HEALTH SPILLOVERS AMONG HOSPITAL PATIENTS

Evidence from Roommate Assignments

OLGA YAKUSHEVA

ABSTRACT
This study examines health spillovers among hospital patients using a natural experiment of conditionally random hospital roommate assignments and electronic medical records data for a sample of 3,793 patients and their roommates. The study finds that positive health spillover effects may play an important role in reducing the amount of care the patient requires during hospitalization, shortening length of hospital stay, and lowering hospitalization cost, with little evidence of negative health effects for most patients. Mechanisms of roommate effects and implications for health-care delivery and policy are also considered.

KEYWORDS: peer effects, hospital roommates, natural experiment
JEL CLASSIFICATION: I1

I. Introduction

Social network effects, or peer effects, are defined as a tendency of an individual to conform to, adopt, or imitate the behaviors of other individuals within some social contexts; they have been reported in many health behaviors and outcomes, including fertility, obesity, mental health, utilization of health-care services, and even mortality (Christakis and Fowler 2007, 2008, 2012; Cacioppo, Fowler, and Christakis 2009; Fowler and Christakis 2008; Eisenberg, Downs, and Golberstein 2012; Eisenberg, Golberstein, and Whitlock 2014; Eagly and Chrvala 1986; Mears, Ploeger, and Warr 1998; Kapinos and Yakusheva 2011; Yakusheva, Kapinos, and Weiss 2011; Yakusheva, Kapinos, and Eisenberg 2014; Yakusheva and Fletcher 2014). In the presence of network externalities, health policy can generate multiplier effects reaching beyond the group initially targeted by the policy (Fowler and Christakis 2008).

A small number of studies examining patient-to-patient peer effects in acute care have been published to date. A descriptive study of patient experiences in multi-occupancy rooms in Norway describes substantial amounts of interaction among patients—including sharing information about their illnesses, treatments, or staff—and even offering help with daily tasks (Album 1989, 2010). In other studies, assigning a preoperative patient to share a room with a postoperative patient significantly reduced anxiety and sped up postsurgical recovery among coronary-bypass patients (Kulik and Mahler 1987; Kulik, Mahler, and

Olga Yakusheva (yakush@umich.edu), School of Nursing and School of Public Health, Institute for Healthcare Policy and Innovation, University of Michigan.

doi:10.1162/ajhe_a_00068
American Journal of Health Economics 3(1): 76–107

Moore 1996). While these findings point to the potential existence of peer effects among patients in acute care settings, they are largely descriptive with limited potential for causal inference.

Empirical estimation of causal peer effects faces several identification challenges first laid out by Manski (1993). Sources of bias include unobserved peer selection (also referred to as homophily) and spurious peer correlations driven by shared exposure to unobserved environmental and contextual effects. Several peer effect studies took advantage of a natural experiment involving some sort of a plausibly random peer assignment (e.g., college roommate assignments, military squadron assignments) to overcome the selection bias and used preassignment peer characteristics to avoid confounding from shared environmental exposures (Carrell, Malmstrom, and West 2008; Carrell, Hoekstra, and West 2011; Carrell, Sacerdote, and West 2013; Duflo, Dupas, and Kremer 2011; Sacerdote 2001; Yakusheva, Kapinos, and Weiss 2011; Yakusheva and Fletcher 2014; Yakusheva, Kapinos, and Eisenberg 2014).

This study contributes to the literature by examining peer effects among hospital patients and their roommates, using a unique data set from a large acute care hospital. The study relies on plausibly random variation in roommate characteristics, which allows me to deal with the above-mentioned biases. I find that an acute care setting could give rise to significant patient-to-patient effects: compared with patients with very sick roommates, patients who shared a room with less acutely ill roommates received care less frequently, had shorter length of stay, and lower costs, without obvious negative effects on clinical condition at discharge or increased risk of being readmitted. The benefits of having a healthy roommate were particularly large for female patients and patients admitted to the hospital in very poor clinical condition. These findings suggest that clinical interventions in acute care may be subject to significant policy multipliers through patient-to-patient spillover effects, and point to potential value of optimal roommate matching in lowering costs of hospital care with little or no negative effect on patient outcomes.

II. Peer Influence in Acute Care

Peer influence, or an endogenous social effect, is defined as a situation where the propensity of an individual to behave in some way varies with the prevalence of that behavior in some reference group containing the individual (Manski 1993). In this study, I will refer to the patient who is being influenced as "the subject patient" or, simply, "the patient," and to the patients who exert influence on the subject patient as "the roommates."

Peer influence is hypothesized as being broadly due to a combination of the social norms effect (i.e., a situation where an individual's private net marginal benefit of engaging in a behavior increases with the number of peers who engage in the same behavior) and the knowledge transfer effect (i.e., a situation where an individual learns from his or her peers about the adoption of positive behaviors or avoidance of negative behaviors) (Case and Katz 1991; Asphjell, Hensvik, and Nilsson 2013; Kuziemko 2006; Yakusheva and Fletcher 2014). For example, the subject patient may feel better when he or she observes the roommates coping well with a similar illness (i.e., social norm effects), or the patient may learn from his or her roommates about disease management techniques (i.e., the knowledge transfer effect). Unique to health, patients may also benefit from healthier roommates if roommates provide informal care or alert hospital staff when the patient may be unaware or unable to do so him- or herself.

In acute care settings, patients can also be *indirectly* affected by their roommates through spillover effects in care delivery, even in the absence of direct (i.e., social norms, knowledge transfer, or informal caregiving) peer effects. For example, better health status of the roommates could contribute positively to a patient's recovery process if patients and their roommates have to compete for limited hospital resources (e.g., clinician time, life-support equipment). In this case, having healthier roommates who utilize fewer hospital resources could free up more of the resources to be available to the patient. Healthier roommates are also less likely to require frequent presence of clinicians in the hospital room, thus resulting in reduced noise levels, fewer interruptions, and improved quality of sleep (Bobrow and Thomas 2000; Brown and Gallant 2006; Duffin 2002; Knutt 2005; Ulrich 2003). These care delivery spillovers are "peer effects" in the sense that an exogenous change in the roommate's clinical condition can have a causal effect on the patient's own clinical condition. However, spillover effects in care delivery are limited to hospitals and similar care settings; therefore, unlike direct peer effects, they aren't indicative of health spillovers as a general phenomenon.

Health spillovers in acute care can be conceptualized as an exogenous shift in the production function of the subject patient's health (see Online Appendix A, http://www.mitpressjournals.org/doi/suppl/10.1162/ajhe_a_00068). An increase in the roommate's health allows the hospital to produce a greater quantity of the subject patient's health for any given quantity of health services provided to the subject patient, thus resulting in an upward shift of the production function of the subject patient's health. Assuming that the patient's health must meet a certain minimum threshold before he or she can be discharged from the hospital, it can be shown that the hospital's optimal response to an upward shift in the subject patient's health production function depends on the extent to which the patient's insurer relies on the fixed-rate method, versus the fee-for-service method, in determining the amount of reimbursement. When the insurer reimburses the hospital based on the fixed-rate method, the reimbursement amount is predetermined based on the patient's diagnosis and reason for admission and does not depend on the actual volume of services provided to the patient. Therefore, under the fixed-rate scenario the hospital's profit-maximization problem is equivalent to cost-minimization, and the hospital's optimal response to an exogenous increase in the production function of the patient's health is to reduce the amount of care provided to the patient. When the patient's insurer uses the fee-for-service reimbursement approach, the hospital is reimbursed based on the volume of services provided to the patient and will optimally provide the amount of services to the patient such that per-patient profit is maximized. Under this scenario, an upward shift in the patient's health production function does not change the hospital's profit-maximizing equilibrium thus leaving the amount of care provided to the patient unchanged; however, the output of the patient's health will be greater as a result of the roommate-induced upward shift in the production function. These predictions are formally derived in Online Appendix A.

In reality, most insurance plans use a combination of fixed-rate and fee-for-service approaches, although private insurance plans are more likely to offer fee-for-service payments while Medicare and Medicaid are almost exclusively fixed-rate payers for hospital expenses. Therefore, while I expect that sharing a room with a healthier patient will result in a combination of a reduction in the quantity of care provided to the patient (e.g., fewer daily contacts with the patient, shorter length of stay, lower costs) and improved health for all patients, the reduction in care quantity effect is expected to dominate for patients with insurance plans that rely more heavily on fixed-rate payment (Medicare, Medicaid), while the increase in health effect might dominate for patients whose insurers rely more heavily on the fee-for-service payment method (private insurance plans).

III. Data

A. DATA SOURCES

The study uses electronic medical records data for adult patients who were hospitalized in semiprivate hospital rooms at a nonprofit tertiary medical center in New Haven, Connecticut. The medical center had 854 adult beds and close to 56,000 annual discharges. The data were extracted from three databases within the study hospital's electronic information system. First, the Record Information Management System (RIMS) included data on utilization (e.g., previous hospitalizations and subsequent readmissions), mortality, patient characteristics (e.g., age, sex, diagnosis, insurance), length of stay, and costs.

Second, the clinical database from the electronic point-of-care system (i.e., Sunrise Clinical Manager (SCM), Allscripts, Chicago, Illinois) provided data on the patient's Rothman Index (RI) score (Rothman et al. 2012; Rothman, Rothman, and Beals 2013; Rothman, Rothman, and Solinger 2013) that I use as a measure of clinical condition. The RI incorporates both objective data on routinely monitored clinical variables including lab results (e.g., protein level in urine, an indicator of renal function, or creatinine blood level, an indicator of kidney function) and vital signs (e.g., blood pressure, heart rhythms). It also incorporates the results of nurse assessments of the patient's overall clinical condition (psychosocial functioning, cognition, mobility, cardiovascular, gastrointestinal, etc.).[1] The RI is recalculated each time a patient's electronic record is updated by a clinician (e.g., a doctor, nurse, laboratory technician, radiologist), with close to 15 daily updates for an average patient. In previous studies, the RI was correlated with a number of health outcomes: when the patient's RI was below 40, the patient had a fivefold increase in the risk of 30-day mortality compared with patients with the RI of 40 or above, and the RI of less than 70 at the time of discharge was associated with close to a threefold increase in the odds of an unplanned 30-day readmission relative to the 70 and above RI category (Rothman, Rothman, and Solinger 2013; Bradley et al. 2013). In my sample, 1 out of every 10 patients whose RI dropped below 40, and 3 out of every 10 people whose RI dropped below 20, died within the next six to eight hours.

