

Does Homelessness Prevention Work? Evidence from New York City's HomeBase Program

Peter Messeri, Mailman School of Public Health, Columbia University, New York

Brendan O'Flaherty, Department of Economics, Columbia University, New York

Sarena Goodman, Department of Economics, Columbia University, New York

July 11, 2012

Abstract: In 2004, New York City established HomeBase in order to reduce the number of families entering its homeless shelters. Families who think they are in danger of becoming homeless can go to HomeBase offices to receive a wide variety of assistance, including financial help and counseling, to keep them out of shelters. HomeBase started in different neighborhoods at different times. We use this variation in start-up to estimate the effect of HomeBase on shelter entries and exits between 2004 and 2008. Our best estimates are that for every hundred families HomeBase enrolled, shelter entries fell by between 10 and 20. HomeBase had no discernible effect on the length of shelter spells. We also demonstrate that foreclosure initiations are associated with shelter entries. This is the first quasi-experimental community-level evaluation of a homelessness prevention program.

KEYWORDS: homelessness, homelessness prevention, foreclosures, housing

Acknowledgments: We are grateful to the New York City Department of Homeless Services for data, financial assistance, and answers to our many questions; especially to Jay Bainbridge, Joanna Weissman, Eileen Lynch, Sara Zuiderveen, Jonathan Kwon, Veronica Neville, and Ellen Howard-Cooper. Scott Auwater and Bronxworks provided us with a detailed look at an operating program. Ingrid Ellen gave us lispens data. We are grateful for support from our institutional partners at CUNY, John Mollenkopf and Ellen Munley; and from the Columbia Center for Homelessness Prevention Studies, especially Carol Caton, Bill McAllister, Sue Marcus, Sue Barrow, Mireille Valbrun and Shoshana Vasheetz. We have received helpful advice from Serena Ng, Kathy O'Regan, Bernard Salanie, Beth Shinn and Till von Wachter; and participants at conferences sponsored by the National Alliance to End Homelessness, NYC Real Estate Economics Group, and the Columbia Population Research Center. Maiko Yomogida, Abhishek Joshi provided excellent research assistance. Financial assistance from the New York City Department of Homeless Services, the National Institute of Mental Health (5 P30MH071430-03), and the National Institute of Child Health and Human Development (1R24D058486-01 A1) is gratefully acknowledged. The errors are our own, as are the opinions expressed.

1.0 Introduction

HomeBase (HB) is a community-based homelessness prevention program that the New York City Department of Homeless Services (DHS) has operated since 2004. The program serves families experiencing housing emergencies in an attempt to avert their entry into the City's shelter system. The goal is to reduce the number of homeless families. The phased introduction of this program in different neighborhoods allows us to estimate the effect that this program had on family shelter entries and shelter spells between 2004 and 2008.

To our knowledge, this is the first quasi-experimental evaluation of homelessness prevention at a *community* level. Early evaluations of homelessness prevention programs asked whether their participants became homeless within some time frame (see Apicello 2008, Apicello et al. 2012, Burt et al. 2005). These evaluations did not convincingly establish counterfactuals to describe the shelter population absent these programs. Specifically, no credible case can be made for how many participants would have become homeless in the absence of these programs, or how long they would have stayed homeless if they had become homeless. We also learn nothing about the effects on non-participants. Two small controlled experiments on homelessness prevention have been attempted. The Western Massachusetts Tenancy Preservation Program compared the housing outcomes for 366 cases they closed with those for 21 individuals who had been wait-listed. Abt Associates is conducting a randomized experiment on a later version of HB in New York; they are currently following a control group of 200 and a treatment group of 200 (Buckley 2010). Results were not available at the time our study was completed.

However, even perfectly designed and executed controlled experiments with large samples of individual participants cannot give us as full a picture of the effects of homelessness prevention as can good community-level evaluations. Experiments estimate program effects for the treated under

conditions that may not reflect program operations under normal operations. Moreover, homelessness prevention programs might affect non-participants. For instance, they could divert limited resources—subsidized or cheap housing, legal assistance, private charity—from non-participants, and hence make non-participants more likely to be homeless. (By definition, successful homelessness prevention efforts increase the demand for housing in the short run, which may raise prices.) Homelessness may also be contagious. The knowledge of one family entering a shelter may encourage others to do so by lessening stigma, by showing them how to do it, or by giving them friends in the shelter. It is equally possible that the HB effect may be contagious. Non-participating families may learn from an HB family how to resolve their housing crises without resorting to shelter.

