and peer fixed effects is -0.018 (p-value 0.16).

If I normalize the average productivity fixed effect of peers working with physicians who are faster than average to be 0 (with standard deviation 0.107), the average productivity fixed effect of peers working with physicians who are slower than average is 0.0001 with standard deviation of 0.105.

If instead of entering the physician's productivity fixed effect on the right-hand-side, I enter physician dummies in the specification

$$\hat{\alpha}_{-j(t)} = \eta_j + \varepsilon_{jt},$$

where the error term is clustered by unique pairings of physician and peer, I am unable to reject the null that the vector $\boldsymbol{\eta}$ is jointly 0 with an F-statistic (53, 1995) of 0.57 (p-value 0.99).

I perform similar analyses with respect to physician preferences and find no relationship between the preferences of peers.

B-7 Peer Effects

This section extends the analysis of peer effects in two ways. First, I estimate direct effects in Section B-7.1, which is effect of working with a peer of higher productivity. Second, I consider how foot-dragging might depend on the relationship between the peer and the index physician in Section B-7.2.

B-7.1 Direct Peer Effects

While the main paper focuses on peer effects on foot-dragging, as an interaction between the presence of a peer and expected future work, I consider the direct effect of peer productivity in this appendix. Consistent with Mas and Moretti (2009), I find that working with a faster peer shortens lengths of stay for the index physician, but I do not find a significant difference in this peer effect between organizational systems.

The effect of peer productivity may occur through a variety of mechanisms. Social incentives have been discussed as a source of positive peer effects (e.g., Mas and Moretti, 2009). Knowledge spillovers are yet another mechanism for peer effects. In addition, in this setting, strategic behavior may also lead to positive peer effects. For example, in the nurse-managed system, a physician working with a less-productive peer will be more likely to get new work unless if he slows down. Moreover, other mechanisms discussed in the main paper could influence the sign and magnitude of peer effects. Free riding by waiting for productive peers to choose work would have a negative influence on peer effects, while dynamic smoothing and matching may have a positive influence if productive peers have complementary skills and availability or a negative influence if less-productive peers have the complementary skills and availability.

Furthermore, social incentives may differ between the self-managed and nurse-managed systems because the rules of the games differ. In the nurse-managed system, peers impose a negative externality by being less productive (i.e., the foot dragging externality), and social incentives may increase efficiency. In contrast, in the self-managed system, peers actually may impose a positive externality by being less productive to others within the team, because they prevent work from being sent to the pod. In addition to the direction of the social incentives, their strength may differ between the two organizational systems as peers work more closely in self-managed teams.

Although peer effects may be difficult to interpret in terms of mechanisms for these reasons, it is still interesting to compare between the two organizational systems in reduced form. It is certainly possible that some physician-peer combinations may perform better in a self-managed setting while others may perform better in a nurse-managed setting. Employing similar methodology as Mas and Moretti (2009), I first estimate physician fixed effects for log length of stay, and then I use the fixed effects of peers as an explanatory variable in a regression of productivity in order to estimate peer effects.

In the first stage, I estimate the following regression on log length of stay Y_{ijkpt} for patient i, physician j, resident-nurse team k, and visit arrival time t:

$$Y_{ijkpt} = \theta_j + \mathbf{M}\phi_{C_j} + \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + \varepsilon_{ijkpt}, \tag{B-7.1}$$

where in addition to patient characteristics $\mathbf{X}_{it}$ and time categories $\mathbf{T}_t$, I control for all possible sets of physicians j, resident-nurse combinations k, physician peers l, and pod locations p in order to address the reflection problem (Manski, 1993). These providers and peers are included in a set of all dummies representing possible combinations $\phi_{C_j} = \{C(j,k,l,p)\}$, where

$$C(j,k,l,p) = \begin{cases} 1 & \text{if physician } j \text{is working with team } k \text{and peer } l \text{in pod } p, \\ 0 & \text{if } i = l, \\ 0 & \text{otherwise.} \end{cases}$$

The parameter of interest is the physician fixed effect θ_j. The standard deviation of estimated fixed effects is 0.11, meaning that physicians one standard deviation above mean productivity have lengths of stay that are 11% shorter than average.

In the second stage, I use the set of fixed effects from (B-7.1) in order to estimate the effect of a peer's productivity on the index physician's outcomes, in particular length of stay. Using the fixed effect θ_{-j} for physician j's peer, I estimate this regression using separate samples for the nurse-managed and self-managed systems:

$$Y_{ijkpt} = \alpha\theta_{-j} + \beta\mathbf{X}_{it} + \eta\mathbf{T}_t + \zeta_p + \nu_{jk} + \varepsilon_{ijkpt}. \tag{B-7.2}$$

The coefficient α represents peer effects from working with a peer with productivity θ_{-j}. A positive α suggests that physicians work faster when working with a more-productive (faster) peer and slower when working with a less-productive (slower) peer. I also estimate a pooled version with

$$Y_{ijkpt} = \alpha_1\theta_{-j} \cdot \mathit{Self}_{it} + \alpha_2\theta_{-j} + \alpha_3\mathit{Self}_{it} + \beta\mathbf{X}_{it} + \eta\mathbf{T}_t + \zeta_p + \nu_{jk} + \varepsilon_{ijkpt}.$$

A positive α_2 suggests positive peer effects in the nurse-managed system; a positive α_1 suggests greater peer effects in the self-managed system (in the case of positive peer effects). All regressions require that there be a peer present in the pod.

Table B-7.1 reports estimates of peer effects. A 1% increase in peer productivity leads to a 0.1% increase in physician productivity. Faster peers have a stronger influence than slower peers in both settings. However, peer effect estimates can be relatively imprecise, compared to foot-dragging results in the main paper and in particular in the self-managed system. I cannot reject that overall peer effects are different between the two organizational systems.

While effects are quantitatively similar, they are less precisely estimated in the self-managed system than in the nurse-managed system, despite similar numbers of observations. This suggests that peer effects and perhaps the interaction between physicians in self-managed teams, through a number of possible mechanisms, are generally less predictable. Although I do not show results here, I do not find any significant effect of peer effects of working with a peer with productivity specific to the index patient. That is, working with a physician who is better at seeing heart patients does not improve the productivity of the index physician seeing the heart patient. In addition, I do not find peer effects on quality outcomes, which suggests that the main (and modest) effect of peers is mostly on length of stay. These facts discount the possibility of learning or being helped by more skilled physicians.

B-7.2 Peer Effects on Foot-dragging by Peer Type

Given that foot-dragging is reduced by the presence of a peer, I also consider different types of social relationships between physicians and their peers. I first consider peers of the same sex, similar age, or same place of residency training as potentially more connected to each other.

Second, I consider peers who are faster (or more productive) than median. The effect of this peer type on foot-dragging may include both social and strategic concerns. To see the strategic concerns, note that slower peers will cause more work to be redirected to physicians unless they slow down as well. Third, I consider peers, by their history of time working with each other, who are more familiar with each other's workplace behavior and more likely to have established reputations with each other.[1] Finally, I consider peers who have at least two years greater tenure than the index physician. Social hierarchy is a common feature in many workplaces, particularly those with professionals, long tenures, or strong work cultures.

For each of these peer types, I estimate regressions of the following form:

$$\begin{aligned} Y_{ijkpt} &= \alpha_1 EDWork_t + \alpha_2 PeerType_{jt}^m \cdot EDWork_t + \alpha_3 PeerType_{jt}^m + \\ &\quad \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + \nu_{jk} + \varepsilon_{ijkpt}, \end{aligned} \tag{B-7.3}$$

separately for nurse-managed and self-managed samples. $PeerType_{jt}^m$ is an indicator that the peer for physician j at time t is of type m. I am interested in the coefficient α_2 as the effect of working with a peer type on foot-dragging, again identified by increases in length of stay with respect to expected future work.

Table B-7.2 reports results for three of the peer types.[2] Senior peers are the only peer type showing a significant effect on foot-dragging. In the nurse-managed system, working with a senior peer decreases foot-dragging by half, from an increase of 0.8% for each patient arriving at the ED to an increase of 0.4%. In a pooled regression shown in the third column, it also appears that senior peers further reduce foot-dragging in the self-managed system.[3] Other peer types – highly productive peers, familiar peers, and connected peers – show no significantly differential effect on foot-dragging. These results suggest that the most important social relationships between peers may be unilateral ones based on hierarchy, as opposed to ones that are based on connectedness or familiarity.

B-8 Bed Location of Assigned Patients

The assignment of patients to beds and to physicians is a crucial aspect of the management of work in the ED. As illustrated in Figure 1 and described in the paper, the fundamental difference between the nurse-managed system and the self-managed system arises from the fact that, in

[1] I use a threshold of at least 60 hours working in the same pod, which is at the 75th percentile, to describe peers that are "familiar" with each other. Given physician turnover and a large number of shift times and locations, it is is relatively uncommon for two ED physicians to have longer histories working together in the same pod.

[2] For brevity, I omit peer types related to social connectedness from Table B-7.2, as they show no effect on foot-dragging.

[3] The pooled regression includes interactions with the self-managed system and a direct effect for the self-managed system. I do not write out this equation above, as Equation (B-7.3) communicates the effect of interest in α_2. Recall that self-management and social incentives may independently reduce foot-dragging, and that there may be some foot-dragging in the nurse-managed system, since the 0 benchmark for no foot-dragging is conservative.

the nurse-managed system, physicians "own" beds and are therefore assigned patients who are assigned to their beds, while physicians share beds in the self-managed system. In this section, I describe in greater detail the use of beds, with particular emphasis on Bravo pod, which switched from a nurse-managed system to a self-managed system in March 2010.

B-8.1 Pod Layout and Bed Use

Figures B-8.1 and B-8.2 show a computer interfaces for Alpha and Bravo pods, respectively. Alpha and Bravo had largely similar physical layouts and stable bed locations over time. However, unlike Alpha, Bravo was divided into two administrative zones when it was a nurse-managed system. Two physicians working in Bravo would each be assigned to a zone, Bravo-1 or Bravo-2, which included a set of beds. The zones were physically contiguous and non-overlapping.

While the physical layout of both pods was stable, there may be slight changes in the electronic designations of beds, listed in Table B-7.1. In addition to physical beds, electronically designated beds include virtual beds, such as hallways and other unwalled areas where stretchers can be rolled into.[4] Tables B-7.2 to B-7.4 show the number of patient visits, as well as the dates spanning the first and last observed visit, associated with each bed in Alpha, Bravo-1, and Bravo-2, respectively. Patients are preferentially assigned to physical beds within walled rooms, since these beds have greater privacy, compared to hallway beds. Finally, although the numbers of possible beds in each Bravo zone were roughly equal, Bravo-1 had fewer beds than Bravo-2 that were routinely filled, in part because Bravo-1 had more virtual beds.[5]

B-8.2 Patient Assignment to Beds and to Physicians

The triage nurse assigns all patients from the waiting room to their initial beds within the ED. Once patients are assigned a bed, they appear on the pod-specific computer interface for either Alpha (Figure B-8.1) or Bravo (Figure B-8.2).[6] Upon being assigned a bed, patients are either technically assigned to physicians, if physicians own beds as in the nurse-managed system, or otherwise must wait for a physician to sign up for their care, when physicians share beds in the self-managed system. In either case, physicians must click on each patient's representation on the computer interface in order to acknowledge that they are assuming the patient's care.

Physicians in the nurse-managed system may not always attend to patients assigned in their zones. Because all patients must be voluntarily acknowledged, physicians working together in

[4] Hallway beds are designated with an "H" and sometimes a number signifying the closest numbered room. Beds 28 to 31 in Bravo-1 are located in unwalled spaces separated by curtains.