1 I describe the RI in more detail in Online Appendix C.

Lastly, the Patient Activity Database (PAD) contained time-stamped records of the patient's location at the time of admission, discharge, and all transfers within the hospital; I use the PAD to match patients with their roommates based on overlaps in the dates and times (to one-quarter of an hour) of in- and out-transfers for each room.

B. SAMPLE

The full electronic data file includes 27,571 adult (i.e., 18 and older) patient discharges that occurred between July 1 and December 31, 2011.[2] From this sample, I eliminate admissions to patient care units that did not have semiprivate (i.e., shared) rooms. Of the 59 patient care units at the study hospital, this eliminates a total of 15,235 admissions to 48 patient care units.[3] The remaining 11 patient care units include three surgical units (i.e., cardiac, gastrointestinal, and orthopedic) and eight general medical units (i.e., two cardiac, one trauma, and five general medical units serving a broad range of patients) and had 12,336 admissions during the study period. At the time that the data were extracted, the Rothman Index (RI)—my clinical condition measure—was computed only for inpatient admissions (not for patients kept at the hospital for clinical observation). Eliminating 1,040 patients on the surgical and general medical units who were at the hospital for clinical observation and 828 admissions with missing clinical data reduces the sample to 10,468 hospital admissions with complete clinical data records on the 11 surgical or general medical units that offered semiprivate rooms. Of these patient admissions, 6,458 patients had at least one roommate for whom I also have complete clinical data. Lastly, to ensure that patients in my sample had sufficient roommate exposure to identify peer effects, I restrict my sample to patients who were hospitalized for at least 24 hours and who spent 50 percent or more of their hospital stay with a roommate. This results in the final sample of 3,793 patients (i.e., 1,950 females and 1,843 males) in 113 semiprivate rooms on the 11 surgical and general medical study units.[4] The descriptive statistics of the full patient sample, the sample restricted to units with semiprivate rooms, and the final patient sample are presented in Table 1.

My roommate sample differs from the general adult patient population at the study hospital for two reasons. First—as I describe earlier—not all patient care units are included in this study, thus resulting in select diagnostic groups being under- or over-represented in my sample relative to their frequency distribution among all adult hospital patients in my data (Online Appendix B). Most notably, obstetric patients (e.g., labor and delivery, pregnancy complications, and delivery-related trauma) represent less than 0.5 percent of all admissions in the final roommate analysis sample, compared with almost

2 The hospital did not grant permission to access data on pediatric patients who were younger than 18 at the time of discharge.

3 This excluded all of the intensive care units, as well as most small specialty units including infectious diseases, burns, neonatal care, cancer and chemotherapy, psychiatric, diagnostic imaging, and dentistry surgery. Obstetric/gynecology units did not have semiprivate rooms and were excluded. Additionally, several of the general medical and surgical units did not have semiprivate rooms and were excluded. However, patients with psychiatric, gynecological, dentistry, and other diagnoses were sometimes admitted to general medical/surgical units instead of the corresponding specialty unit. I retained these observations in my sample.

4 I discuss the robustness of my results to gradually relaxing this restriction in Section V.

TABLE 1. Descriptive characteristics of the sample

Variable	All adult patients, $N = 27{,}571$ Mean [std. dev.]	Patients on units with shared rooms, $N = 10{,}468$ Mean [std. dev.]	Final roommate sample, all, $N = 3{,}793$ Mean [std. dev.]	Final roommate sample, females, $N = 1{,}950$ Mean [std. dev.]	Final roommate sample, males, $N = 1{,}843$ Mean [std. dev.]
Subject patient's characteristics:					
Female	0.57 [0.49]	0.51 [0.50]	0.51 [0.49]	—	—
Age	54.36 [19.81]	60.21 [19.06]	60.53 [19.18[c]]	62.16 [19.89]	58.81 [18.24]
Insurance type: Private	0.34 [0.47]	0.28 [0.45]	0.27 [0.44[c]]	0.25 [0.43]	0.31 [0.46]
Insurance type: Medicaid	0.23 [0.42]	0.21 [0.41]	0.21 [0.41]	0.21 [0.41]	0.22 [0.41]
Insurance type: Medicare	0.38 [0.48]	0.49 [0.50]	0.48 [0.49[c]]	0.52 [0.50]	0.43 [0.50]
Insurance type: None	0.09 [0.29]	0.11 [0.32]	0.11 [0.26]	0.11 [0.26]	0.12 [0.27]
Type of admission: Surgical	0.18 [0.38]	0.18 [0.38]	0.28 [0.44[a]]	0.28 [0.45]	0.28 [0.45]
Type of admission: Medical	0.49 [0.50]	0.62 [0.49]	0.72 [0.44[a]]	0.72 [0.45]	0.72 [0.45]
Prior hospitalization w/in 30 days	0.17 [0.36]	0.14 [0.35]	0.13 [0.33[c]]	0.13 [0.33]	0.14 [0.34]
Rothman Index: Admission	—	74.53 [18.13]	77.51 [14.59[b]]	77.01 [14.54]	78.00 [14.63]
Rothman Index: Discharge	—	77.17 [17.06]	79.31 [13.95[b]]	78.33 [14.15]	80.34 [14.63]

TABLE 1. *Continued*

Variable	All adult patients, $N = 27{,}571$ Mean [std. dev.]	Patients on units with shared rooms, $N = 10{,}468$ Mean [std. dev.]	Final roommate sample, all, $N = 3{,}793$ Mean [std. dev.]	Final roommate sample, females, $N = 1{,}950$ Mean [std. dev.]	Final roommate sample, males, $N = 1{,}843$ Mean [std. dev.]
Rothman Index: Start of spell	—	—	77.71 [13.96]	77.15 [13.94]	78.28 [13.97]
Rothman Index: End of spell	—	—	78.84 [13.72]	77.92 [13.80]	79.81 [13.57]
Length of stay	5.55 [9.56]	5.68 [9.75]	4.12 [4.08[c]]	4.11 [4.00]	4.12 [4.17]
Total cost of hospitalization	17,433.60 [31,855.90]	18,266.24 [30,164.86]	12,206.10 [12,243.20[a]]	11,857.30 [11,646.90]	12,575.20 [12,836.90]
Death during hospitalization	0.02 [0.13]	0.01 [0.10]	0.00 [0.05]	0.00 [0.06]	0.00 [0.05]
Unplanned readmission w/in 7 days	0.06 [0.24]	0.07 [0.25]	0.06 [0.24[c]]	0.06 [0.23]	0.07 [0.26]
Unplanned readmission w/in 30 days	0.16 [0.37]	0.15 [0.36]	0.14 [0.35[b]]	0.14 [0.34]	0.14 [0.35]
Average daily care frequency	—	13.41 [7.52]	14.39 [6.15[c]]	14.13 [5.84]	14.67 [6.46]
Proportion of stay in shared room	—	—	0.82 [0.14]	0.82 [0.14]	0.81 [0.15]
Number of roommate spells	—	—	1.89 [1.27]	1.89 [1.31]	1.89 [1.23]
Length of roommate spell, days	—	—	1.79 [1.16]	1.81 [1.13]	1.77 [1.19]

TABLE 1. *Continued*

Variable	All adult patients, $N = 27,571$ Mean [std. dev.]	Patients on units with shared rooms, $N = 10,468$ Mean [std. dev.]	Final roommate sample, all, $N = 3,793$ Mean [std. dev.]	Final roommate sample, females, $N = 1,950$ Mean [std. dev.]	Final roommate sample, males, $N = 1,843$ Mean [std. dev.]
Roommate's characteristics:					
Female	—	—	0.51 [0.49]	—	—
Age	—	—	61.27 [16.08]	63.00 [16.52]	59.44 [15.41]
Private insurance	—	—	0.26 [0.37]	0.24 [0.36]	0.30 [0.38]
Medicaid	—	—	0.21 [0.34]	0.20 [0.34]	0.22 [0.35]
Medicare	—	—	0.49 [0.42]	0.55 [0.42]	0.45 [0.42]
Surgical admission	—	—	0.12 [0.29]	0.11 [0.28]	0.13 [0.31]
Medical admission	—	—	0.64 [0.45]	0.64 [0.46]	0.66 [0.46]
Prior hospitalization w/30 days	—	—	0.14 [0.29]	0.14 [0.30]	0.14 [0.29]
Rothman Index at start of spell	—	—	74.53 [13.80]	73.92 [13.48]	75.17 [14.09]

Notes: Shown are means and standard deviations for the full sample of all adult patients hospitalized between June 1, 2011, and December 31, 2011 (column 1), the sample of all patients on units with semiprivate rooms included in the study (column 2), and the final roommate sample (all patients, column 3; female patients, column 4; and male patients, column 5). The lowercase superscript letters denote significant differences between column 3 and column 2. $^{a}p < 0.01$, $^{b}p < 0.05$, $^{c}p < 0.10$.

10 percent of all admissions in the full adult hospital population. Additionally, cancers and psychiatric diagnoses represent only 2 percent and 3.5 percent my final analysis sample, respectively, compared with their shares in the full adult patient population being close to 8 percent each. Among over-represented diagnoses in my sample are cardiovascular diseases (30 percent of the final sample relative to 23 percent in the full population), trauma and poisoning (11 percent in the sample versus 6 percent in the hospital population), and back problems (close to 4.5 percent in the final sample versus 1.5 percent in the hospital population). As a result, the patients in my sample tend to be slightly older and more likely to be on Medicare, fewer of them are females, and more are admitted for medical reasons and not for surgery (Table 1, columns 2 and 3 versus column 1).