Furthermore, a thorough evaluation of a homelessness prevention program should estimate not only how many non-participants enter or avoid homelessness because of the program, but also how long they stayed homeless, or would have. Finally, the two individual-level studies we are aware of use applicants for the program as the control group. This procedure estimates the effect of the program correctly only if the program's existence does not change the incidence of the problems that the program is trying to alleviate. Our methods do not require this strong assumption.

Community-level studies like ours can estimate all these effects, but individual-level studies, even experiments, cannot. We estimate the effect of HB on both participants and non-participants within communities receiving HB services. Because we use communities as experimental controls, we are also able to estimate whether HB shifts the prevalence of homelessness from eligible to ineligible communities. The cost, however, is that we can estimate only the sum of effects at the community-level, not the effects on participants or non-participants separately.

Because we have enough detailed data on the regular operations of the New York City family shelter system, we are able to construct multiple plausible counterfactuals, and therefore we are able to

come much closer to a complete analysis of the effects of a prevention program than previous work. However, because we consider a range of plausible counterfactual scenarios, we present a range of point estimates as our best approximation of “the” HomeBase effect.

We estimated the effect of HB on shelter entries along two different dimensions. One dimension is *geography*: we use New York City community districts (CDs) and census tracts (CTs) as alternative geographic designations of neighborhoods. Census tracts are more plentiful and exhibit greater variation in intensity of HB services, but community districts have cleaner definitions of when HB was operating and because they are larger than census tracts, allow us to net out at least some of the potential role that HB might play in increasing or decreasing entries by non-participants. The other dimension is *how we quantify HB operations*. In some equations, the key explanatory variable is a measure of whether HB had begun operation in a particular geography—HB **coverage**— and in other equations the key explanatory variable is the number of families that HB serves—HB **service**. The coverage equations are simpler and more direct, but the great variation in the actual level of services that different offices provide means that they obscure a great deal of the heterogeneity in how many entries HB might divert in a month. The service equations resolve this heterogeneity issue but have to be estimated by instrumental variables techniques, and raise questions about the timing of HomeBase effects.

To preview study findings, HB was associated with a statistically significant reduction in shelter entries in the neighborhoods in which it was operating. Between the start of limited operations in November 2004 and November 2008, we estimated that HomeBase reduced shelter entries by 10 to 20 for every 100 cases opened. There was great heterogeneity in the effect of HB on shelter entries. Alternative statistical models resulted in a large range of point estimates related to both choice of geographic unit of analysis, CT or CD, and counterfactual scenario. HB appears to avert entries, at least

on net, not divert or delay them. We also present evidence that suggests that this reduction does not come at the expense of additional entries in nearby neighborhoods.

We also look at the rate of shelter exits for families who entered the system when HB was operating. A priori, an effective prevention program might have conflicting effects on spell length. On one hand, it might disproportionately divert the families with the least serious problems and hence those who would have had the shortest spells; this would raise the average length of shelter spells. On the other hand, if HB delayed shelter entry for some families and these families ended their homelessness when their luck turned better, independent of what occurred in the shelter, then HB would shorten shelter spells. We find essentially no effect; HB appears to reduce the length of shelter spells but the effect is small and statistically insignificant.

Besides estimating the effects of HB, we also examine the impact of foreclosures on entries to homeless shelters. These are the first results we are aware of that attempt to link foreclosures to homelessness. Specifically, we consider *lispendens* (LP) filings, the first step in the lengthy foreclosure process in New York that may typically take 12 to 18 months from LP filing to auction. A hundred *lispendens* filings were associated with 3-5 additional shelter entries over the next 18 months.

The plan of the paper is the following. We begin with a history of HomeBase and a description of the data. The third section provides summary statistics and information about trends in shelter entries and HomeBase cases. The fourth section estimates the coverage effect, using both community-district and census-tract data, and the fifth section estimates the service effect. Section 6 addresses the questions of robustness and spillovers—both to non-participants in the same month and to the same participant in later months. Section 7 analyzes exits and spell length.