[5] Relatedly, Bravo-1 was located closer to the Bravo doors and was reserved to hold patients who would likely exit the ED sooner. In practice, Bravo-1 had patients with shorter unadjusted lengths of stay. Patients were not necessarily less severe in Bravo-1; there was a wider distribution of Emergency Severity Index scores in Bravo-1 versus Bravo-2.

[6] Patients also appear on the computer interface while they are still in the waiting room. Physicians can tab through each location in the ED, including pods and the waiting room, to observe the patients in each location, their brief clinical information, and the health care providers (if any) assigned to them.

the nurse-managed system may still agree to see patients who are not strictly in their zone. This is particularly the case for patients waiting to be seen in virtual beds in the hallway when the pod is busy and for patients near the border between zones, even though all beds (virtual and physical) belong to only one zone. In this sense, the nurse-managed system is never purely managed by the triage nurse, and a small proportion of patients may be assigned by physician discretion.

In addition, patients routinely may change beds during their ED stay. Although the median visit involves only one bed, 44% of patients switch beds at least once. The mean number of beds occupied during a patient visit is 1.85 beds. The 95th percentile for the number of beds occupied is four. Although the triage nurse has sole discretion for the assignment of the initial bed, physicians and nurses may have some input regarding changing beds after initial pod arrival. Physicians are rarely reassigned when patients change beds; in fact, physicians are expected to care for patients during their entire stay and routinely remain present for two to three hours past their end of shift, unless if a patient is expected to stay significantly longer (e.g., greater than three hours) past the end of the responsible physician's shift.

B-9 Additional Results

In this appendix, I present the following additional empirical results, as well as a brief discussion of each set of results:

- Figure B-9.1: Presents the foot-dragging coefficient on expected future work interacted with pod identities and four-month intervals.

- Figure B-9.2: Presents the correlation between censuses and new patient assignment in both pods over time using a local linear regression (Figure 5 in the main text), with additional confidence intervals on estimates.

- Figure B-9.3: Presents the correlation between censuses and new patient assignment in both pods over time using a local linear regression, but extends the use in Equation (7.1) of vestigial "Bravo-1" and "Bravo-2" shift labels until June 2010. Confidence intervals are also shown.

- Figure B-9.4: Presents the correlation between censuses and new patient assignment in both pods over time, but estimates Equation (7.1) with a kernel regression. Confidence intervals are also shown.

- Table B-9.1: Presents responses to expected future work, estimated by Equation (5.2), for outcomes other than length of stay.

Figure B-9.1 estimates foot-dragging coefficients, which indicate the response of length of stay to expected future work, interacted with pod identities and four-month intervals. Specifically, I estimate

$$Y_{ijkpt} = \sum_{\tau \in T} \alpha_{p\tau} I_{t\in\tau} EDWork_t + \beta \mathbf{X}_{it} + \eta \mathbf{T}_t + \zeta_p + \nu_{jk} + \varepsilon_{ijkpt},$$

where $\tau \in T$ is a four-month interval. Coefficients of interest, $\{\alpha_{p\tau}\}_{p\in\{0,1\},\tau\in T}$, are specific to pod and four-month interval, but are of course foot-dragging coefficients on expected future work, $EDWork_t$, as measured by the number of patients arriving at the ED in the hour prior to the index patient's bed arrival. I therefore can control for the same rich set of time categories (hour of the day, day of the week, and month-year-interaction) and pod identities. Foot-dragging (i.e., increases in length of stay as expected future work increases) does not immediately disappear in Bravo when Bravo changes from a nurse-managed system to a self-managed system. Foot-dragging is indistinguishable from 0 during all periods in Alpha.

Figures B-9.2 presents the same results as in Figure 5 in the main text, plotting the coefficient on censuses a linear probability model in Equation (7.1) of new patient assignment. The only difference between these two Appendix Figures and Figure 5 is that confidence intervals are plotted for Bravo and Alpha pods, respectively. For most months, I can reject that the central estimate for one pod is the same as the central estimate for the other pod, but the two confidence intervals otherwise overlap. Also, all central estimates for Alpha pod are below those for Bravo pod in the pre-regime period.

Figure B-9.3 plots coefficients and confidence intervals for Bravo pod estimated from the following regression:

$$Y_{ijt} = \alpha Census_{jt} + \beta \mathbf{ShiftTime}_{jt} + \gamma ZoneLabel_{jt} + \eta_j + \nu_{it} + \varepsilon_{ijt}, \qquad \text{(B-9.1)}$$

which differs from Equation (7.1) only in that it allows for "Bravo-1" and "Bravo-2" shift labels to matter even after the regime change. The reason for this change relative to Equation (7.1) is to consider the possibility that physicians with higher censuses are more likely to be assigned new patients because they are following the norm that they are responsible for beds in Bravo-2, which is the larger zone. Figure B-9.3 is largely the same as Figure 5 in the main text, except that the positive coefficients immediately at the regime change are now reduced. However, there is still a positive coefficient at April 2010 that is close to being significant at the 5% level.

Figures B-9.4 estimates Equation (7.1) by using a kernel regression. I use triangular kernels with 45-days on each side, which may be truncated if sufficiently close to March 1, 2010. Results are largely similar to the baseline results, shown in Figure 5 in the main text, which are estimated using a local linear regression. I can reject that the central estimate for one pod is the same as the central estimate for the other pod, but the two confidence intervals otherwise overlap. Also, all central estimates for Alpha pod are below those for Bravo pod in the pre-regime period.

Figure B-1.1: Average Patient Age

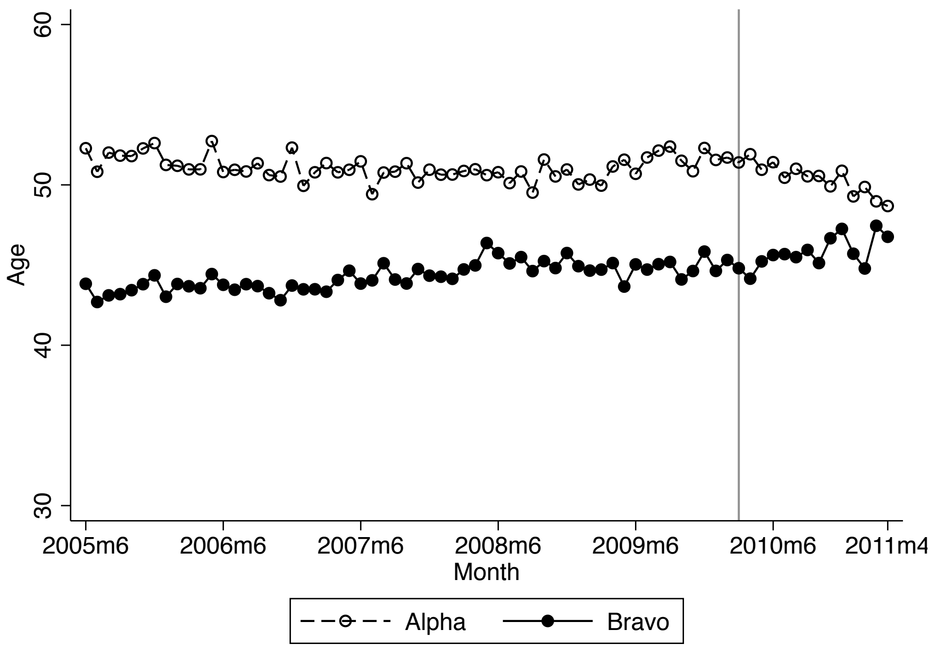

Note: This figure shows average age for patients in each pod, month, and year. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.2: Average Emergency Severity Index

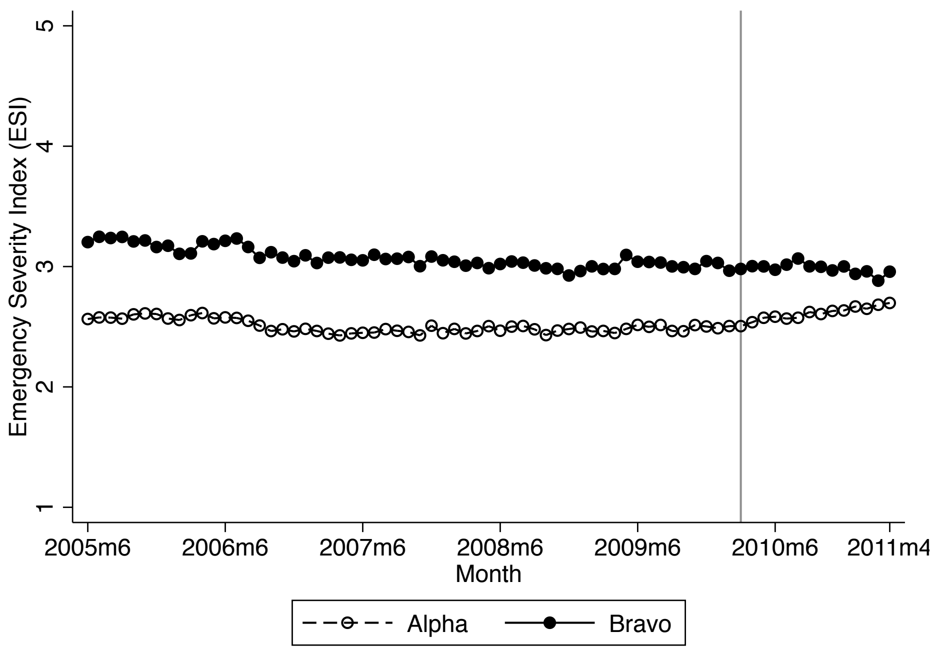

Note: This figure shows average Emergency Severity Index (ESI) for patients in each pod, month, and year. The ESI is a integer from 1 to 5, based on an algorithm accounting for patient pain, mental status, vital signs, and medical condition (Tanabe et al., 2004). An ESI of 1 is the most severe category, while an ESI of 5 is the least severe category. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.3: Average Number of Elixhauser Indices

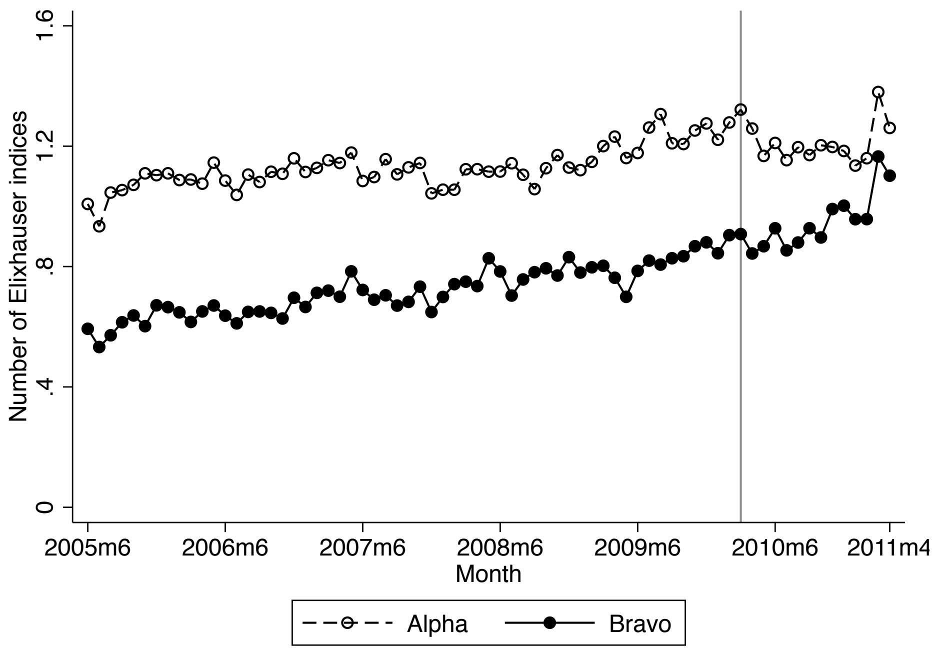

Note: This figure shows the average numbers of Elixhauser indices for patients in each pod, month, and year. Each Elixhauser index represents a clinical condition relevant for predicting patient outcomes, including congestive heart failure, diabetes, hypothyroidism, AIDS, metastatic cancer, and drug abuse (Elixhauser et al., 1998). Elixhauser indices are captured based on coding entered by health care providers in the medical record. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.4: Average Predicted Log Length of Stay