Second—in addition to selection resulting from not including units without semiprivate rooms—another source of selection in my final sample stems from excluding (by design) patients hospitalized alone in private rooms who tend to be in poorer health than patients in shared rooms. Compared with all patients on the selected units (Table 1, column 2), patients in my final roommate sample (Table 1, column 3) have higher RI scores at the time of admission and are less likely to have had another hospitalization during the prior 30 days; they also have better outcomes (e.g., higher RI scores at discharge, lower likelihood of death, shorter lengths of stay, lower costs of hospitalization, and lower likelihood of readmission).

C. PATIENT OUTCOMES

I use five main outcome measures: daily frequency of care (average number of times new activity was recorded in the electronic point-of-care system), length of stay, cost of hospitalization, clinical condition at discharge, and readmission within 7 and 30 days post-discharge. I describe the variables in detail in Online Appendix C.

Clinical conditions at discharge and readmission are used as health measures, and my a priori expectation is that having a healthier roommate will result in improved clinical condition at discharge and reduced likelihood of readmission. I use daily care frequency, length of stay, and cost to examine the roommate effect on the quantity of care. I expect that sharing a room with a healthier roommate will result in lower care frequency, shorter hospitalization, and lower costs.

D. ROOMMATE HEALTH MEASURE

I use the RI score of the patient's roommate as the main roommate exposure variable. I use electronic time stamps (date, hour, minute) from the patient activity database to determine the exact time when a new patient-roommate match is formed (using the later of the patient's or the roommate's respective in-transfer time stamps) and dissolved (using the earlier of the patient's or the roommate's out-transfer time stamps). I define the time interval between the formation and the dissolution of the patient-roommate match as a roommate exposure spell.

I then extract the last RI score of the patient's roommate that was recorded prior to the start of the roommate spell. Frequent RI updates allow me to establish pre-exposure measurements with the average lead of approximately 35 minutes prior to the beginning of the roommate spell; the maximum lead is less than two hours. I examine the effect of a unit change in the roommate's average RI as a continuous variable, and I also create a three-level categorical roommate RI variable for having a roommate in the "low," "medium," or "high" RI category corresponding to the lowest (−5.13–71.9, mean 58.77), middle (72–84.5, mean 78.46), and highest (84.5–99.6, mean 90.70) one-third of the RI distribution.

E. ROOM ASSIGNMENTS

The assignment of patients to rooms was conducted at the hospital through computerized bed assignment software. A nurse manager responsible for new beds logged into the system and was presented with a screen showing a diagram of the real-time status of all of the beds on the designated unit, including their location on the floor, type of room, and the proximity to the nurse station (Online Appendix D). Color coding was used to show which beds were available or would soon become available, as well as the gender of the patient occupying the other bed and whether the patient required isolation (for semiprivate rooms). The average occupancy rate during the study period was over 90 percent. Since the hospital is located in an urban area with a large proportion of low-income uninsured patients, many hospital encounters were unplanned admissions through the emergency room. The designated unit was determined depending on the patient's diagnosis and care needs at the time of admission or transfer from the intensive care unit (e.g., a heart surgery patient's order would call for the patient to be transferred to the cardiac surgery unit). However, when a bed on the designated unit was not immediately available, the patient was assigned a "virtual bed" (i.e., the grey squares underneath the chart) and placed in the hallways until a regular bed became available. Several of the hospital administrators I spoke with stated that bed assignments for patients with noncontagious diseases and for patients who were not immune-deficient themselves (those not requiring isolation) were largely based on space availability; however, patients were almost always assigned to rooms with patients of the same gender, and sicker patients requiring more care were placed in rooms located closer to the nurse station located on the hospital unit whenever possible.

IV. Analysis

A. THE MODEL

I estimate a variant of the standard linear-in-means peer influence model (Manski 1993) where the outcomes of the subject patient are regressed on the leave-out peer group average health of the patient's roommates. The conventional linear-in-means equation is the following:

$$Y_{ik} = \alpha_0 + \beta_0 \bar{Y}_{(i)k} + \gamma_0' X_i + u_{ik}. \qquad (1).$$

In equation 1, Y_{ik} is the outcome of patient i in room k (i.e., care frequency, length of stay, cost, discharge RI, and readmission) and $\bar{Y}_{(i)k}$ is the leave-out peer mean computed as the average of the roommates' RIs excluding patient i. The vector of patient characteristics, X_i, includes the subject patient's own RI score at admission, gender (equal to one if the patient is female), age (coded as a categorical variable: 18–44, the reference category / 45–54 / 55–64 / 65–74 / and 75 or older), type of insurance (coded as a categorical variable: private/Medicare/Medicaid/uninsured), prior hospitalization history (an index variable for having a record of an earlier hospitalization within 30 days prior to the focus admission), and type of admission (an index variable for being admitted for surgery, versus an admission for general medical care). Because patients who spent a larger proportion of their hospital stay in a private room tend to be sicker, I also control for the proportion of the hospital stay spent in a shared room (i.e., a continuous variable from 0.5 to 1). The social spillover effect, or peer influence, is captured by coefficient β_0, which can be used to calculate the social multiplier or the total effect of a unit change in the peer group average clinical condition, $1/(1-\beta_0)$ (Glaeser, Sacerdote, and Scheinkman 2003).

I examine the effect of the roommate's health on the subject patient's outcomes of hospitalization for the full roommate sample first, and then for male and female patients separately. The gender differences are of interest for several reasons. First, acute care is characterized by a natural gender division of patients because of same-sex patient room assignments and gender-specific clinical care needs. Second, in conducting preliminary analyses for this study, I rejected the hypothesis that the peer effect coefficient estimates were the same for male and female patients. Lastly, gender differences in peer influence are apparent in a number of earlier empirical peer effects studies (Eagly and Chrvala 1986; Mears, Ploeger, and Warr 1998; Tamres, Janicki, and Helgeson 2002; Trogdon, Nonnemaker, and Pais 2008; Yakusheva, Kapinos, and Weiss 2011; Yakusheva, Kapinos, and Eisenberg 2014).

There are two issues that I need to consider in estimation. First, using the standard linear-in-means approach (equation 1) for these data has limitations because the number of roommates varies from patient to patient. The average peer group size is 1.89 roommates (not including the patient), with about one-half of the sample matched to one roommate and close to 10 percent of the sample matched to three or more (up to 30) roommates over the course of hospitalization. Patients exposed to few roommates are more likely to be in the tails of the average roommate RI distribution (i.e., to be exposed to very sick or very healthy roommates), while patients with more roommates are more likely closer to the mean of roommate RI. Because the size of the peer group is endogenous (i.e., healthier patients tend to have shorter hospitalizations, limiting the number of roommates), this could bias my estimates. I include a set of fixed effects for the number of roommate spells in all models; however, the results are robust to not including them.[5, 6]

5 Because the number of roommate spells is positively correlated with the length of stay, it is possible that—by controlling for the number of roommate spells—I somewhat underestimate the roommate effect on the length of stay and costs. However, my estimates are robust to controlling or not controlling for the number of roommate spells.

6 I also estimated the model for subsamples of patients with the same number of roommate spells (grouping them, for sample size considerations, as 1/ 2/ 3 +). The results within each of the subgroups are consistent with the overall effects for both males and females; in the pooled regression with interaction effects for the number of roommate spells, the results are also robust and none of the interaction terms are statistically significant. The results are available upon request.

The second issue is that patients are exposed to roommates for different lengths of time. For example, consider patient *P* who shares the room with roommate *R1* on days one and two and with roommate *R2* on days three, four, and five of his five-day hospital stay. Using the average of the clinical condition scores of the roommates as the main peer measure (as it is commonly done in the peer effects literature) would ignore the fact that the second exposure spell was longer than the first. To account for this feature of the data, I estimate a dyadic peer influence model where a patient is linked to each of the roommates individually. For the patient in the example, this approach results in two dyadic observations (*P* and *R1*), (*P* and *R2*). I weight the data by the standardized proportion of stay spent with each roommate. In the above example, the first observation (*P* and *R1*) would receive the weight 2/5, and the second observation (*P* and *R2*) would receive the weight 3/5 (Yakusheva and Fletcher 2014). When expanded into dyads of roommate spells, my sample of 3,793 patients has 7,166 patient-roommate observations (i.e., 3,688 female observations and 3,478 male observations).

The equation of my main estimation model is the following:

$$Y_{ijk} = \alpha_1 + \beta_1 Y_{jk} + \gamma_1' X_i + d_i + l_k + t_{ijk} + u_{ijk} \quad (2).$$

As before, Y_{ijk} is the outcome patient i matched with roommate j in room k (i.e., care intensity, length of stay, cost, discharge RI, and 7-day and 30-day readmission); Y_{jk} is roommate j's RI immediately prior to the beginning of patient i's exposure to roommate j in room k; X_i is the vector of patient i's controls (i.e., own RI at admission, gender, age, insurance, type of hospitalization, proportion of stay in shared room, and fixed effects for the number of roommates). As I discuss in detail in the next section, I also include fixed effects for the patient's diagnosis, d_i, the room where the i-j spell occurred, l_k, and the calendar week, day, and time when the i-j spell began, t_{ijk}. The observations are weighted as described above and the standard errors are clustered at the patient level. I estimate linear models for all outcomes, including linear probability models for the two readmission outcomes.

The social spillover effect is captured by coefficient β_1, which reflects the statistical effect of a unit change in the roommates' average lagged-to-baseline RI on the focus patient's frequency of care, length of stay, costs, clinical condition at discharge, and likelihood of readmission. If positive patient-to-patient health spillovers are present in my sample, I expect my findings to include some combination of negative coefficient estimates in the models for the effect of the roommate health on the subject patient's care frequency, length of stay, costs, and readmission, and a positive peer effect estimate in the model for the subject patient's clinical condition at discharge. Because of the inverse frequency weighting, the peer effect estimate is equivalent to the conventional social spillover estimate β_0 in model 1.