2.0 Background

2.1 HomeBase history

For planning and land use purposes, New York City is divided into 59 community districts (CDs), each with about 140,000 people (the largest CD, Flushing, would have been the 75th largest city in the United States in 2010). The Department of Homeless Services (DHS) began HomeBase in November 2004 by selecting non-profit agencies in six CDs to operate the program. These CDs were not chosen randomly, although they were dispersed over four of New York City's five boroughs: DHS chose CDs that were heavily represented among the last addresses of shelter entrants. (This is a weakness in our study; an ideal community-level study would randomize selection of communities in which the treatment was available.) We will refer to these CDs as the "big six," because the majority of HB participants during the period we studied came from these six CDs. During the 22-month pre-program period (January 2003 through October 2004), the average number of shelter entries from the big six CDs was about 2.3 times as large as the citywide average. These six included the CD with the most shelter entries in the pre-program period, and even the big six CD with the smallest number of family shelter entries was ranked 14th among all 59 CDs. Only residents of the big six CDs were eligible to receive HB services between the start of operations in November 2004 and June 2007.

DHS expanded HB citywide in two phases. In July 2007, 31 more CDs were included in HB, and the remaining 22 CDs started in January 2008. The average CD that started in July 2007 had 21% fewer shelter entries in the pre-program period than the citywide average, and the average CD that started in January 2008 had 7% fewer. We will refer to these two sets of CDs as the "2007 cohort" and the "2008 cohort" respectively. After January 2008, HB operated everywhere in the city. In spring 2007, the initial contracts in the big six CDs were nearing completion, and because HB was going citywide, operators

knew that the contracts for these particular CDs would be smaller. The intensity of service in the big six CDs declined in spring 2007 and remained low for the rest of the period (see Table 2). The addition of each successive cohort expanded the set of families eligible to receive HB services.

By “families,” we mean any group of individuals who live together, and pregnant women. We include both families with children and families with no children (“adult families”). This is consistent with the definitions that DHS uses. HB also served single adults unaccompanied by children. Because we matched HB cases to family shelter entries, single adult cases were not part of our analysis.

HB helps families overcome immediate problems and obstacles that could result in loss of housing. Families experiencing difficulties voluntarily apply at an HB office located in their neighborhood, or are referred from shelter intake centers (but only when an HB office was operating in the family’s CD of origin). HB centers disseminate information through outreach activities in neighborhoods that are known to be the origin of large numbers of family shelter entries. HB case managers have wide discretion in matching services to the specific problems of eligible families. Services include family and landlord mediation, legal assistance, short-term financial assistance, mental health and substance abuse services, child care, and job search assistance. Our data do not allow us to distinguish among services or service-delivery strategies. The neighborhood base of operations—the fact that each participating family has an address, if only tenuously—is the key to our analysis.

HB offices are instructed to provide services only to eligible families. Eligibility includes an income criterion (200 percent of the federal poverty level during our study period) and, through December 2007, a residence criterion.¹ Although HB offices were supposed to serve only families living in their designated CDs, they did not always abide by this rule. About 6 percent of HB participants received

¹ Since 2008 13 HB centers provide services to families living within their respective catchment areas.

services before HB was operating in the CDs in which they appeared to live. Presumably they travelled to an HB office outside the CD in which they lived (DHS officials sometimes instructed HB staff to accept these ineligible families). Our models include a parameter to isolate the HB effect for families residing in ineligible CDs from the effect on officially eligible clients.

2.2 Data

Our primary dependent variables are family entries and exits for the New York City family homeless shelter system, disaggregated by time and geography. We study only eligible families who spent one or more days in the shelter system. Eligibility for the family shelter system essentially means a demonstration that a family has no reasonable alternative place to live. Determining eligibility often takes over a week, and families who are found ineligible often re-apply. We do not have information on families who entered the shelter system and were subsequently found ineligible or who left before an eligibility determination was made. HomeBase may have affected their number as well as the number of eligible entrants.

DHS supplied us with a list of all eligible families who entered the shelter system between January 2003 and December 2008. Because information for December 2008 appeared to be incomplete, we excluded it from our analysis of entries and of HB operations. Information for each family included family composition, their entry date, their exit date (if it was before December 31, 2008), their type of exit, and their last address. Names were redacted. The Center for Urban Research at CUNY attached a community district and a census tract to each address, and then redacted the address before sending us the data.