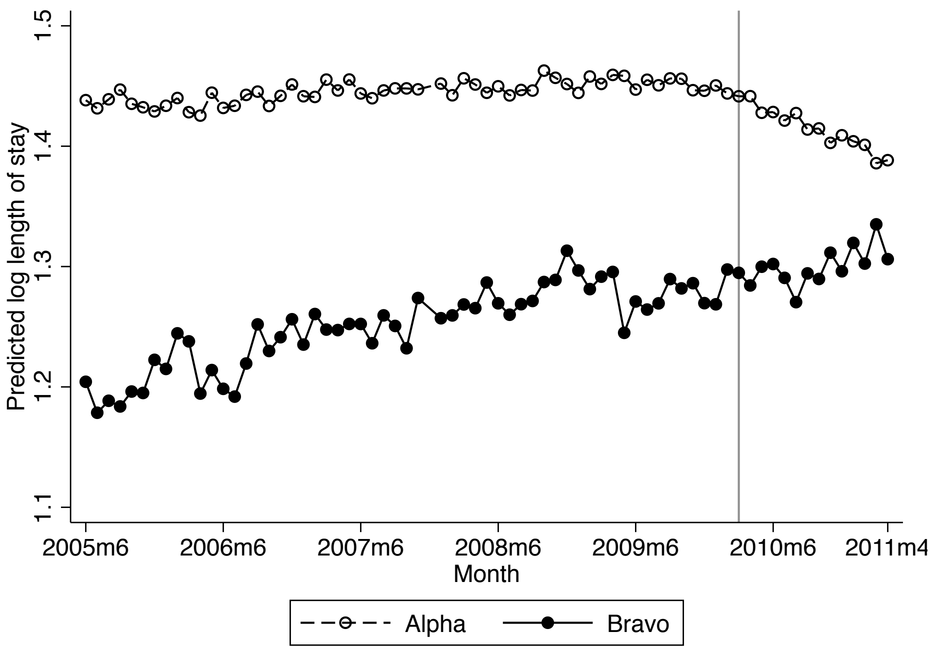

Note: This figure shows average predicted log length of stay for patients in each pod, month, and year. In the first step, Equation (B-1.1) is estimated for patients going to Alpha pod in 2005. Coefficients are then used to calculated predicted log lengths of stay for all patients. In the second step, averages for predicted log lengths of stay are computed for each pod, month, and year. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.5: Average Physician Productivity

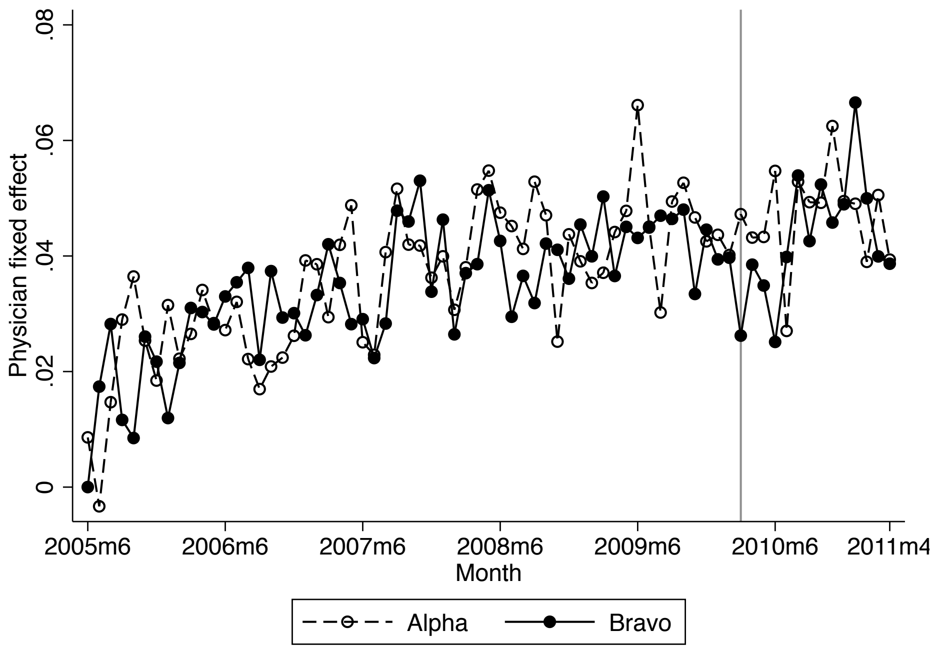

Note: This figure shows average productivity for physicians in each pod, month, and year. Physician productivity is first estimated as fixed effects in a regression of length of stay, described in Section B-1.2. In the second step, physician fixed effects are averaged for each pod, month, and year. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.6: Average Peer Productivity

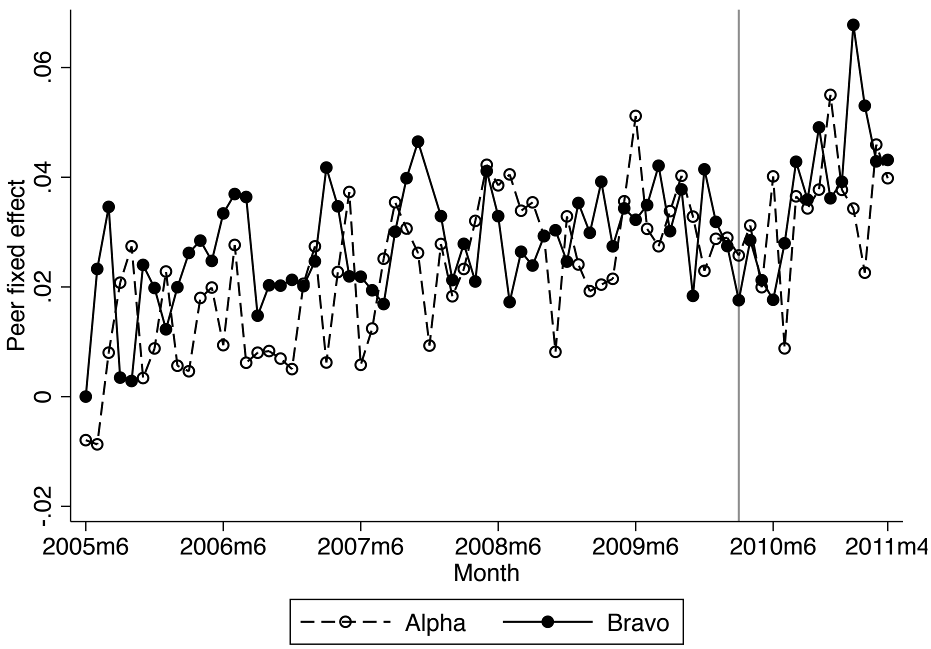

Note: This figure shows average productivity for peers (when present) in each pod, month, and year. Physician productivity is first estimated as fixed effects in a regression of length of stay, described in Section B-1.2. In the second step, fixed effects for peers are averaged for each pod, month, and year. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.7: Average Physician-peer Productivity Differences

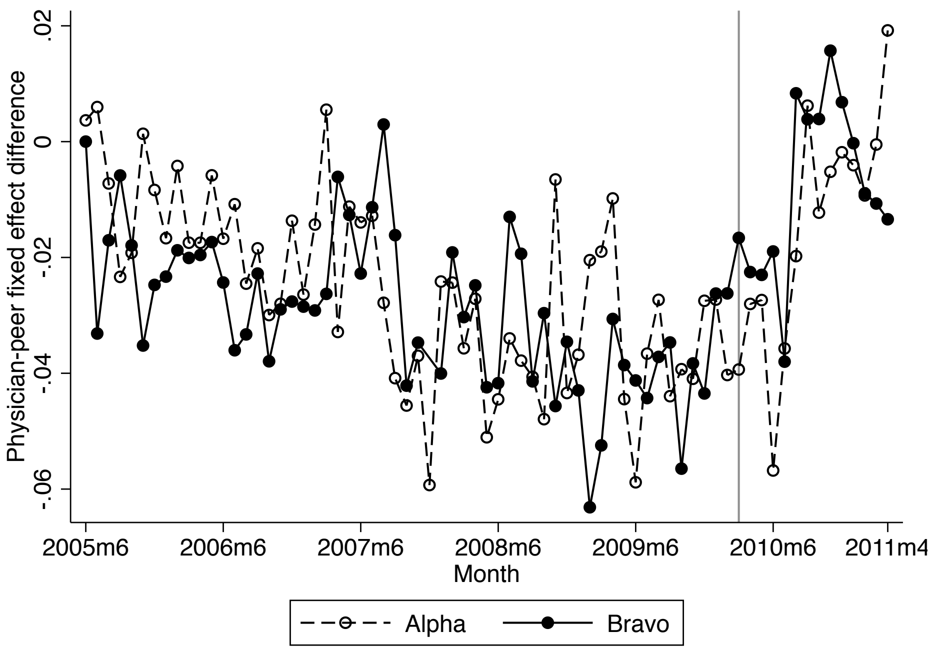

Note: This figure shows average differences in productivity between physicians and peers (when peers are present) in each pod, month, and year. Physician productivity is first estimated as fixed effects in a regression of length of stay, described in Section B-1.2. In the second step, differences in fixed effects between physicians and peers are averaged for each pod, month, and year. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.8: Average Physician Tenure

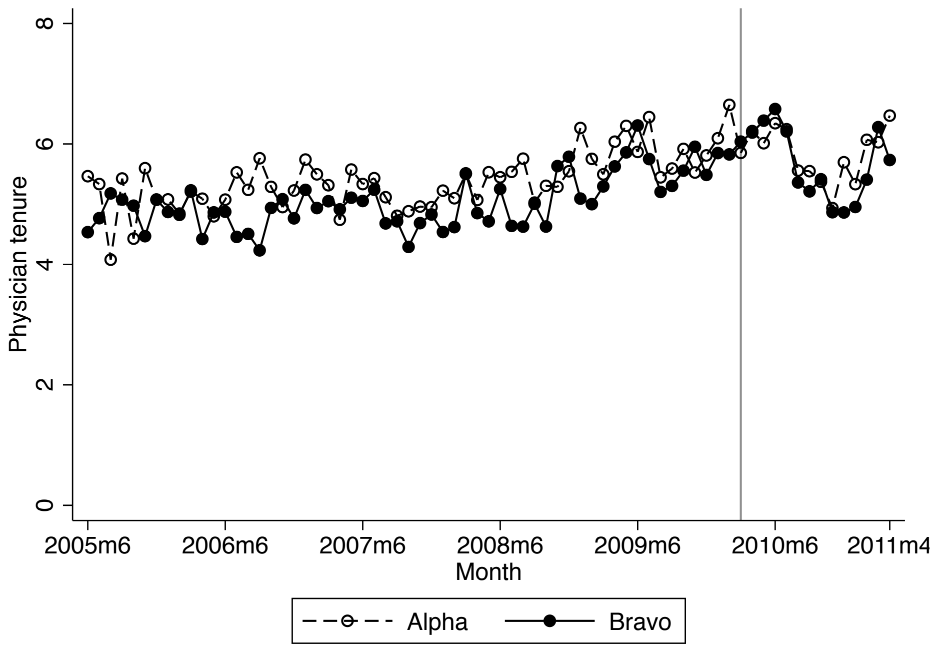

Note: This figure shows average tenure for physicians in each pod, month, and year. Physician tenure is measured by the difference between the patient visit and the physician date of hire. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.9: Average Peer Tenure