B. SOURCES OF BIAS AND INTERNAL VALIDITY

The challenge in empirical investigation of causal peer effects is that regressing a patient's outcome on the outcomes of the roommates is likely to produce a positive estimate of the peer effect coefficient even in the absence of true underlying peer influence, thus leading to an incorrect attribution of observed positive correlations in group-level outcomes to peer effects. As first pointed out by Manski (1993) and further developed by others (Angrist 2014; Boozer and Cacciola 2001; Hanushek et al. 2003; Moffitt 2001), nonexperimental peer effect studies face significant threats to internal validity because of potential unobserved peer selection and spurious correlation in outcomes. Unobserved selection bias may arise in my study because—although hospital patients do not choose their roommates—room assignment process matches patients based on their gender, diagnoses, and care needs. Clinical condition scores of patients hospitalized in the same room could also be spuriously correlated because of exposure to common shocks (staffing differences during holidays or weekends, outbreaks of contagious diseases, etc.). Not accounting for unobserved selection and common shocks could result in an upward bias in the estimate of the peer effect coefficient in model 2.

To address unobserved selection, this study relies on a room assignment process that is plausibly random after conditioning on a set of observed patient characteristics that are used by the nurse manager in determining the patient's room assignment (gender, diagnosis, room type that meets the patient's care needs). To account for these factors, I create a vector of diagnosis fixed effects and a vector of room fixed effects. The diagnosis fixed effects categorize patients based on the diagnostic groups (1–207) that are defined by the Agency for Healthcare Research and Quality (AHRQ) Clinical Classification Software (CCS); and the room fixed effects control for the room (1–113) in which the roommate exposure spell occurred. In order to reduce unobserved selection from potential systematic variance in patient characteristics that may occur over time, I also include three sets of time indicators—the week of the year (1–26 during the six-month study period), the day of the week (1–7), and the hour of day (1–24)—for the start time of the roommate exposure spell. Conditional on diagnosis, room, and week/day/hour group means, the remaining variation in peer characteristics is assumed to be exogenous.

I conduct several randomization checks to examine the plausibility of the conditionally random roommate assignment. Table 2 presents a set of balancing tests that I obtain using univariate regressions with and without controls for the diagnosis, room, and week/day/hour fixed effects (I split the results by gender in Online Appendix E). The estimates are corrected for the negative small-cell bias (i.e., two to three patients per room) using the Guryan, Kroft, and Notowidigdo (2009) correction. Consistent with the positive roommate matching process described to me by hospital administrators, the results show strong positive unconditional correlations in almost all of the clinical and demographic measures between the subject patients and the patients' roommates (column 1). For example, having roommates who are admitted for surgery increases the patient's own chances of having been admitted for surgery by nearly 45 percent;[7] for every point increase in the roommate's average RI, the patient's own admission RI increases by 0.104 points, and for each additional year of the roommate's age, the patient is 0.18 of a year (65 days) older.

7 Because there are units that designated specifically for surgical patients, one might expect this correlation to be closer to unity; however, patients admitted for surgery often spend only part of their stay on a designated surgical unit and are transferred to general medical units once their needs for care become more routine.

TABLE 2. Balancing tests

	Fixed effects: None	Fixed effects: Diagnosis only	Fixed effects: Room only	Fixed effects: Admission week/day/time only	Fixed effects: All
Female	0.985[a] (0.002)	0.983[a] (0.003)	0.983[a] (0.003)	0.984[a] (0.003)	0.982[a] (0.003)
Admission RI	0.104[a] (0.012)	0.078[a] (0.011)	0.025[b] (0.012)	0.101[a] (0.012)	0.020[c] (0.011)
Age	0.178[a] (0.015)	0.132[a] (0.014)	0.050[a] (0.015)	0.168[a] (0.015)	0.031[b] (0.014)
Insurance: Private	0.063[a] (0.015)	0.040[a] (0.015)	0.002 (0.015)	0.050[a] (0.015)	−0.008 (0.015)
Insurance: Medicaid	0.053[a] (0.015)	0.032[b] (0.015)	−0.004 (0.015)	0.043[a] (0.015)	−0.016 (0.015)
Insurance: Medicare	0.119[a] (0.015)	0.091[a] (0.014)	0.025[c] (0.015)	0.108[a] (0.015)	0.018 (0.014)
Type of admission: Surgical	0.447[a] (0.021)	0.230[a] (0.019)	0.018 (0.021)	0.395[a] (0.023)	0.008 (0.016)

Notes: $N = 3{,}793$. Shown are the individual correlations between the characteristics of the patient and the average of his/her roommates. In column 1, I show unconditional correlations; columns 2–4 show univariate regressions with controls for diagnosis, room, and week/day/time fixed effects, respectively, and column 5 shows correlations after all of the fixed effects are included. The estimates are corrected for small-cell bias at the room level using the Guryan, Kroft, and Notowidigdo (2009) correction. Standard errors are in parentheses. $^{a}p < 0.01$, $^{b}p < 0.05$, $^{c}p < 0.10$.

Not surprisingly, there is a close to one-to-one correlation between the patient's and the roommate's gender.[8]

Adding the fixed effects makes the correlations smaller, with the room fixed effects alone reducing the correlation coefficients between the patient's and his/her roommates' characteristics by close to 90 percent of their unconditional magnitudes. Once the full set of fixed effects is included, only two of the conditional baseline correlations (RI and age) in the pooled sample remain statistically significant but small (the patient's admission RI increases by 0.02 points for a 1 point increase in the roommate's RI, and the patient is 0.03 years (10 days) older when the roommate is 1 year older). These residual correlations might result from higher-order nonlinear gender-specific selection effects. (Note that controlling for the full set of fixed effects does not affect the magnitude of the gender correlation.) None of the conditional correlations are significant when the sample is split by gender (Online Appendix E), with the exception of a negative correlation in the type

8 Of the 7,166 patient-roommate matches, there are 41 female subject / male roommate matches and 36 male subject / female roommate matches.

of admission for females (female patients are 3.96 percent less likely to be admitted for surgery when their roommates are admitted for surgery) and in the indicator for being privately insured among males (male patients are 5.3 percent less likely to be privately insured if their roommates are privately insured). There also appear to be small negative correlations across the other insurance categories in the magnitude of 1–3 percentage points for male patients. After discussing these results with one of the nurse managers who was routinely responsible for bed assignments, I was unable to determine how this negative correlation could have resulted from a purposeful roommate matching process. It should also be noted that the p-values reported in the table are not corrected for multiple tests; a Bonferroni-type correction would increase the standard errors.

In an attempt to further explore the plausibility of conditionally random roommate assignments, I estimate linear regressions of the roommate's lagged-to-baseline RI on the patient's preassignment characteristics (i.e., patient's RI at admission, age, insurance type, prior hospitalization history, and whether the patient was admitted for a surgical procedure), using the patient-roommate spell as the unit of analysis and conditional on the full set of fixed effects. Statistical significance of the patient's characteristic in predicting the roommate's start-of-spell RI, conditional on the full set of fixed effects, could indicate a systematic assignment bias and would, therefore, be inconsistent with conditionally random roommate assignment. The results of the analysis are presented in Table 3 (split by gender in Online Appendix F) and show that overall, the patient's preassignment characteristics individually did not explain variance in the roommate's RI at the time of assignment (column 1) and did not predict having a roommate in the low or high RI category. Although a few of the coefficients are individually significant, there does not appear to be a systematic pattern of significance across genders and models; furthermore, the signs of the significant coefficients are not consistent with unobserved selection (positive or negative). I also tested the joint significance of the individual preassignment patient characteristics in predicting the roommate's RI at the time of assignment, and the coefficients were jointly nonsignificant in all of the models I tested. Overall, these results are consistent with conditionally balanced assignment of roommates in my study.

An additional consideration pertaining to unobserved selection involves roommate assignments that are made later in the patient's hospitalization. These later assignments could potentially incorporate unobservable information learned by the nurses while interacting with the patient during hospitalization, thus essentially becoming nonrandom. To examine the extent of this potential issue, I estimate intent-to-treat models by limiting the sample to only the first roommate for each patient and then reestimating the models after limiting the sample even further by eliminating all patients who spent two or more hours at the hospital before their first roommate assignment. I find that the intent-to-treat estimates are similar in magnitude to the estimates obtained using my full sample of all roommate spells, as I discuss in more detail in Section V. The plausibly random assignment of hospital roommates (conditional on observables) helps ensure internal validity of my estimates of peer influence.

My approach to the second source of bias, due to spurious correlations in outcomes that may result from common shocks, entails using a lagged baseline roommate RI measure that was recorded prior to the time when the patient became exposed to the

TABLE 3. Randomness regressions of the roommate's start of spell RI on preassignment characteristics of the patient

	Roommate's RI	Roommate's RI category is "low"	Roommate's RI category is "high"
Patient characteristics:			
Female	−0.481	−0.015	0.011
	(0.492)	(0.015)	(0.015)
RI, start of spell	0.016	0.000	0.000
	(0.018)	(0.001)	(0.001)
Age: 45–54	−0.669	0.033	−0.027
	(0.733)	(0.023)	(0.023)
Age: 55–64	−0.215	0.008	0.006
	(0.727)	(0.022)	(0.023)
Age: 65–74	0.654	−0.020	−0.007
	(0.934)	(0.028)	(0.029)
Age: 75 +	−0.747	0.008	−0.026
	(0.950)	(0.030)	(0.029)
Insurance: Medicare	−1.119	0.040[c]	−0.011
	(0.721)	(0.022)	(0.022)
Insurance: Medicaid	−1.037	0.026	−0.020
	(0.646)	(0.020)	(0.020)
Insurance: None	−0.671	0.005	−0.019
	(0.687)	(0.021)	(0.020)
Prior hospitalization	0.924	−0.011	0.010
	(0.660)	(0.020)	(0.020)
Observations	7,166	7,166	7,166
p-value: Joint significance of all patient characteristics	0.124	0.455	0.792

Notes: $N = 3{,}793$. Shown are the results of regressing the roommate's RI at the beginning of the spell (left-hand-side variable) on the patient's predetermined characteristics (right-hand-side variables) with controls for diagnosis, room, and calendar fixed effects. Lack of individual and joint statistical significance is consistent with random assignment. Observations are weighted using roommate exposure weights; standard errors are adjusted for clustering at the patient level. Robust standard errors are in parentheses. [a]$p < 0.01$, [b]$p < 0.05$, [c]$p < 0.10$.

roommates. Assuming orthogonal-to-baseline peer group variation (conditional on observables at baseline), this estimation approach minimizes confounding from any common shocks during the roommate spell that may cause a simultaneous spurious change in the outcomes of the patient and his or her roommates.