The operation of HB provides us with several independent variables—some of them measures of coverage, some measures of service. The simplest measure of coverage comes from a listing of when HB began operation in each CD, and the exact location of each HB office. We also have a file of all HB

participants from inception to December 2008. Like the shelter entrant file, this file has names and addresses redacted, but it includes the community district and census tract of the family's current address (not the office where they sought services). We also know eligibility status and date of enrollment. From this file, we compute the number of familyHB cases opened each month by the CDs and CTs they live in. (However, we are not able to tell whether the families that HB served later enter shelters; subsequent work by Abt Associates will examine this issue.)

For methodological and substantive reasons, we stratified CDs and CTs by number of entries during the 22-month pre-program period. In particular, we would like to know if HB worked differently in neighborhoods with many shelter entrants than it does in neighborhoods with few shelter entries. We assumed that the ordering of neighborhoods by characteristics that determine risk of homelessness—the supply of affordable rental apartments, socioeconomic status of residents, family and employment stability—were relatively stable during the six-year period of this study. Therefore, we measured a CD's or CT's use of family shelter as the count of families entering the shelter system by CD or CT of residence at time of entry for the 22 months preceding the start of HB program in November 2004. We assumed this stratification applied for the HB program period. A comparison of shelter entries before and after the start of the HB programs substantiates this assumption. Thus, the Pearson correlation between the count of families entering the shelter system by census tract for the 22 months just prior to the start of HB services and first four years of HB operations was 0.93. Even eliminating the many "low use census tracts", those with only 0, 1, or 2 families entering shelters during the pre-program period, the correlation only slightly declined to 0.90 for 894 census tracts with between 3 and 67 family entries during the pre-program period.

We divided the CDs into a high-shelter-use stratum that included 21 CDs that experienced above average number of entries during the pre-program period, and a low-use stratum that contained the

remaining 38 CDs. Figure 1 presents the geographical layout of the high- and low-use CDs, highlighting the big six (all of which are in the high-use CD stratum). The much more numerous CTs were divided into three shelter use strata. CTs with 0, 1, or 2 shelter entries during the pre-program period were grouped into a low use stratum, those with 3 to 18 shelter entries formed a moderate-use stratum, and those with 19 to 67 shelter entries formed a high-use stratum. Most census tracts, 995, fell into the low shelter use group. There are 702 moderate-use and 192 high-use census tracts. Figure 2 shows the distribution of NYC CTs by shelter use. This map illustrates that the moderate and high use census tracts are heavily concentrated in high use CDs, but the correlation is not perfect.

Finally, the Furman Center at New York University provided us with a file of all *lispendens* (LP) filings in New York City by date and address for the period 2001-2008. An LP filing is the first step in a foreclosure, although many LPs do not result in foreclosures². After the LP is filed, the foreclosure and resale process often takes a long time—a year or more.

3.0 Summary Statistics

Table 1 presents summary statistics. During the study period, 45,088 families entered the NYC shelter system and HB centers opened 10,987 cases. There were 115,320 LP filings during this period. Shelter entries averaged 10.7 families per month from each CD. On average 3.6 HB cases were opened per month in a CD during months in which HB was operating. The number of LP filings averaged 17.5 per CD-month. There were 1,889 NYC census tracts with at least one family shelter entry or HB case during the study period.

3.1 Trends in shelter entries

HomeBase began operations during an extraordinarily dynamic period in the use of the City's family shelter system. During the study period, 2003-2008, an average of 635 families entered the

² Of *lispendens* filings in New York City in 2007, only 14 percent had ended with bank ownership or third party auction by 2009; 54 percent had had no subsequent legal transactions. See Furman Center 2010.

shelter system each month, but the mean obscures pronounced seasonal and year-to-year fluctuations (see Figure 3). In any calendar year, entries tend to peak in the late summer months and ebb during the winter and early spring. The amplitude of seasonal variation is large, on the order of a 200 to 300 change in family monthly entries. The start of HomeBase in November 2004 occurred midway through a period of relatively low shelter use that extended through 2005: entries fluctuated between 400 and 600 each month. After 2005 family entries began to climb. The numbers of monthly family entries increased to between 600 and 800 during 2006, 2007, and the first half 2008. By the second half of 2008, monthly family entries jumped to over 1,000 as the burst of the housing bubble and the full force of the Great Recession began to take hold.