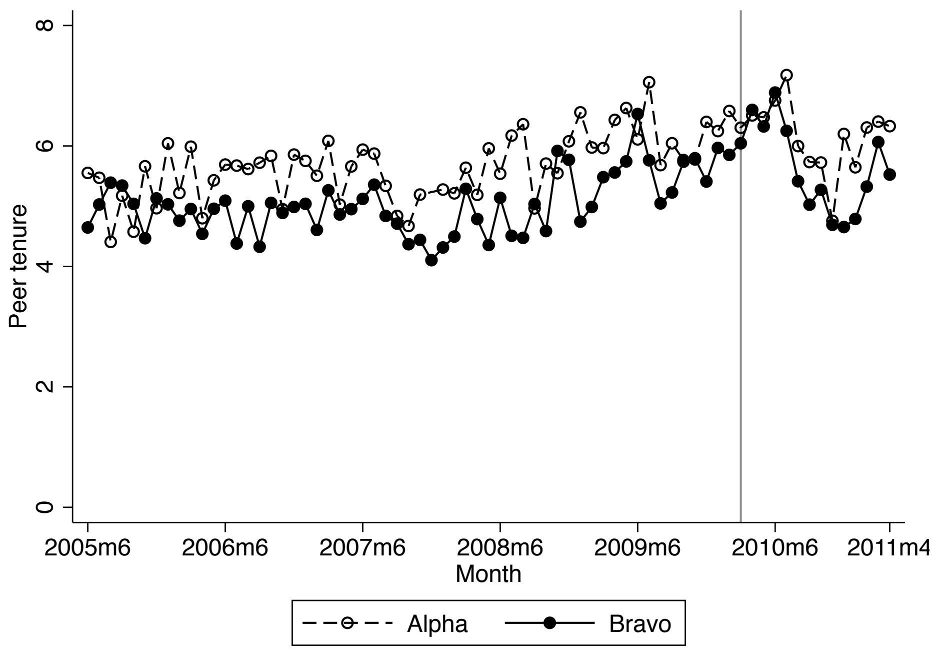

Note: This figure shows average tenure for peers (when present) in each pod, month, and year. Physician tenure is measured by the difference between the patient visit and the physician date of hire. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.10: Average Physician-peer Differences in Tenure

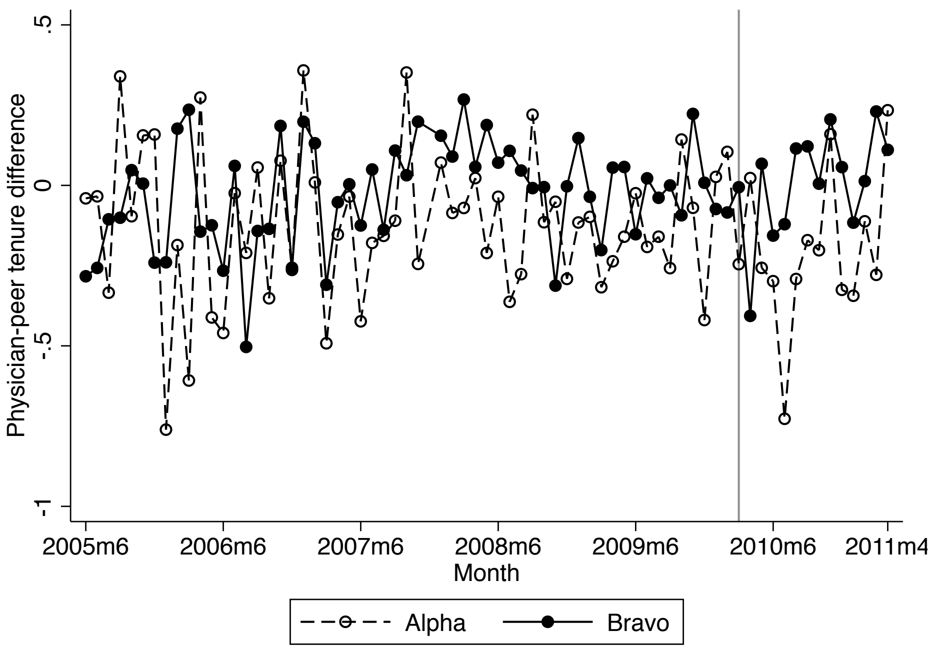

Note: This figure shows average differences in tenure between physicians and peers (when peers are present) in each pod, month, and year. Physician tenure is measured by the difference between the patient visit and the physician date of hire. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-1.11: Unadjusted Average Log Length of Stay by Pod and Month-year

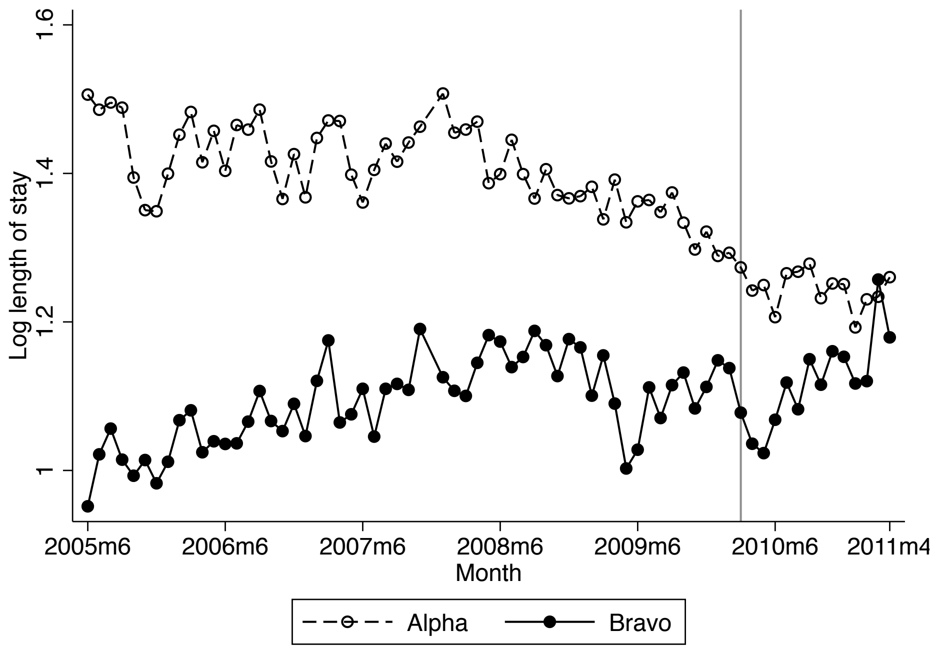

Note: This figure shows average log length of stay for each pod, month, and year, where the average is normalized to 0 for Bravo pod at the beginning of the panel. No patient characteristics or provider identities are controlled for, unlike Figure 4 in the main paper. Alpha pod averages are plotted with hollow circles; Bravo pod averages are plotted with solid circles. The vertical gray line indicates the month of the regime change of Bravo from a nurse-managed system to a self-managed system, in March 2010. Alpha was always self-managed.

Figure B-5.1: Overall Patient Volume by Hour

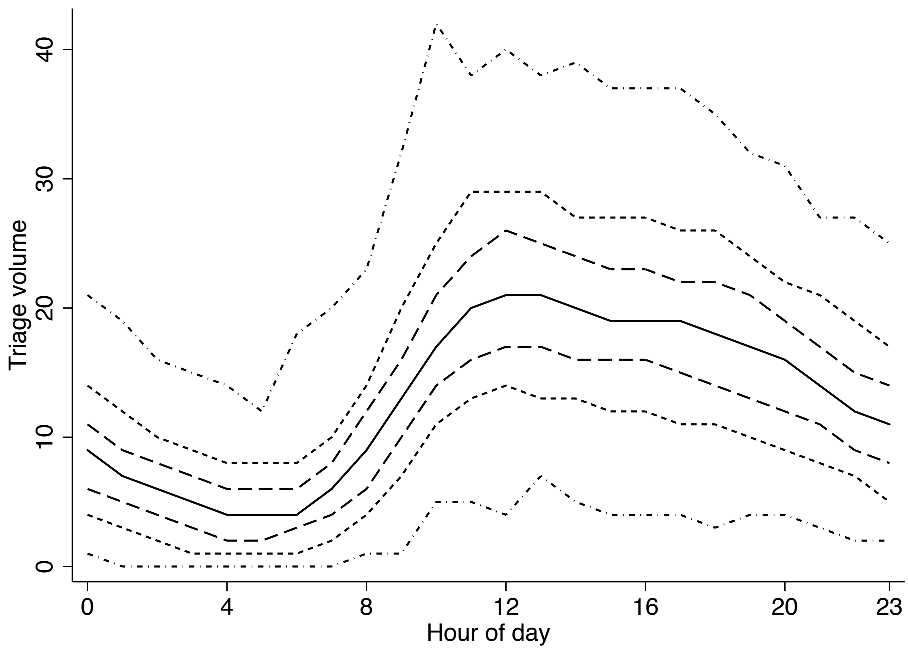

Note: This figure shows plots of overall patient volume by hour of the day. Patient volume is defined as the number of patients arriving to triage (the waiting room) in each hour. Summary statistics are then generated across these hourly observations for each hour of the day. The solid line plots median patient volume; the long-dashed lines plot 20th and 80th percentiles; the short-dashed lines plot the 5th and 95th percentiles; the dash-dotted lines plot minimum and maximum values.

Figure B-5.2: Overall Patient Volume by Day of Week

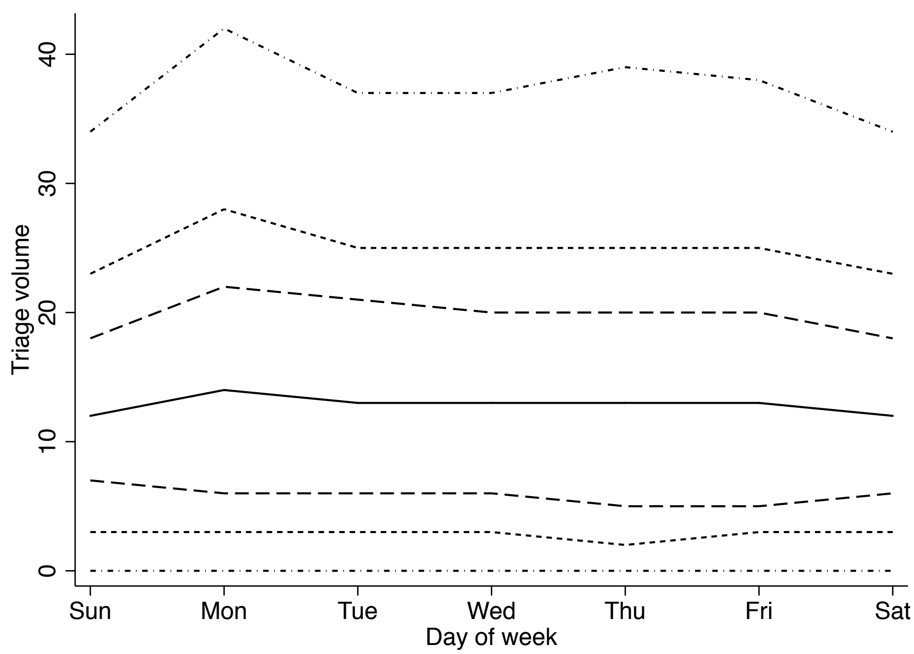

Note: This figure shows plots of overall patient volume by day of the week. Patient volume is defined as the number of patients arriving to triage (the waiting room) in each hour. Summary statistics are then generated across these hourly observations for each day of the week. The solid line plots median patient volume; the long-dashed lines plot 20th and 80th percentiles; the short-dashed lines plot the 5th and 95th percentiles; the dash-dotted lines plot minimum and maximum values.