Lastly, a recent critique to empirical investigation of peer effects by Angrist (2014) discusses a threat to internal validity of peer effect estimates that may exist even when random peer group formation is present. In particular, Angrist (2014) shows that the estimate of the peer effect coefficient in a linear-in-means model with covariates, similar to my main estimation model, is approximately equal to the ratio of the two-stage least squares estimates of the effect of individual covariates (X) on outcomes (Y), instrumented with peer group dummies, to the simple ordinary least squares estimates of the effect of the individual covariates on the outcomes. This mechanical relationship leads to potential misattribution of any divergence between the 2SLS and OLS estimates (for any number of reasons including measurement error, omitted variable bias, nonlinearity, or local average treatment effects) to peer effects, thus posing threat to internal validity even in randomized peer group settings. This problem is exacerbated in studies with few large peer groups (e.g., schools, grades, or dormitories), which can produce a strong first stage despite randomization (Angrist 2014).

My study has two features that make divergence between 2SLS and OLS estimates, for reasons other than peer effects, unlikely. First, an average patient in my study is exposed to fewer than two roommates (half of the patients have only one roommate and 90 percent of the patients have three or fewer roommates); therefore, my group sizes are small. Second, because patients are admitted and discharged at different times, and because my pool of roommates is not restricted by the requirement that at least half of the hospital stay is spent in a semiprivate room in the same way the patient sample is, the number of different peer groups is nearly equal to the number of subject patients (there are only 152 clusters where all patients are included in the analysis as both the subject patient and the peer, and where none of the patients have any other peers outside of the cluster). Therefore, my study has a multiple small peer group structure, leading to a many weak IV scenario that minimizes bias from divergent 2SLS and OLS estimates (Angrist 2014).

Overall, the plausibly random roommate assignment, lagged peer variable specification, and a many weak IV peer group structure are all features of a robust research design (Angrist 2014) that is less likely to generate biased estimates than many earlier peer effect studies.

V. Results

A. PEER INFLUENCE

I present the results of our main regression model 2 of roommate effects on the quantity of care received by the patient during hospitalization (i.e., care intensity, length of stay, and cost) and the patient's health at discharge (i.e., RI, readmission) in Table 4 (and, split by gender, in Online Appendix Table G1). For each of the outcomes, I list the coefficient of the continuous roommate RI variable and the coefficients of the "medium RI" and "high RI" dummies ("low RI" is the omitted category). The top panel shows the estimates for the pooled sample, the middle panel is for female patients, and the bottom panel is for male patients.

The results are consistent with positive health spillovers: for every 1 point increase in the roommate's RI, the subject patient's number of daily RI updates decreases by 0.0802 ($p < 0.001$), the length of stay and costs decrease by 0.18 percent ($p < 0.001$) and 0.13 percent ($p = 0.019$), and the clinical condition score at discharge improves by 0.0144

TABLE 4. Roommate effects on hospitalization outcomes

Variables	Daily quantity of care (Mean = 14.39)		Log of length of stay (Mean = 1.14)		Log of total cost (Mean = 9.08)		Discharge RI (Mean = 79.31)		Readmission, 7 days (Mean = 0.06)		Readmission, 30 days (Mean = 0.14)	
Roommate RI:												
RI	-0.080^{a}		-0.002^{a}		-0.001^{b}		0.014^{c}		0.000		0.001	
	(0.006)		(0.001)		(0.001)		(0.008)		(0.000)		(0.000)	
Categorical roommate RI:												
Medium RI		-1.786^{a}		−0.026		−0.018		0.875^{a}		−0.007		−0.005
		(0.202)		(0.019)		(0.023)		(0.304)		(0.009)		(0.012)
High RI		-2.725^{a}		-0.082^{a}		-0.069^{a}		0.545^{c}		−0.005		0.004
		(0.213)		(0.019)		(0.022)		(0.298)		(0.009)		(0.013)
Observations	7,166	7,166	7,166	7,166	7,166	7,166	7,166	7,166	7,166	7,166	7,166	7,166
R^2	0.301	0.294	0.544	0.545	0.496	0.497	0.677	0.677	0.129	0.129	0.153	0.153

Notes: $N = 3{,}793$. All models control for the patient's characteristics and primary diagnosis fixed effects, room fixed effects, and week/day/hour fixed effects. Observations are weighted with standardized roommate exposure weights; standard errors are adjusted for clustering at the patient level. Robust standard errors are in parentheses. $^{a}p < 0.01$, $^{b}p < 0.05$, $^{c}p < 0.10$.

($p = 0.071$) points. Compared with patients whose roommates are in the low RI category (RI < 71.9), patients with roommates in the high RI (>84.4) category have 2.725 lower daily frequency of RI updates, 8.2 percent shorter lengths of stay, and 6.87 percent lower costs. All of the effects are statistically significant at $p < 0.001$ level. The reduction in the quantity of care is similar for male and female patients (Online Appendix G). The magnitudes of the coefficients of the roommate's RI variable are about five times smaller than the magnitudes of the patient's own admission RI variable (Online Appendix H).

The reduction in the daily frequency of care (i.e., 2.725 fewer RI updates for patients with a roommate in the highest RI category, compared with patients with a roommate in the lowest RI category) represents an 18.9 percent reduction relative to the sample average of 14.39 daily RI updates. Accompanied by an 8.2 percent (equivalent to 8.11 hours per patient) reduction in the length of stay, these estimates suggest that the total care quantity provided to the patient over the course of hospitalization is approximately 27 percent lower when the patient's roommate is in the highest clinical condition category versus the lowest clinical condition category. This represents a significant reduction in the overall quantity of care that is equivalent to half of the difference in the total quantity (from admission to discharge) of all RI updates that I observe between patients younger than 45 and patients aged 55 and older in my sample. It is also similar to the difference in the total quantity of RI updates between patients who are admitted for a planned surgical procedure and patients admitted for an acute medical problem requiring emergency surgery.

The 6.87 percent reduction in the cost of hospitalization (equivalent to $838.55 per patient) is similar in magnitude to the 8.2 percent reduction in the length of stay, but smaller than the estimated reduction in the total quantity of care. This is expected, given that the cost accounting software uses a predetermined cost commutation algorithm and does not log the actual labor hours or products and services used in the treatment of a patient. The cost estimates are derived from patient population resource utilization averages (by diagnosis, length of stay, etc.) and are not considered to be very accurate in measuring the actual cost of care delivery at the individual patient level (Azoulay et al. 2007).

The roommate effects on health, as measured by clinical condition at discharge and likelihood of being readmitted within the first 7 and 30 days after leaving the hospital, are small in magnitude and positive for female, but not for male, patients. Female patients leave the hospital in a slightly better clinical condition when they share a room with a healthier patient, with an approximately 1 RI point difference in clinical condition at discharge between the low roommate RI category and the two higher roommate RI categories. I also observe a 2.18 percent reduction in the probability of readmission within 7 days of discharge for females who had a relatively healthy roommate, although this readmission effect is imprecisely estimated and dissipates by 30 days post-discharge.

I do not observe significant positive or negative effect on the clinical condition score at discharge for male patients. The peer effect coefficients are small in magnitude and nonsignificant in both the linear and the semiparametric model specifications. In fact, males who shared a room with the healthiest roommates are 3.1 percent more likely to have an unplanned hospitalization during the first month after going home, compared with patients with the sickest roommates. This could mean that the lower quantity of care and quicker discharge might have a slight negative effect on the patient's health, increasing the risk of readmission, even though the patient's clinical condition score at the time of discharge does not reflect it. However, the potentially negative roommate effect on health for male patients is still small in comparison with the reduction in the amount of care, length of stay, and costs. Given the nearly 8 percent reduction in the length of stay and 6.77 percent lower cost of hospitalization for male patients, having a healthy roommate is expected to lower the overall cost of care, even accounting for the 3.1 percent chance of additional cost from a subsequent unplanned hospitalization.

Overall, these effects are consistent with theory that having a healthy roommate shifts the patient's health production function upward, thus leading the provider (hospital) to optimally reduce the quantity of care provided to the patient while still discharging the patient in the same or better clinical condition. In also appears that, overall, the reduction in care quantity is much stronger than the increase in health, suggesting that the hospital is responding to predominantly fixed-rate reimbursement incentives. Indeed, nearly 50 percent of my sample are insured by Medicare under a fixed-rate payment system, and close to 20 percent are Medicaid patients (with only a little over a quarter of the sample being privately insured and potentially eligible for fee-for-service reimbursement).

In Table 5, I examine whether the roommate effect estimates vary by the type of insurance (I split the results by gender in Online Appendix I). I estimate a version of model 2 with interaction terms between the roommate's clinical condition and the subject patient's insurance type (private, Medicare, Medicaid, self-pay/uninsured). I expect that Medicare and Medicaid patients will primarily be affected along the reduced quantity of care margin, while privately insured patients will be primarily affected along the increased health margin. I find that the effect on the daily care frequency is similar across the insurance types. However, the reductions in the length of stay and cost do, in fact, appear to be stronger for patients insured by Medicare than for patients who are privately insured. Uninsured patients, who represent only 11 percent of the sample, are impacted by roommates in a way that is similar to privately insured patients; however, the effects are imprecisely estimated. The differences are particularly pronounced among male patients (Online Appendix I) whose length of stay and cost are affected little if they are privately insured, with much larger significant negative peer effects if they are insured by Medicare or Medicaid (with effect sizes about three time as large as they are in the main noninteracted model). I do not, however, observe the expected positive effect on discharge RI or probability of readmission for privately insured patients; in fact, it is Medicaid patients whose clinical condition appears to benefit the most from having a healthy roommate for both males and females. As a group, Medicaid patients are more likely to have complex care needs and lack social support, which could make them more susceptible to roommate effects.