Of particular note for evaluating HomeBase was the sharp downward disjuncture in family entries that followed the start of the program in November 2004. Family entries attained a study period nadir for the seven months following the start of HomeBase operations in the big six CDs. The sharp fall-off in entries was in part due to seasonal declines, but a regression model that adjusts for both seasonal variation and annual change in family entries estimated this disjuncture to be a highly significant 75 (95% C.I.=27, 122) drop in monthly family shelter entries averaged over 50 months of partial and later full coverage HomeBase operations. It is not possible to draw strong causal inferences about the effects of HomeBase from a citywide aggregated time series, because of the possibility that changes in economic and housing market factors that coincided with the start of HomeBase operations were also influencing the observed secular changes in shelter entries. Figure 3 also illustrates the very “noisy” background in shelter entries that presents significant challenges in modeling the counterfactual condition that would have obtained in the absence of the HB program. For a more robust estimate of treatment effects, we turn to more detailed analyses at the CD and CT levels.

3.2 Trends in HB Cases

Annual counts of shelter entry and HB cases presented in Table 2 confirm the intended restriction of HB services to the relatively small number of very high use CDs during the early years of the HB program. During the two pre-program years 3,088 or 23% of all family shelter entries resided in the “Big Six” CDs. During the first two full years of HB operations, 2005 and 2006, the number of families entering the shelter system from the six CDs served by HB was substantially less than the number of HB cases opened in these communities. For every 10 shelter entries, approximately 16 HB cases were opened. The successive expansion of the HB program to all of New York starting in July 2007 sharply diluted HB services. The combined effect of dispersal of a constant level of HB services over a greatly expanded geographic area during a period of rising shelter entry resulted in a reversal of the early period with roughly three shelter entries for each HB case, for both the entire city and the Big Six CDs.

4.0 Shelter entries: Coverage analysis

4.1 Methods

4.1.1 CD level

In this section we treat HB operation as a series of binary variables, and see how changes in shelter entries in CDs subject to the “treatment” of HB operation differ from changes in untreated CDs. Specifically, the simplest equation is:

$$S_{ct} = \alpha + \beta H_{ct} + \gamma_c + \delta_t + \varepsilon_{ct}. \quad (1)$$

Here c indexes CDs and t indexes months, S_{ct} denotes the number of shelter entrants from CD c in month t , γ_c is a CD fixed effect, and δ_t is a month fixed effect. The key independent variable is H_{ct} , a

dummy variable equal to one if and only if HomeBase is officially operating in CD c in month t . The coefficient β is an estimate of the HB effect, the average number of shelter entries averted in a CD-month of HB operation. A negative β indicates that the treatment works: HomeBase reduces shelter entries.

Simple equation (1) can be improved in several ways. First, foreclosures are month- and CD-specific events that may affect shelter entries, since most households affected by foreclosure in New York City were probably renters (Furman Center 2010) or may reflect CD-specific housing market trends. Let F_{ct} denote the number of LP filings in CD c in month t . We add several lags of this variable to equation (1): contemporaneous, 3-month lag, 9-month lag, 12-month lag, 15-month lag, and 18-month lag. We use these lags because the time between the filing of an LP and the resolution of the foreclosure, including the displacement of residents, is often long.

Second, the official definition of when HB was operating in a CD does not properly account for treatment, since some families received HB services before HB was operating in their CD, as we noted in the discussion of the small number of apparently ineligible families receiving services. To address this problem, we add a second dummy variable P_{ct} , equal to one if and only if some resident of CD c has received HB services during or before month t . During normal operations of HB, both H_{ct} and P_{ct} equal one, and so we will be interested in the sum of their coefficients.

Third, treatment effectiveness may change as time goes on. One possibility is that HB offices become more proficient as they acquire more experience. In that case, later months of experience would be associated with greater reductions in shelter entries. On the other hand, HB may delay shelter entries rather than avert them entirely. Then the first months of operation would show the greatest reduction in total entries; after that, delayed entries would offset new reductions. Participants who come to HomeBase in later months when the program is better known may also differ systematically

from participants who made their way to HomeBase when it was little known. Another possibility is that as time passes, enthusiasm wanes and treatment intensity declines. To explore these possibilities, we include a dummy variable R_{ct} , equal to one if and only if in month t HB has been officially operating in CD c for more than two months.