Figure B-5.3: Overall Patient Volume by Month and Year

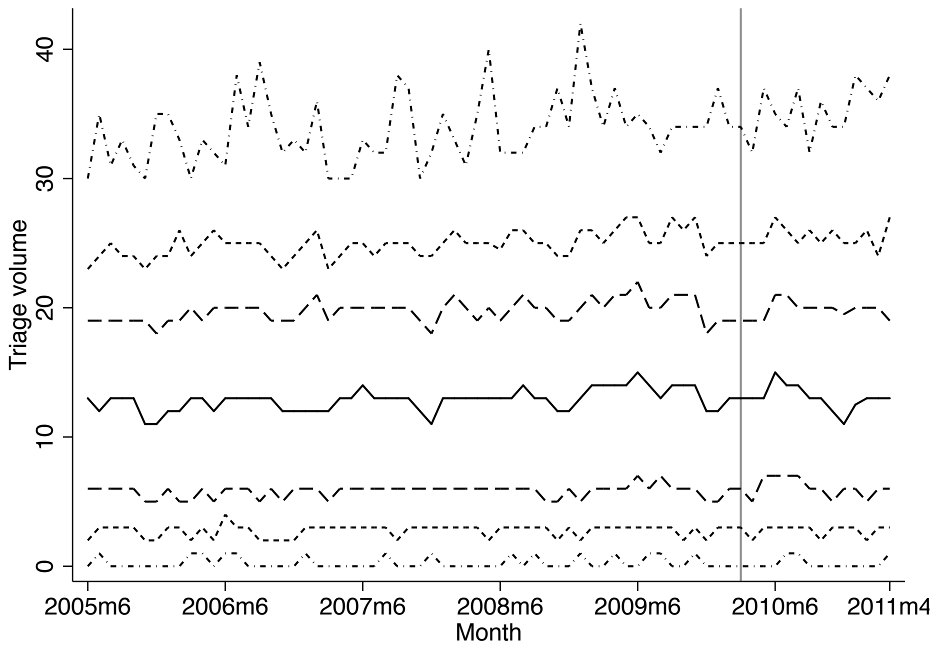

Note: This figure shows plots of overall patient volume by month and year. Patient volume is defined as the number of patients arriving to triage (the waiting room) in each hour. Summary statistics are then generated across these hourly observations for each month-year interaction. The solid line plots median patient volume; the long-dashed lines plot 20th and 80th percentiles; the short-dashed lines plot the 5th and 95th percentiles; the dash-dotted lines plot minimum and maximum values.

Figure B-5.4: Overall and Pod-specific Patient Volume

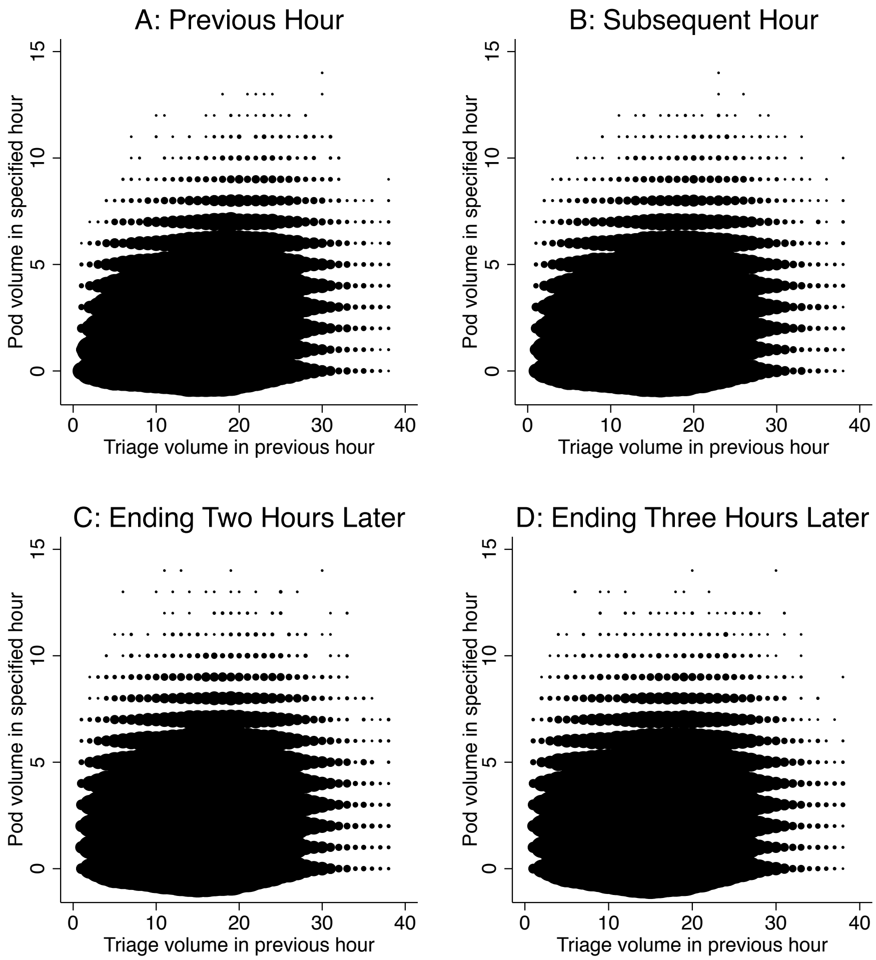

Note: This figure shows plots of overall and pod-specific patient volume. Circle sizes are proportional to the number of observations matching the specified overall and pod volumes. In all panels, overall volume is the number of patients arriving in triage in the hour prior to an index patient's pod arrival. Depending on the panel, pod-specific patient volume is shown for a specified hour, ranging from the same hour prior to an index patient's pod arrival to the hour ending three hours after the patient's pod arrival. The median time from triage to pod is about 30 minutes, while the median time from pod to discharge order is about 3.4 hours. Correlation coefficients are 0.21, 0.16, 0.13, and 0.12 for Panels A, B, C, and D, respectively.

Figure B-6.1: Density of Patient Volume to ED for High- and Low-Productivity Physicians

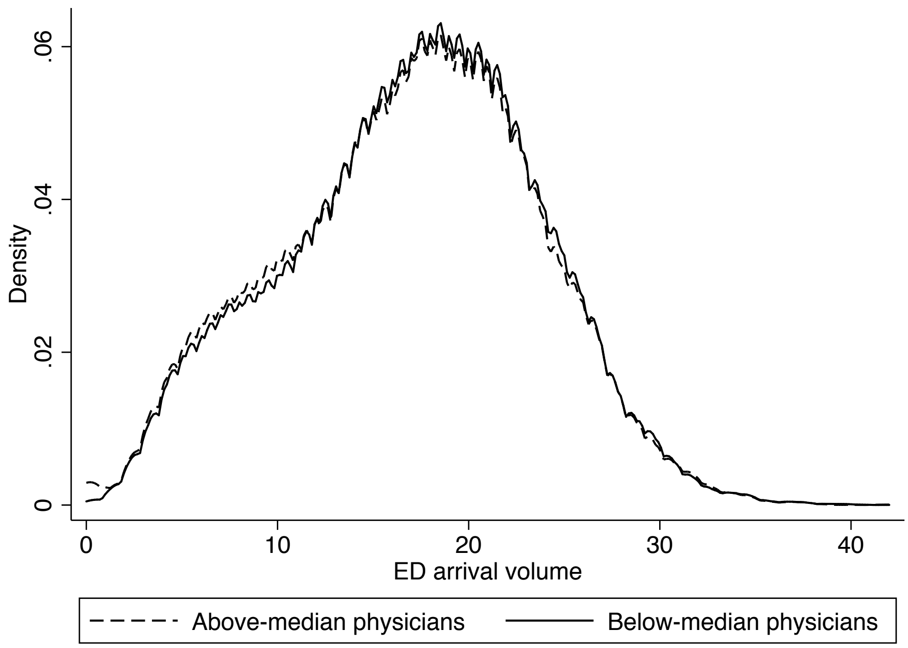

Note: This figure shows a kernel density plot of the ED arrival volume (number of patients arriving at the ED) during each hour while physicians of above- and below-median productivity are working. Dashed and solid lines show the densities for above- and below-median physicians, respectively. Physician productivity is estimated by fixed effects in a regression of length of stay, controlling for all possible interactions of team members (physician assistant or resident and nurse), coworker, pod location; patient demographics (age, sex, emergency severity index, Elixhauser comorbidities); ED arrival volume; and time dummies (month-year combination, day of the week, and hour of the day). The average difference in productivity between physicians of above- and below-median productivity is 0.28, meaning that physicians with above-average productivity take 28% less time than those with below-average productivity.