Lastly, I examine difference in roommate influence by the subject patient's own clinical condition at admission by estimating model 2 with interaction effects between the roommate's RI and the subject patient's own clinical condition categories at admission ("low"/"medium"/"high"). Significant differences between the magnitudes of the peer effect across the patient clinical condition can point to the types of patient for whom the benefit from a purposeful roommate assignment approach might be the largest. The estimates are presented in Table 6 (by gender in Online Appendix J), and they show that

TABLE 5. Roommate effects on hospitalization outcomes, by patient's type of insurance

Patient characteristic	Daily quantity of care (Mean = 14.39)	Log of length of stay (Mean = 1.14)	Log of total cost (Mean = 9.08)	Discharge RI (Mean = 79.31)	Readmission, 7 days (Mean = 0.06)	Readmission, 30 days (Mean = 0.14)
Insurance type						
d(y)/d(R's RI) |	−0.083[a]	−0.001	−0.001	0.002	0.000	0.000
Insurance = "Private"	(0.013)	(0.001)	(0.001)	(0.019)	(0.001)	(0.001)
d(y)/d(R's RI) |	−0.081[a]	−0.003[a]	−0.002[c]	0.018	0.000	0.000
Insurance = "Medicare"	(0.009)	(0.001)	(0.001)	(0.018)	(0.000)	(0.001)
p-value of difference	[0.887]	[0.316]	[0.530]	[0.515]	[0.944]	[0.615]
d(y)/d(R's RI) |	−0.072[a]	−0.001	−0.003[b]	0.065[a]	0.000	−0.001
Insurance = "Medicaid"	(0.014)	(0.001)	(0.001)	(0.021)	(0.001)	(0.001)
p-value of difference	[0.527]	[0.931]	[0.320]	[0.020]	[0.559]	[0.729]
d(y)/d(R's RI) |	−0.088[a]	−0.001	0.000	0.020	0.001	0.001
Insurance = "None"	(0.014)	(0.001)	(0.001)	(0.033)	(0.001)	(0.001)
p-value of difference	[0.791]	[0.918]	[0.709]	[0.631]	[0.442]	[0.402]

Notes: $N = 3{,}793$. Shown are the predicted peer effect sizes from a variant of model 2 with an interaction term of the roommate's RI (R's RI) with the subject patient's insurance type. All models control for the other characteristics of the subject patient and primary diagnosis fixed effects, room fixed effects, and week/day/hour fixed effects. Observations are weighted with standardized roommate exposure weights; standard errors are adjusted for clustering at the patient level. Robust standard errors of the coefficients are in parentheses. The p-values of the test of equivalency of the margins with the private insurance category are in brackets. [a]$p < 0.01$, [b]$p < 0.05$, [c]$p < 0.10$.

TABLE 6. Roommate effects on hospitalization outcomes, by patient's clinical condition

Patient characteristic	Daily quantity of care (Mean = 14.39)	Log of length of stay (Mean = 1.14)	Log of total cost (Mean = 9.08)	Discharge RI (Mean = 79.31)	Readmission, 7 days (Mean = 0.06)	Readmission, 30 days (Mean = 0.14)
Roommate's RI:						
d(y)/d(R's RI) |	−0.080[a]	−0.003[a]	−0.001	0.045[b]	0.000	0.001
Own RI = "Low"	(0.010)	(0.001)	(0.001)	(0.020)	(0.000)	(0.001)
d(y)/d(R's RI) |	−0.067[a]	−0.001	−0.002[c]	−0.005	0.000	0.000
Own RI = "Med"	(0.010)	(0.001)	(0.001)	(0.015)	(0.000)	(0.001)
p-value of difference	[0.329]	[0.073]	[0.899]	[0.039]	[0.536]	[0.642]
d(y)/d(R's RI) |	−0.092[a]	−0.002[b]	−0.001	0.008	−0.001	−0.001
Own RI = "High"	(0.010)	(0.001)	(0.001)	(0.011)	(0.000)	(0.001)
p-value of difference	[0.403]	[0.580]	[0.979]	[0.092]	[0.087]	[0.057]

Notes: $N = 3{,}793$. Shown are the predicted peer effect sizes from a variant of model 2 with an interaction term of the roommate's RI (R's RI) with the subject patient's own RI category at admission. All models control for the other characteristics of the subject patient and primary diagnosis fixed effects, room fixed effects, and week/day/hour fixed effects. Observations are weighted with standardized roommate exposure weights; standard errors are adjusted for clustering at the patient level. Robust standard errors of the coefficients are in parentheses. The p-values of the test of equivalency of the margins with the "low" own RI category are in brackets. [a]$p < 0.01$, [b]$p < 0.05$, [c]$p < 0.10$.

patients in the lowest RI category derive the most health benefit from having a healthier roommate in terms of reduced length of stay, lower costs, and improved discharge RI. The benefits of having healthy roommates are particularly strong for the sickest female patients—they gain 0.112 points in their clinical condition score at discharge, while their daily care frequency is 0.877 lower, and hospitalizations are 0.26 percent shorter and 0.26 percent less costly, for every 1 point increase in the roommate's RI. However, males in the lowest RI category—although they, too, require less care during hospitalization and go home earlier—appear to be affected negatively by healthy roommates as their clinical condition at discharge is slightly worse and they are more likely to have an unplanned readmission after being discharged (Online Appendix J).

To explore asymmetric peer effects further, I examine predicted outcomes using a two-way stratification ([patient's own admission RI category] $\times$ [roommate's admission RI category]) for the pooled sample of male and female patients. I convert the length of stay and cost from their log-transformed values to hours and dollars for ease of interpretation. (The results of the two-way stratification of the predicted care frequency, length of stay, cost, and discharge RI by the patient's own and the roommate's clinical condition category are presented in Online Appendix K. Because readmissions are relatively infrequent, the two-way stratification of the predicted readmission rate is noisy and I do not present these results.)

The stratified estimates suggest that the patient's outcomes might be most sensitive to the roommate's clinical condition when the patient's own health is the poorest. The sickest patients have 19.3 hours shorter hospitalizations, $1,700 lower costs, and are discharged with 2-point-higher RIs when their roommates are in the healthiest category, relative to when their roommates are in the sickest category; for the healthiest patients, the gains of having healthy roommates are 7-hour-shorter hospitalizations, $1,130 lower costs, and close to one-third of a point higher RIs, on average. The differences between the effect of the roommates on sick and healthy patients' length of stay and cost are significant at the 5 percent and the 10 percent level, respectively; however, the differences in the roommate effect on daily quantity of care and discharge RI are not statistically significant and unlikely to be clinically meaningful.

The findings of the asymmetric effects of the roommates' health status on the outcomes of hospitalization of sick versus healthy patients suggest that a negative roommate matching mechanism—whereby high-acuity patients are assigned to rooms with low-acuity patients—may have a potential for improving overall health outcomes, because the gains to the sick roommate from being matched with a healthy roommate appear to significantly outweigh the losses to the healthy roommate from having a sick roommate. This approach echoes the earlier hospital roommate study that found benefits from matching preoperative patients with postoperative patients (Kulik, Mahler, and Moore 1996).

According to my estimates, negative matching could produce a net reduction of 12.3 hours in the length of stay (the 19.3-hour gain to the sick roommate, minus the 7-hour loss to the healthy roommate) and $570 lower costs (the $1,700 gain to the sick roommate, minus the $1,130 loss to the healthy roommate) per roommate pair, or 6.15 hours and $285 per roommate, relative to positive matching. For the 1,787 positively matched patients in my sample (846 low/low and 941 high/high matches), purposeful reassignment based on negative matching has the potential of reducing the overall number of inpatient days by 458 days and the hospital costs by $509,295, during the six-month study period, or over 900 days and $1 million annually. However, as the gender-specific results (Online Appendix J) suggest, the benefits of negative matching are more likely to occur in female than in male roommate pairs.

An additional potential to improve patient outcomes in acute care lies in the possibility that some very ill patients may benefit from being purposely assigned to semiprivate rooms, in contrast to the current practice of preferential assignment of such patients to private rooms. While my study focuses on patients in semiprivate rooms only and my findings cannot be generalized outside of my study population, the possibility of refining the room assignment process beyond just relying on procedure and diagnosis to consider a more granular set of patient characteristics needs to be carefully explored in further studies.

I conduct an extensive series of robustness and falsification tests that I discuss briefly here and in detail in Online Appendix L. In the robustness tests (Table L1), I examine the extent to which roommate assignments that occurred later in hospitalization may have involved "unfriending" or purposeful matching on unobservable patient characteristics, which the hospital staff may be learning about through personal interactions with the patient. I examine the robustness of my results to restricting the sample to only the first roommate per patient, only patients who were assigned the first roommate within two hours of admission, and only patients who did not have a prior hospitalization. The estimates are robust to these alternative sample inclusion criteria. In a series of falsification tests (Table L2 in the Online Appendix), I test for false roommate influences from "non-roommates" (i.e., patients who were hospitalized in the same room but at different times, at the same time but in different rooms) and from randomly selected roommates using Monte-Carlo simulations. I find that peer influences from non-roommates are small, noisy, and have the opposite signs to my main set of estimates, and Monte-Carlo simulation estimates are very close to precisely estimated zeros. The estimates suggest that unobserved selection is an unlikely driver of the findings presented in Table 4.

B. MECHANISMS OF PEER INFLUENCE

As discussed earlier, roommate influences in acute care can be of a direct nature (i.e., social norms, learning, and informal patient-to-patient caregiving) and of an indirect nature (i.e., rivalry for access to limited care resources and noise/inconvenience spillovers). While I are unable to distinguish between all of these mechanisms, I examine the extent to which my findings might be driven by the two types of indirect influences. First, the analysis of peer influences among patients hospitalized concurrently on the same unit but in different rooms (so that they are never directly exposed to each other) can shed light on the extent to which rivalry for care resources might be driving roommate influences. The premise is that given limitations imposed by the fast-paced and complex acute care environment, the amount of care resources that are available to a patient might be reduced following an elevation in the acuity level of another patient (regardless of whether the other patient is in the same or in a different room). If peer influences among hospital roommates are driven not by direct patient-to-patient influence but by rivalry for limited care resources, I expect that peer effects from patients who are not roommates will also be positive. However, I do not find peer influences from non-roommates, with all of the coefficients being small in magnitude, nonsignificant, or in the opposite direction of the main set of findings (bottom panel of Table L2 in the Online Appendix).

Second, if noise, inconvenience, and sleep interruptions associated with having a very sick roommate play an important role in generating positive roommate effects from a healthy roommate, one might expect that the magnitude of the peer influence coefficient in model 2 would be greater for roommate exposure spells that occur during the night shift than for roommate spells that occur during the day. To examine this, I use the timing of the start of each roommate spell and the length of the spell to generate an indicator equal to one for roommate spells during which the night-shift hours (8 p.m. to 8 a.m.) comprised more than 50 percent of the entire duration of the spell. I then estimate a modification of model 2 with an interaction term between the roommate's clinical condition measure and the night-shift indicator, and estimate the marginal effect of a roommate's clinical condition for the primarily day-shift spells and primarily night-shift spells, separately. I then test the equality of the day- and night-shift margins. A finding of a greater roommate effect during the primarily night-shift spells would be consistent with the presence of care delivery spillovers due to noise and sleep interruptions. If roommate effects are larger when the patients are exposed to the roommate primarily during the day, direct roommate effects are more likely to underlie the observed peer influence. The estimates do not support the negative care delivery spillovers mechanism of peer influence (Table 7). On the contrary, the magnitudes of the roommate influence effects are greater for roommate spells that occur mostly during the day shift (for both genders). These findings point to direct patient-to-patient influences as the main mechanism of roommate influence in my data.

VI. Limitations and Future Extensions

Although my study has several strong features, including plausibly random variation in peer characteristics and a multiple small peer group design, the nonexperimental nature of my study imposes limits on the internal validity of the estimates. The data come from patients hospitalized in shared rooms at a single hospital; therefore, the external validity of the results is also limited to similar study settings and patient populations. Measurement error[9] and the inclusion of an extensive set of fixed effects may have introduced attenuation bias and reduced precision, further reducing the predictive validity of my findings. An intervention study, one that is motivated by the results presented here and involves a

9 The measure of health (i.e., the Rothman Index) may not accurately reflect the patient's overall health status (more discussion is in Online Appendix C). The readmission variables include only readmissions to the same hospital and could miss as many as 15 percent of all readmissions according to the hospital's internal estimates. Electronic updates in the patient's RI score do not capture care that is not logged; nor do I observe the event itself that triggers a particular update. Therefore, my quantity of care variable is a proxy and not a direct measure of the amount of care received by a patient on any given day. The cost of hospitalization was obtained from the hospital's internal accounting systems and is not a precise measure of the actual costs of hospitalization at the individual patient level.

TABLE 7. Roommate effects on hospitalization outcomes, day/night analysis

Predicted margins	Daily quantity of care (Mean = 14.39)	Log of length of stay (Mean = 1.14)	Log of total cost (Mean = 9.08)	Discharge RI (Mean = 79.31)	Readmission, 7 days (Mean = 0.06)	Readmission, 30 days (Mean = 0.14)
d(y)/d(R's RI) | Day	−0.085[a]	−0.002[a]	−0.002[a]	0.015	0.000	0.000
	(0.007)	(0.001)	(0.001)	(0.010)	(0.000)	(0.000)
d(y)/d(R's RI) | Night	−0.068[a]	−0.001	0.000	0.011	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.013)	(0.000)	(0.001)
p-value of difference	[0.155]	[0.196]	[0.165]	[0.779]	[0.529]	[0.638]

Notes: $N = 3{,}793$. Shown are the predicted peer effect sizes from a variant of model 2 with an interaction term of the roommate's RI (R's RI) with an indicator for whether exposure was mostly during the night hours. Robust standard errors of the effect sizes are in parentheses. The p-values of the test of equivalency of the margins with the day exposure category are in brackets. $^{a}p < 0.01$, $^{b}p < 0.05$, $^{c}p < 0.10$.

random peer group manipulation on a large diversified patient sample, should be the next step toward establishing definitive causal evidence of patient-to-patient effects in acute care.

Future studies of patient-to-patient effects in acute care should explore several important extensions and directions. First, careful examination of the types of patient-roommate matches with the largest expected net benefit is critical to designing interventions with the highest potential to improve outcomes and reduce costs. In this initial exploration, negative matches where the sickest female patients shared the room with the healthiest female patients were associated with a net improvement in outcomes. This suggests that the patient's clinical condition is one of the observable patient characteristics that can be used to improve the roommate assignment process, but there could be others. Second, future research efforts should focus on identifying the mechanisms of peer influence. While data on frequency and the type of roommate interactions and other social support (family members, friends) were not available in this study, such data would allow for a direct study of the mechanisms of transmission of peer influence, which can further inform effective roommate manipulation interventions.

Finally, my findings of positive roommate effects point to the possibility that some patients may do better with a roommate than being alone in a private room. Several earlier studies of hospital patients show that 10–30 percent of patients prefer sharing a room with another patient because they enjoy the shared experience and wish to avoid feelings of isolation; patients tend to be elderly, have limited social support outside of the hospital, and suffer from medical conditions associated with a poor prognosis (e.g., oncology patients) (Hill-Rom 2002; White 2003; Contemporary Longterm Care 1997; Kirk 2002; Pease and Finlay 2002). One earlier study compared patient outcomes between private and shared rooms and also found evidence of clinical benefits of the shared room for the elderly patients with delirium, including lower mortality, fewer falls, and less frequent use of pain medication (Flaherty et al. 2003). While my study does not investigate this issue empirically—because the data lack an exogenous source of variation for identifying the effect of a private versus semiprivate room assignment—an intervention study of a random assignment of clinically eligible patients to private versus semiprivate rooms could produce robust evidence of health benefits from roommate interactions that may naturally arise in a shared room setting. Such evidence would be highly policy-relevant in the context of the current industry-wide shift toward all-private-room hospital designs (Knutt 2005), especially considering significant cost advantages of semiprivate rooms.

VII. Summary and Conclusions

This study examines social spillovers in health among hospital patients using quasi-random assignments of hospital roommates as a natural experiment. I find that patients sharing a room with healthier roommates have lower daily frequency of care, are discharged from the hospital sooner, and have lower costs of hospitalization, with little to no negative health effects. On the contrary, female patients who have relatively healthy roommates are discharged in better clinical condition and have a slightly lower probability of readmission, relative to female patients whose roommates are more acutely ill.

These findings are consistent with my theoretical conceptualization that views the roommate effect as an exogenous change in the production function of patient's health. In response to a roommate-induced shift in the production function of the subject patient's health, the hospital is able to adjust the amount of care provided to the subject patient and lower the costs, with little or no reduction, and in some cases improvement, in the output of the subject patient's health. I find that the care reduction effect is the strongest for publicly insured patients, which further supports the idea that the effects I observe could be a manifestation of a hospital's profit-maximizing (or cost-minimizing) behavior in response to a positive roommate effect on the subject patient's health.

In exploring the mechanisms of peer influence, I find that the roommate effects are not driven by rivalry for limited care resources and do not appear to be due to noise spillovers; instead, the effects appear to be direct patient-to-patient influence, presumably through social norms, learning, or informal care mechanisms. These findings are important from the policy perspective because they are indicative of policy multipliers in health care and suggest that the societal benefit of treating a patient may significantly exceed the patient's private benefit typically used in clinical cost-effectiveness computations.

The study points to significant potential benefits that could result from negatively matched bed assignments that pair sick patients with healthy roommates. My back-of-the-envelope computations suggest that reassigning positively matched roommates could save the hospital in my study over 900 inpatient days and $1 million annually in the cost of labor and other resources. However, more research is needed to explore the mechanisms of roommate influence and examine generalizability of these results, especially in the light of the Angrist (2014) critique of nonexperimental peer effect studies. Additionally, hospital administrators need to be aware of potential negative care quantity effects that appear to be potentially harmful for some patients (e.g., high-acuity male patients in my study). A carefully tested, comprehensive strategy, one that combines a purposeful roommate matching approach with care processes ensuring that the needs of patients are not being overlooked, has the potential to improve patient outcomes and reduce costs of hospital care.

ACKNOWLEDGEMENTS

I would like to express my gratitude to Elizabeth Bradley for invaluable guidance and support during this project, and to Jason Fletcher for helpful feedback during the early stages of the project. I am also thankful to Joan Rimar for her help in identifying and locating the data sources, and to many Yale New Haven Hospital employees for assistance in extracting the data and providing many helpful and informative conversations. I am grateful to Michael Rothman for all of the help and guidance he was able to provide about the clinical and administrative aspects of the Rothman Index computations and software. I would also like to thank participants at the departmental seminar in the Department of Health Management and Policy at the Yale School of Medicine, the Social Networks Research Group at Yale University, the Economics Department at Yale University, the University of Michigan Nursing School, and the George Washington University Nursing School. This paper was presented at the ASSA 2014 meeting in Philadelphia, Pennsylvania, and at the ASHEcon 2014 meeting in San Diego, California. Work on this project was partly funded through AHRQ (5 T32 HS 017589) postdoctoral training grant. This project was also funded through a faculty research grant from the Marquette University College of Business Administration.