Figure B-8.1: Computer Schematic of Alpha Pod

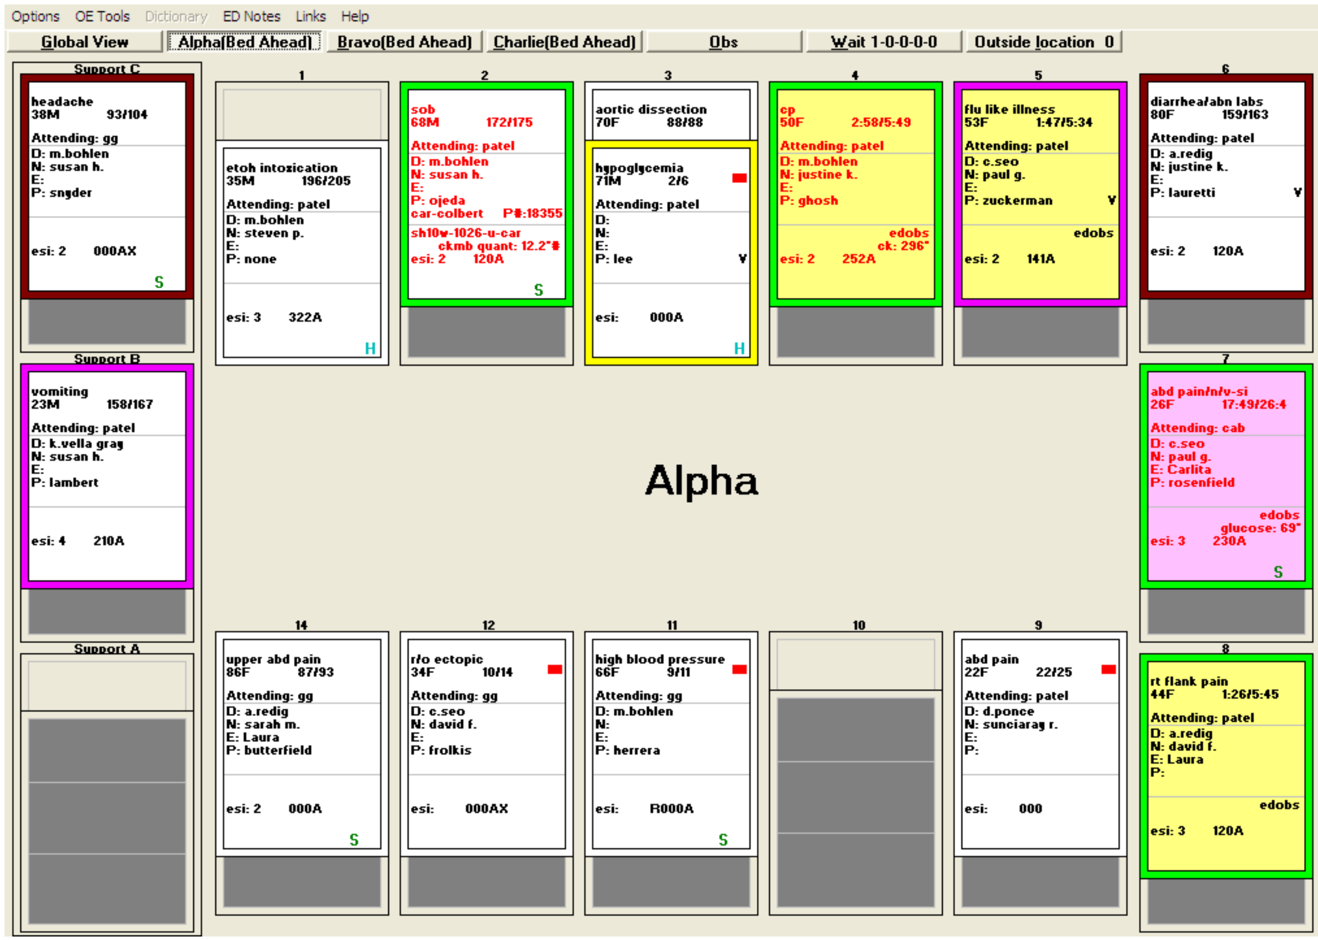

Note: This figure shows a computer screen layout of Alpha pod, which is both a spatial representation of the physical pod and the interface for physicians to select patients, examine the electronic medical record, and enter orders. Slots represent beds, with two beds per room, and filled beds are represented by slots with information. Colors represent various patient states, for example, whether an order needs to be taken off or whether the patient has been ordered for discharge. Identifying patient information has been removed here, but when displayed, such information includes patient last name, chief complaint, age, sex, physician, resident, nurse, emergency severity index, and minutes in ED (including triage) and in pod.

Figure B-8.2: Computer Schematic of Bravo Pod

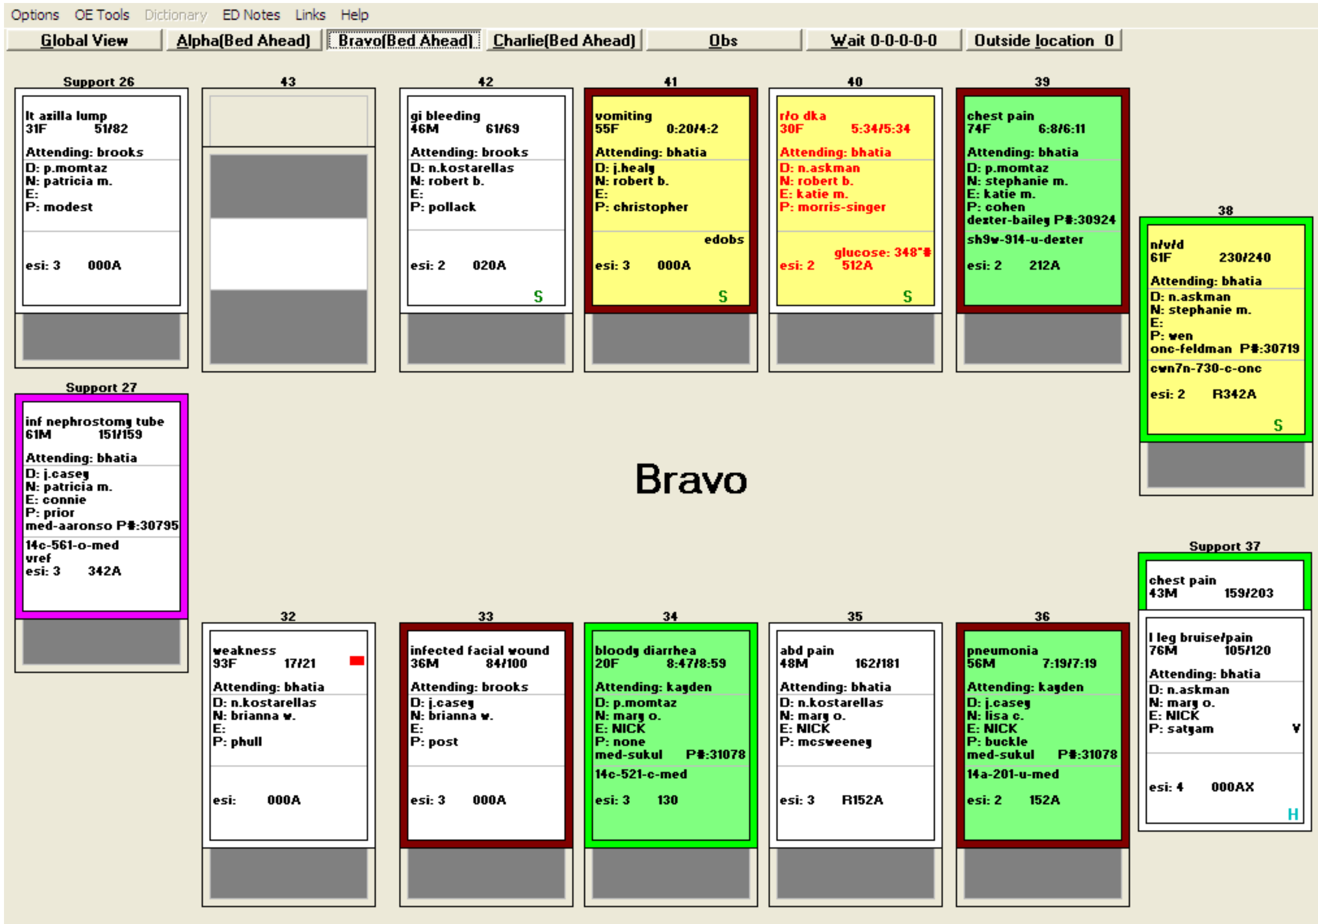

Note: This figure shows a computer screen layout of Bravo pod, which can be compared to Figure 2 of Alpha pod in the main paper. As in Figure 2, this figure is both a spatial representation of the physical pod and the interface for physicians to select patients, examine the electronic medical record, and enter orders. Slots represent beds, with two beds per room, and filled beds are represented by slots with information. Prior to the regime change in March 2010, beds in Bravo was divided into two zones – Bravo-1 and Bravo-2 – each owned by a different physician. The location and numbering of beds were stable throughout the entire study period. Colors represent various patient states, for example, whether an order needs to be taken off or whether the patient has been ordered for discharge. Identifying patient information has been removed here, but when displayed, such information includes patient last name, chief complaint, age, sex, physician, resident, nurse, emergency severity index, and minutes in ED (including triage) and in pod.

Figure B-9.1: Event Study of Foot-dragging in Both Pods

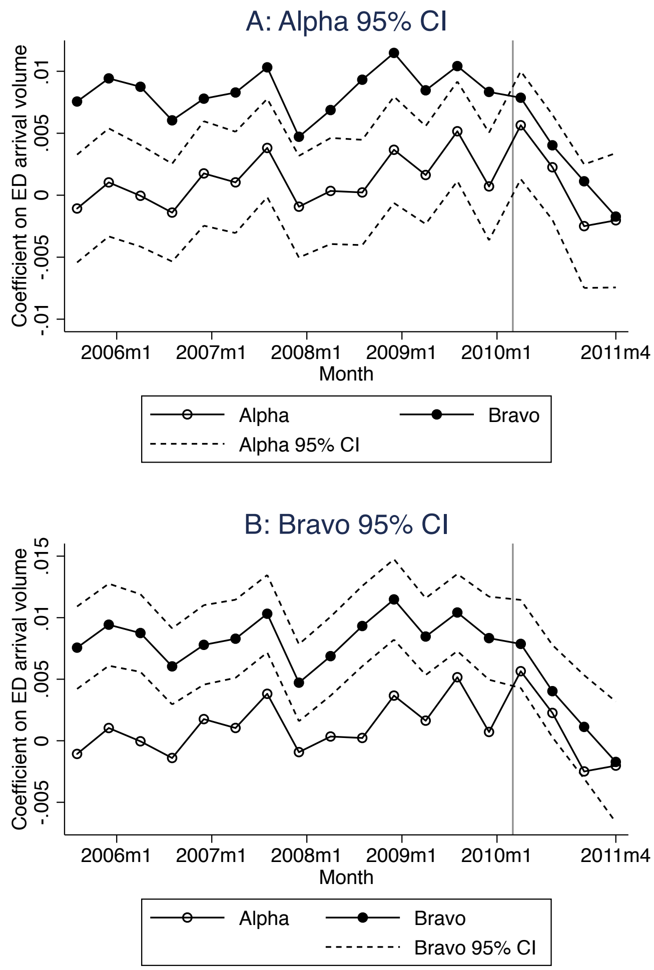

Note: This figure shows foot-dragging in both pods, as estimated by the log-length-of-stay coefficient for expected future work (measured by ED arrival volume, defined as the hourly rate of patients arriving at triage when the index patient arrives at the pod) interacted with pod identities and four-month interval dummies. These coefficients are plotted as an event study before and after the regime change of Bravo pod from a nurse-managed system to a self-managed system in March 2010, shown with a vertical gray line. Hollow and solid circles plot estimates for Alpha and Bravo pods, respectively. 95% confidence intervals, shown in dotted lines, are plotted in Panels A and B for Alpha and Bravo pods, respectively.

Figure B-9.2: New-patient Assignment Probability with Confidence Intervals

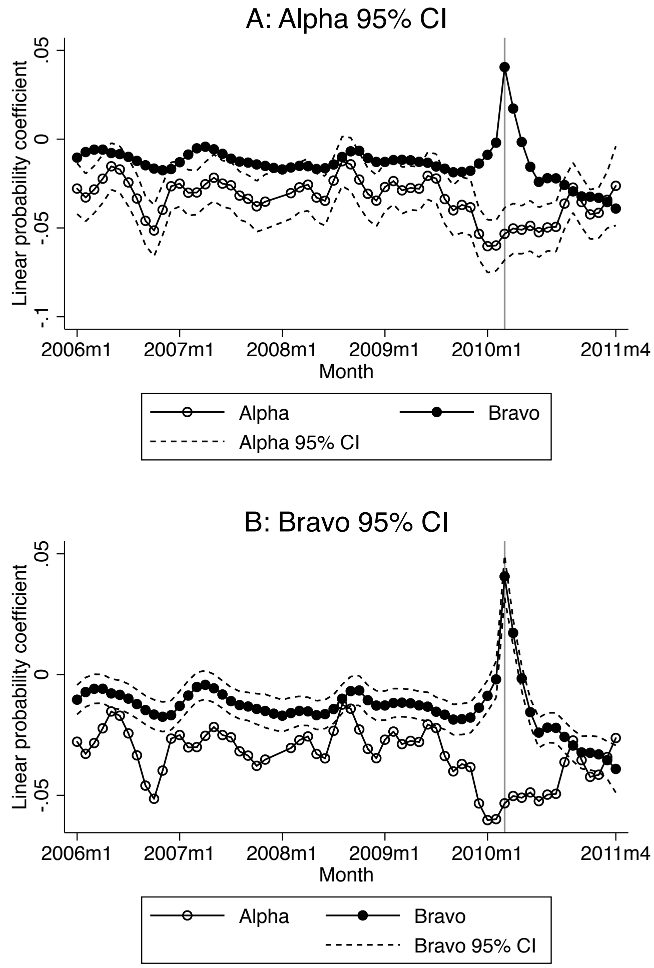

Note: This figure shows the new-patient assignment probability, as a function of relative censuses for physicians within each pod. The plotted coefficient estimates from Equation (7.1) represent the average effect on assignment probability of each additional patient on a physician's census relative to his peer's census. Hollow circles show coefficient estimates for Alpha pod, which was always self-managed. Solid circles show the coefficient estimates for Bravo pod, which switched to a self-managed system in March 2010, shown with a vertical gray line. Coefficients are estimated in a local linear regression using a triangular kernel with 90 days on each side, which may be truncated if sufficiently close to March 1, 2010 on either side. Panels A and B show confidence intervals for Alpha and Bravo coefficients, respectively.