REFERENCES

Album, Dag. 1989. "Patients' Knowledge and Patients' Work. Patient: Patient Interaction in General Hospitals." *Acta Sociologica* 32 (3): 295–306.

———. 2010. "Close Strangers: Patient-Patient Interaction Rituals in Acute Care Hospitals." In *The Contemporary Goffman*, edited by Michael H. Jacobsen, 352–72. New York: Routledge.

Angrist, Joshua D. 2014. "The Perils of Peer Effects." *Labour Economics* 30: 98–108.

Asphjell, Magne K., Lena Hensvik, and Peter J. Nilsson. 2013. "Businesses, Buddies and Babies: Fertility and Social Interactions at Work." Center for Labor Studies, Uppsala University, Department of Economics, Working Paper Series No. 2013: 8.

Azoulay, Arik, Nadine M. Doris, Kristian B. Filion, Joanna Caron, Louise Pilote, and Mark J. Eisenberg. 2007. "The Use of the Transition Cost Accounting System in Health Services Research." *Cost Effectiveness and Resource Allocation* 5 (11). doi:10.1186/1478-7547-5-11.

Bobrow, Michael, and Julia Thomas. 2000. "Multibed versus Single-Bed Rooms." In *Building Type Basics for Healthcare Facilities*, edited by Richard Kobus, Ronald L. Skaggs, Michael Bobrow, Julia Thomas, Thomas M. Payette, Sho-Ping Chin, and Stephen A. Kliment, 145–57. New York: John Wiley & Sons.

Boozer, Michael, and Stephen E. Cacciola. 2001. "Inside the 'Black Box' of Project Star: Estimation of Peer Effects Using Experimental Data." Yale Economic Growth Center Discussion Paper No. 832.

Bradley, Elizabeth, Olga Yakusheva, Leora Horowitz, Jason M. Fletcher, and Heather Sipsma. 2013. "A Simple Tool to Identify Patients at Increased Risk of Readmission within 30 Days." *Medical Care* 51 (9): 761–66.

Brown, Katherine K., and Dennis Gallant. 2006. "Impacting Patient Outcomes through Design: Acuity Adaptable Care/Universal Room Design." *Critical Care Nursing Quarterly* 29 (4): 326–41.

Cacioppo, John T., James H. Fowler, and Nicholas A. Christakis. 2009. "Alone in the Crowd: The Structure and Spread of Loneliness in a Large Social Network." *Journal of Personality and Social Psychology* 97 (6): 977–91.

Carrell, Scott E., Mark Hoekstra, and James E. West. 2011. "Is Poor Fitness Contagious? Evidence from Randomly Assigned Friends." *Journal of Public Economics* 95 (7–8): 657–63.

Carrell, Scott E., Frederick V. Malmstrom, and James E. West. 2008. "Peer Effects in Academic Cheating." *Journal of Human Resources* 43 (1): 173–207.

Carrell, Scott E., Bruce I. Sacerdote, and James E. West. 2013. "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." *Econometrica* 81 (3): 855–82.

Case, Anne C., and Lawrence F. Katz. 1991. "The Company You Keep: The Effects of Family and Neighborhood on Disadvantaged Youth." NBER Working Paper No. 3705.

Christakis, Nicholas A., and James H. Fowler. 2007. "The Spread of Obesity in a Large Social Network over 32 Years." *New England Journal of Medicine* 357 (4): 370–79.

———. 2008. "The Collective Dynamics of Smoking in a Large Social Network." *New England Journal of Medicine* 358 (21): 2249–58.

———. 2012. "Social Contagion Theory: Examining Dynamic Social Networks and Human Behavior." *Statistics in Medicine* 32 (4): 556–77.

Contemporary Longterm Care. 1997. "A Room of One's Own." *Contemporary Longterm Care* 20 (8): 14.

Duffin, Christian. 2002. "Private Rooms in Hospitals Would Hasten Recovery." *Nursing Standards* 16 (37): 8.

Duflo, Esther, Pascaline Dupas, and Michael Kremer. 2011. "Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya." *American Economic Review* 101 (5): 1739–74.

Eagly, Alice H., and Carole Chrvala. 1986. "Sex Differences in Conformity: Status and Gender Role Interpretations." *Psychology of Women Quarterly* 10 (3): 203–20.

Eisenberg, Daniel, Ezra Golberstein, and Janis L. Whitlock. 2014. "Peer Effects on Risky Behaviors: New Evidence from College Roommate Assignments." *Journal of Health Economics* 33:126–38.

Eisenberg, Daniel, Marilyn F. Downs, and Ezra Golberstein. 2012. "Effect of Contact with Treatment Users on Mental Illness Stigma: Evidence from University Roommate Assignments." *Social Science & Medicine* 75 (6): 1122–27.

Flaherty, Joseph H., Syed H. Tariq, Srinivasan Raghavan, Sanjeev Bakshi, Asif Moinuddin, and John E. Morley. 2003. "A Model for Managing Delirious Older Patients." *Journal of the American Geriatrics Society* 51:1031–35.

Fowler, James H., and Nicholas A. Christakis. 2008. "Estimating Peer Effects on Health in Social Networks." *Journal of Health Economics* 27 (5): 1386–91.

Glaeser, Edward L., Bruce I. Sacerdote, and Jose A. Scheinkman. 2003. "The Social Multiplier." *Journal of the European Economic Association* 1 (2–3): 345–53.

Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo. 2009. "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments." *American Economic Journal: Applied Economics* 1 (4): 34–68.

Hanushek, Eric A., John F. Kain, Jacob M. Markman, and Steven C. Rivkin. 2003. "Does Peer Ability Affect Student Achievement?" *Journal of Applied Economics* 18 (5): 527–44.

Hill-Rom. 2002. *The Patient Room of the Future*. Batesville, IN: Hill-Rom Publications.

Kapinos, Kandice, and Olga Yakusheva. 2011. "Environmental Influences on Young Adult Weight Gain: Evidence from a Natural Experiment." *Journal of Adolescent Health* 48 (1): 52–58.

Kirk, Steve. 2002. "Patient Preferences for a Single or Shared Room in a Hospice." *Nursing Times* 98 (50): 39–41.

Knutt, Elaine. 2005. "Healthcare Design: Build for the Future." *Health Service Journal* 115 (5940): 35–37.

Kulik, James A., and Heike I. Mahler. 1987. "Effects of Preoperative Roommate Assignment on Preoperative Anxiety and Recovery from Coronary-Bypass Surgery." *Health Psychology* 6 (6): 525–43.

Kulik, James A., Heike I. Mahler, and Philip J. Moore. 1996. "Social Comparison and Affiliation under Threat: Effects on Recovery from Major Surgery." *Journal of Personality and Social Psychology* 71 (5): 967–79.

Kuziemko, Ilyana. 2006. "Is Having Babies Contagious? Estimating Fertility Peer Effects between Siblings." Unpublished manuscript, Harvard University. http://www0.gsb.columbia.edu/mygsb/faculty/research/pubfiles/5799/fertility_11_29_06.pdf.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60 (3): 531–42.

Mears, Daniel P., Matthew Ploeger, and Mark Warr. 1998. "Explaining the Gender Gap in Delinquency: Peer Influence and Moral Evaluations of Behavior." *Journal of Research in Crime and Delinquency* 35 (3): 251–66.

Moffitt, Robert A. 2001. "Policy Interventions, Low-Level Equilibria, and Social Interactions." In *Social Dynamics*, edited by S. N. Durlauf and P. H. Young, 45–82. Cambridge, MA: MIT Press.

Pease, Nikki Jane F., and Ilora G. Finlay. 2002. "Do Patients and Their Relatives Prefer Single Cubicles or Shared Wards?" *Palliative Medicine* 16 (5): 445–46.

Rothman, Michael J., Steven I. Rothman, and Joseph Beals. 2013. "Development and Validation of a Continuous Measure of General Patient Condition Spanning the Hospital Acuity Spectrum Using Common Electronic Medical Record Data." *Journal of Biomedical Informatics* 46 (5): 837–48.

Rothman, Steven I., Michael J. Rothman, and Alan B. Solinger. 2013. "Placing Clinical Variables on a Common Linear Scale of Empirically Determined Risk: A Step Toward Construction of a General Patient Condition Score from the Electronic Health Record." *BMJ Open* 3 (5): 1–8.

Rothman, Michael J., Alan B. Solinger, Steven I. Rothman, and G. Duncan Finlay. 2012. "Clinical Implications and Validity of Nursing Assessments: A Longitudinal Measure of Patient Condition from Analysis of the Electronic Medical Record." *BMJ Open* 2 (4): e000849.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *Quarterly Journal of Economics* 116:681–704.

Tamres, Lisa, Denise Janicki, and Vicki S. Helgeson. 2002. "Sex Differences in Coping Behavior: A Meta-analytic Review." *Personality and Social Psychology Review* 6 (1): 2–30.

Trogdon, Justin G., James Nonnemaker, and Joanne Pais. 2008. "Peer influences in adolescent overweight." *Journal of Health Economics* 27 (5): 1388–99.

Ulrich, R. S. 2003. "Creating a Healing Environment with Evidence-Based Design." American Institute of Architects Academy of Architecture for Health Virtual Seminar: Healing Environments.

White, Robert D. 2003. "Individual Rooms in the NICU: An Evolving Concept." *Journal of Perinatology* 23:S22–24.

Yakusheva, Olga, and Jason Fletcher. 2014. "Learning from Teen Childbearing Experiences of Close Friends: Evidence Using Miscarriages as a Natural Experiment." *Review of Economics and Statistics* 97 (1): 29–43.

Yakusheva Olga, Kandice Kapinos, and Daniel Eisenberg. 2014. "Estimating Heterogeneous and Hierarchical Peer Influences on Body Weight Using Roommate Assignment as a Natural Experiment." *Journal of Human Resources* 49 (1): 234–61.

Yakusheva, Olga, Kandice Kapinos, and Marianne Weiss. 2011. "Peer Effects and the Freshman 15: Evidence from a Natural Experiment." *Economics and Human Biology* 9 (2): 119–32.