Figure B-9.3: New-patient Assignment Probability, Alternative Specification

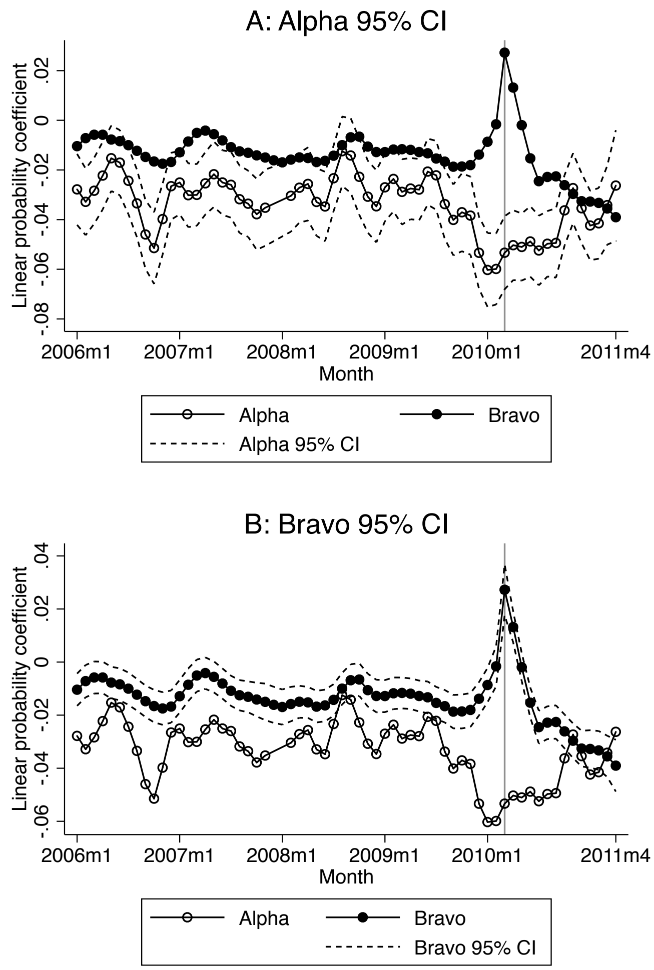

Note: This figure shows the new-patient assignment probability, as a function of relative censuses for physicians within each pod. Coefficients are estimated using Equation (B-9.1), slightly modified from Equation (7.1) by allowing for zone-based norms to continue in Bravo after the regime change. Coefficients represent the average effect on assignment probability of each additional patient on a physician's census relative to his peer's census. Hollow circles show coefficient estimates for Alpha pod, which was always self-managed. Solid circles show the coefficient estimates for Bravo pod, which switched to a self-managed system in March 2010, shown with a vertical gray line. Coefficients are estimated in a local linear regression using a triangular kernel with 90 days on each side, which may be truncated if sufficiently close to March 1, 2010. Panels A and B show confidence intervals for Alpha and Bravo coefficients, respectively.

Figure B-9.4: New-patient Assignment Probability, Kernel Regression

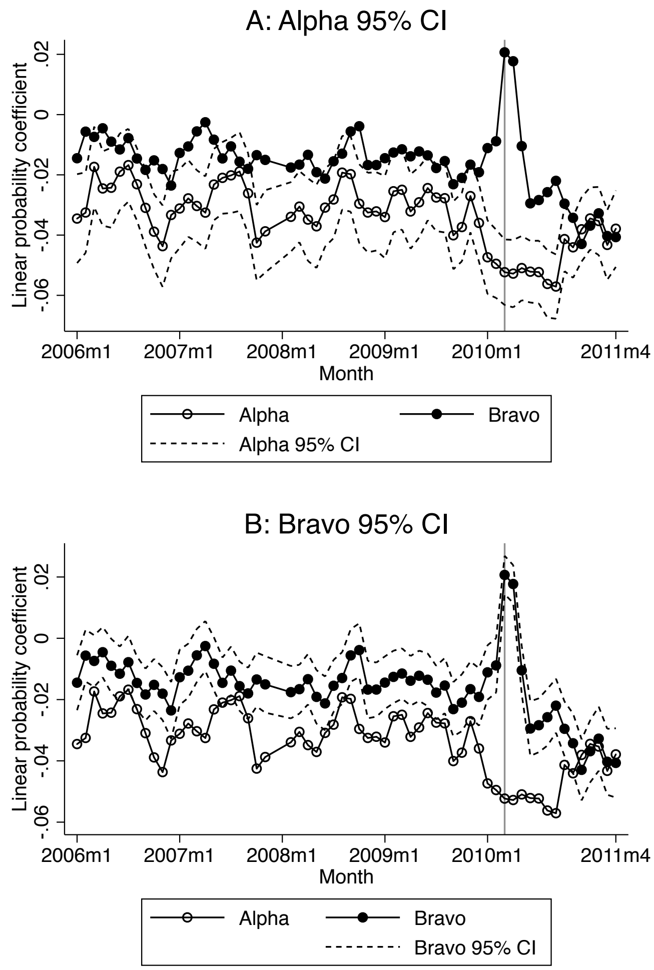

Note: This figure shows the new-patient assignment probability, as a function of relative censuses for physicians within each pod. The plotted coefficient estimates from Equation (7.1) represent the average effect on assignment probability of each additional patient on a physician's census relative to his peer's census. Hollow circles show coefficient estimates for Alpha pod, which was always self-managed. Solid circles show the coefficient estimates for Bravo pod, which switched to a self-managed system in March 2010, shown with a vertical gray line. Coefficients are estimated in a kernel regression using a triangular kernel with 45 days on each side, which may be truncated if sufficiently close to March 1, 2010. Panels A and B show confidence intervals for Alpha and Bravo coefficients, respectively.

Table B-1.1: Average Patient Characteristics Assigned to Alpha or Bravo Pod

Patient characteristic	Alpha	Bravo
Age	49.9	44.5
	(19.3)	(18.8)
Emergency severity index	2.58	3.06
	(0.725)	(0.785)
White	0.527	0.458
	(0.499)	(0.498)
Black or African-American	0.231	0.247
	(0.421)	(0.431)
Spanish speaking	0.086	0.112
	(0.280)	(0.315)
Female and age $<$ 35 years	0.269	0.249
	(0.374)	(0.432)

Note: This table reports average patient characteristics for patients being assigned to Alpha and Bravo pods. Alpha pod was always opened 24 hours, while Bravo pod always closed at night. The emergency severity index (ESI) ranges from 1 to 5, and a lower ESI represents a more severe patient. Standard deviations are reported in parentheses.

Table B-6.2: Average Patient Characteristics Available to Physicians with Preference for or against Patient Type

	Nurse-managed system		Self-managed system	
Patient characteristic	Physicians with preference for	Physicians with preference against	Physicians with preference for	Physicians with preference against
Age	43.6	43.8	48.8	48.8
	(18.4)	(18.5)	(19.6)	(19.6)
Emergency severity index	3.11	3.12	2.73	2.73
	(0.774)	(0.771)	(0.773)	(0.776)
White	0.448	0.445	0.512	0.509
	(0.497)	(0.497)	(0.500)	(0.500)
Black or African-American	0.250	0.249	0.233	0.232
	(0.433)	(0.432)	(0.423)	(0.422)
Spanish speaking	0.114	0.114	0.097	0.096
	(0.318)	(0.317)	(0.297)	(0.294)
Female and age $<$ 35 years	0.266	0.261	0.185	0.184
	(0.442)	(0.440)	(0.388)	(0.387)

Note: This table reports average patient characteristics available to physicians with a preference for or against a given patient type, while working in a nurse-managed or self-managed system. In the nurse-managed system, patients are assigned; on the self-managed system, patients are made available by the triage nurse. I construct measures of physician preference by (1) estimating a linear probability model of patient choice using observations in the behavioral system and including physician-specific coefficients on patient characteristics, and (2) selecting physicians with high or low coefficients as having a preference for or against each patient characteristic, respectively. Once I arrive at measures of physician preference for each patient characteristic, I describe the average for that patient characteristic for patients available to choose from for physicians with preferences for or against that characteristic. The emergency severity index (ESI) ranges from 1 to 5, and a lower ESI represents a more severe patient. Standard deviations are reported in parentheses.

Table B-7.1: Peer Effects on Log Length of Stay

	(1)	(2)	(3)	(4)	(5)	(6)
Peer effect	0.115***	0.087*	0.106***			
	(0.035)	(0.049)	(0.032)			
Peer effect, faster than median				0.142***	0.079	
				(0.047)	(0.069)	
Peer effect, slower than median				0.039	0.106	
				(0.098)	(0.124)	
Peer effect × self-managed			-0.053			
			(0.053)			
Peer effect, faster than median × self-managed						0.109*
						(0.062)
Peer effect, slower than median× self-managed						-0.081
						(0.115)
Peer effect, faster than median × nurse-managed						0.128***
						(0.045)
Peer effect, slower than median× nurse-managed						0.048
						(0.093)
Self-managed			-0.081***			-0.061**
			(0.029)			(0.024)
Sample	Nurse-managed	Self-managed	Pooled	Nurse-managed	Self-managed	Pooled
Number of observations	129,274	129,798	259,072	129,274	129,798	259,072
Adjusted R-squared	0.485	0.378	0.404	0.485	0.378	0.404
Sample log length of stay mean	1.006	1.263	1.135	1.006	1.263	1.135

Note: I estimate physician peer effects depending on the organizational structure. I first estimate fixed effects in a regression of productivity on physician identities, adjusting for all possible interactions with resident, nurse, peer, and pod location. I then use these fixed effects for observations in the mechanical system where there is another physician in the same pod present to estimate a regression of log length of stay with peer effects. Models (1) and (4) are estimated for observations from nurse-managed teams, models (2) and (5) are estimated for observations from self-managed teams, and models (3) and (6) estimated from pooled observations. Models (4) to (6) split peer effects for peers that are slower or faster than median. All regressions in both stages are adjusted for hour of the day, day of the week, month-year dummies, Elixhauser score dummies, emergency , ED arrival volume, patient demographics, and patient triage time. All models are clustered for physician-resident-nurse trio identities.

Table B-7.2: Effect of Peer Relationships on Foot-dragging

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Log length of stay			Log length of stay			Log length of stay		
Type of peer	Greater tenure than physician			Above-median productivity			$\geq$ 60 hours prior work together		
ED volume	0.008***	0.003***	0.009***	0.008***	0.004***	0.009***	0.006***	0.004***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
ED volume × peer type	-0.004***	0.000	-0.003***	-0.002*	-0.001	-0.001	0.002	-0.002	0.002*
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Self-managed × ED volume			-0.008***			-0.007***			-0.005***
			(0.001)			(0.001)			(0.001)
Self-managed × ED volume × peer type			0.002			0.001			-0.004*
			(0.002)			(0.002)			(0.002)
Self-managed × peer type			-0.031**			-0.027**			0.005
			(0.014)			(0.013)			(0.014)
Sample	Nurse-managed	Self-managed	Pooled	Nurse-managed	Self-managed	Pooled	Nurse-managed	Self-managed	Pooled
Adjusted R-squared	0.445	0.364	0.376	0.445	0.364	0.376	0.445	0.364	0.376
Sample mean log length of stay (log hours)	0.936	1.251	1.097	0.926	1.179	1.064	0.926	1.179	1.064
Sample mean ED volume	17.30	15.08	16.16	17.12	14.22	15.53	17.12	14.22	15.53

Note: This table reports effect of expected future work, interacted with a peer type, as by Equation (B-7.3). Expected future work is measured by ED arrival volume ("ED volume" for brevity), defined as the number of patients arriving at ED triage during the hour prior to the index patient's arrival at the pod. All observations require the presence of a peer. Peers with greater tenure, those with high productivity (faster-than-median fixed effect for lengths of stay), and familiar peers (peers with at least 60 hours of history, approximately the 75th percentile, working with the index physician) are considered. ED volume is demeaned. Coefficients for the direct effect of the peer type and pooled coefficients for self-managed are omitted for brevity. All models control for time categories (month-year, day of the week, and hour of the day dummies), pod (when applicable), patient demographics (age, sex, race, and language), patient clinical information (Elixhauser comorbidity indices, emergency severity index), triage time, and physician-resident-nurse interactions. The nurse-managed, self-managed, and pooled samples had 121,024, 126,264, and 247,288 observations, respectively. All models are clustered by physician. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table B-7.1: Bed Locations

Pod or zone	Bed locations
Alpha	1, 1H, 2, 2H, 3, 3H, 4, 4H, 5, 5H, 6, 6H, 7, 7H, 8, 8H, 9, 9H, 10, 10H, 11, 11H, 12, 12H, 14, 14H, A Hall-A, A Hall-AH, A Hall-B, A Hall-BH, A Hall-C, A Hall-CH, A-HALL
Bravo-1	25, 25H, 26, 26H, 27, 27H, 28, 29, 30, 31, 32, 32H, 43, 43H, E Hall-A, E Hall-AH, E Hall-B, E Hall-BH
Bravo-2	33, 33H, 34, 34H, 35, 35H, 36, 36H, 37A, 37AH, 37B, 37BH, 38, 38H, 39, 39H, 40, 40H, 41, 41H, 42, 42H, B-HALL

Note: This table details bed locations for beds considered in each pod or zone. Alpha and Bravo pods are separate locations in the ED. Prior to March 2010, Bravo pod was divided into two non-overlapping zones – Bravo-1 and Bravo-2 – and physicians in Bravo were scheduled to care for patients in one of these two zones. Alpha pod is shown in Figure 2 in the main paper; Bravo pod is shown in Figure B-8.2. Some of the bed locations are visible in these figures, while others are not annotated, as they may represent a "virtual bed," such as the hallway space in front of a bed, for example. All bed locations correspond to unique electronic designations for a physical space in the ED. There was no change in the physical arrangement of non-virtual beds in either pod throughout the sample period.

Table B-7.2: Alpha Bed Visits

Bed	Dates observed	Number of visits
1	6/1/2005-4/30/2011	15,560
1H	1/22/2008-4/30/2011	11,576
2	6/1/2005-4/30/2011	2,220
2H	1/22/2008-4/30/2011	11,625
3	6/1/2005-4/30/2011	1,976
3H	1/22/2008-4/30/2011	12,279
4	6/1/2005-4/30/2011	3,112
4H	1/22/2008-4/30/2011	12,769
5	6/1/2005-4/30/2011	5,358
5H	1/22/2008-4/30/2011	5,380
6	6/1/2005-4/30/2011	15,945
6H	1/21/2008-4/30/2011	4,885
7	6/1/2005-4/30/2011	13,385
7H	1/21/2008-4/29/2011	4,210
8	6/1/2005-4/30/2011	13,103
8H	1/21/2008-4/28/2011	4,642
9	6/1/2005-4/30/2011	12,722
9H	1/21/2008-4/30/2011	3,638
10	6/1/2005-4/30/2011	11,995
10H	1/21/2008-4/30/2011	1,642
11	6/1/2005-4/30/2011	10,634
11H	1/21/2008-4/30/2011	1,279
12	6/1/2005-4/30/2011	10,423
12H	1/22/2008-4/29/2011	1,137
14	6/1/2005-4/30/2011	11,819
14H	1/21/2008-4/30/2011	4,493
A Hall-A	1/21/2008-4/20/2011	6,354
A Hall-AH	1/21/2008-4/17/2011	1,692
A Hall-B	1/21/2008-4/20/2011	5,652
A Hall-BH	1/21/2008-4/15/2011	1,429
A Hall-C	1/21/2008-4/19/2011	6,432
A Hall-CH	1/21/2008-4/19/2011	1,735
A-HALL	6/1/2005-1/21/2008	24,800

Note: This table shows the number of visits for which a patient is recorded to spend some time in a given bed in Alpha pod. Date ranges are also given for the first and last of these visits. Alpha pod is shown in Figure 2 in the main paper. Some of the bed locations are visible in this figures, while others are not annotated, as they may represent a "virtual bed," such as the hallway space in front of a bed, for example. Hallway beds are designated with an "H" and sometimes a number signifying the closest numbered room. All bed locations correspond to unique electronic designations for a physical space in the ED. There was no change in the physical arrangement of non-virtual beds throughout the sample period.

Table B-7.3: Bravo-1 Bed Visits

Bed	Dates observed	Number of visits
25	6/1/2005-2/15/2011	7,415
25H	1/22/2008-4/19/2011	1,229
26	6/1/2005-4/30/2011	9,043
26H	1/22/2008-4/28/2011	1,642
27	6/1/2005-4/30/2011	9,812
27H	1/22/2008-4/30/2011	1,474
28	6/1/2005-2/7/2010	5,719
29	6/1/2005-2/3/2010	7,051
30	6/1/2005-3/29/2010	8,593
31	6/1/2005-1/26/2010	9,808
32	6/1/2005-4/30/2011	9,343
32H	1/22/2008-4/30/2011	2,278
43	6/1/2005-4/30/2011	9,970
43H	1/22/2008-4/30/2011	2,073
E Hall-A	1/24/2008-4/5/2010	1,936
E Hall-AH	1/22/2008-4/5/2010	1,768
E Hall-B	1/24/2008-4/5/2010	1,136
E Hall-BH	1/22/2008-4/5/2010	845

Note: This table shows the number of visits for which a patient is recorded to spend some time in a given bed in Bravo-1 zone within pod. Date ranges are also given for the first and last of these visits. Bravo-1 and Bravo-2 designate adjoining but non-overlapping zones within Bravo, each containing beds that are owned by a physician working in shifts labeled with "Bravo-1" or "Bravo-2," respectively. Bravo pod is shown in Figure B-8.2. Some of the bed locations are visible in this figures, while others are not annotated, as they may represent a "virtual bed," such as the hallway space in front of a bed, for example. Hallway beds are designated with an "H" and sometimes a number signifying the closest numbered room. Beds 28 to 31 are spaces without walls but separated by curtains. All bed locations correspond to unique electronic designations for a physical space in the ED. There was no change in the physical arrangement of non-virtual beds throughout the sample period.

Table B-7.4: Bravo-2 Bed Visits

Bed	Dates observed	Number of visits
33	6/1/2005-4/30/2011	11,613
33H	1/22/2008-4/29/2011	3,982
34	6/1/2005-4/30/2011	12,186
34H	1/22/2008-4/29/2011	4,504
35	6/1/2005-4/30/2011	11,186
35H	1/22/2008-4/30/2011	2,323
36	6/1/2005-4/30/2011	10,669
36H	1/22/2008-4/30/2011	1,905
37A	6/1/2005-4/19/2011	11,199
37A H	1/23/2008-4/19/2011	1,682
37B	6/1/2005-4/19/2011	10,357
37B H	1/22/2008-4/18/2011	1,036
38	6/1/2005-4/30/2011	11,492
38H	1/22/2008-4/29/2011	1,956
39	6/1/2005-4/30/2011	13,067
39H	1/22/2008-4/30/2011	5,715
40	6/1/2005-4/30/2011	12,772
40H	1/22/2008-4/28/2011	2,678
41	6/1/2005-4/30/2011	13,047
41H	1/22/2008-4/30/2011	2,726
42	6/1/2005-4/30/2011	12,915
42H	1/22/2008-4/30/2011	4,279
B-HALL	6/1/2005-1/21/2008	14,771

Note: This table shows the number of visits for which a patient is recorded to spend some time in a given bed in Bravo-2 zone within pod. Date ranges are also given for the first and last of these visits. Bravo-1 and Bravo-2 designate adjoining but non-overlapping zones within Bravo, each containing beds that are owned by a physician working in shifts labeled with "Bravo-1" or "Bravo-2," respectively. Bravo pod is shown in Figure B-8.2. Some of the bed locations are visible in this figures, while others are not annotated, as they may represent a "virtual bed," such as the hallway space in front of a bed, for example. Hallway beds are designated with an "H" and sometimes a number signifying the closest numbered room. All bed locations correspond to unique electronic designations for a physical space in the ED. There was no change in the physical arrangement of non-virtual beds throughout the sample period.

Table B-9.1: Effect of Expected Future Work on Other Outcome and Process Measures

	Outcomes				Process Measures			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	30-day mortality	Hospital Admissions	14-day return visits	Log total costs	Total order count	Laboratory order count	Medication order count	CT or MRI ordered
ED volume	-0.0001	-0.0006*	-0.0002	-0.0022*	-0.0117*	-0.0041	-0.0052**	0.0002
	(0.0001)	(0.0004)	(0.0002)	(0.0012)	(0.0071)	(0.0042)	(0.0023)	(0.0003)
ED volume × self-managed	0.0002	-0.0006	0.0001	-0.0011	-0.0094	-0.0027	-0.0024	-0.0002
	(0.0003)	(0.0004)	(0.0003)	(0.0027)	(0.0101)	(0.0060)	(0.0033)	(0.0005)
Self-managed	-0.001	0.004	-0.011	-0.0062	-0.526*	-0.060	-0.109	-0.021
	(0.006)	(0.012)	(0.009)	(0.0642)	(0.280)	(0.167)	(0.090)	(0.013)
Adjusted R-squared	0.338	0.461	-0.038	0.526	0.548	0.489	0.402	0.295
Sample mean outcome	0.019	0.265	0.066	6.684	13.267	5.337	2.688	0.246

Note: This table shows the effect of expected future work on other outcomes and process measures, estimated by Equation (5.2). Log total costs include all direct costs incurred from the encounter, including any from admissions. Expected future work is measured by ED arrival volume ("ED volume" for brevity), defined as the number of patients arriving at ED triage during the hour prior to the index patient's arrival at the pod. Models do not include pod-level prior patient volume; results are unchanged when including this. All models control for time categories (month-year, day of the week, and hour of the day dummies), pod, patient demographics (age, sex, race, and language), patient clinical information (Elixhauser comorbidity indices, emergency severity index), triage time, and physician-resident-nurse interactions. Results are insensitive to controlling for pod-level patient volume. All models are clustered by physician. All models have 289,132 observations, except for model (4), which has 269,905 observations. The sample mean patient volume is 15.53 for all models, except for the model (4), for which it is 15.48. * significant at 10%; ** significant at 5%; *** significant at 1%.