To summarize, our fullest model of entries as a function of HB coverage is:

$$S_{ct} = \alpha + \beta_1 H_{ct} + \beta_2 P_{ct} + \beta_3 R_{ct} + \gamma_c + \delta_t + \sum_s \varphi_s F_{c(t-s)} + \varepsilon_{ct}. \quad (2)$$

Equation (2) is linear. While the functional form of the estimating equation does not matter for the argument that the coefficients on HB operation give average reductions in shelter entries associated with HB, it does matter for the interpretation of the fixed effects and hence the construction of the counterfactual (what would have happened if HB were not operating in the CD-months in which it was operating?). In particular, the form of equations (1) and (2) forces the month fixed effect to be the same in absolute magnitude in every CD: a city-wide shock like a change in shelter eligibility rules or a recession increases or decreases shelter entries in every CD by the same amount.

We address this problem in two different ways, neither of which is totally satisfactory. One approach is to estimate equation (2) in logarithmic form:

$$\ln(S_{ct} + 1) = \alpha + \beta_1 H_{ct} + \beta_2 P_{ct} + \beta_3 R_{ct} + \gamma_c + \delta_t + \sum_s \varphi_s F_{c(t-s)} + \varepsilon_{ct}. \quad (2')$$

The implicit counterfactual in equation (2') is that city-wide shocks operate in multiplicative fashion: they cause identical percentage changes in shelter entries in every CD. The possible drawback of equation (2') is that it forces entries averted to be a constant fraction of entries not averted. By contrast, linear equation (2) is less restrictive about the relationship between entries averted and entries not averted. Equation (2) provides a better model of HB operations if those operations are small relative to flows into the shelter system: if HB offices run out of resources to treat cases, for instance,

not cases to treat, or good cases to treat. If, for instance, an HB office is incapable of averting more than a small fixed number of cases in a month, no matter how many families are entering the shelter system, then the coefficient on HB operations in (2') will be biased down.

An alternative way to construct the counterfactual is to stratify the sample and estimate equation (2) separately for CDs that usually have many entries and CDs that usually have few. This allows us to estimate different sets of month effects for the CDs that are usually large and for the CDs that are usually small. We used the two strata based on pre-program shelter entries. Stratifying the sample and estimating two equations has an additional advantage in interpretation. Although we do not attempt to perform a cost-benefit analysis, anyone who wanted to use our results for this purpose could be guided by the stratified results in deciding whether to expand or contract the program in high- or low-use CDs.

Stratifying the sample and thereby multiplying the number of month fixed effects, however, has a downside. The month fixed effects are supposed to reflect city-wide shocks, but if they are not constrained to be equal or proportional across the two strata, they may not reflect the same shocks at all. Adding sets of month fixed effects almost automatically reduces the size and significance of coefficients that measure within CD and month variation, like those on HB operations and foreclosures. In the limit, if there were 59 strata so that each CD had its own set of month fixed effects, the effects of HB and foreclosures would be unidentified. Thus when we stratify, we will have to examine the resulting month fixed effects and see how well the month effects from one equation are correlated with those from the other—essentially, whether they are picking up the same city-wide shocks.

Including 18-month lags of foreclosures forces us to ignore 18 months of data, starting in 2003. We have estimated these equations over the full period without foreclosure lags. Results are not

materially different. We have also estimated these equations dropping the CD-months in which unofficial entries appear. Again, the results are not materially different.

4.1.2 CT level

Except for adjustments related to the change in unit of analysis, the CT level equations that estimate the average treatment effect of the HB program are identical to CD level coverage equations (1), (2) and (2'). Corresponding to the change in unit of analysis, the dependent variable is now monthly counts of family shelter entries for each census tract and CT fixed effects substitute for CD fixed effects. All models also include monthly fixed effects. The HB coverage variables are identical to those applied to the CD level analysis. That is to say, all CT's are assigned their CD's values for official and unofficial start of HB services. As in the CD-level analysis, the CT-level coverage equations are estimated in both linear additive and multiplicative forms. Because of the large number of CT-months with no shelter entries and limited range of shelter entries at the CT level, we substituted a Poisson regression model for the loglinear CD model. We also estimated CT-models stratified by three shelter use levels. The stratified Poisson model contains only one set of city-wide monthly fixed effects that are allocated proportionate to the CTs level of shelter use, whereas the stratified linear models estimated separate monthly effects for each stratum.

While there may be good reasons to prefer a linear additive model when examining HB operations at the CD-level, there are equally good reasons for preferring a loglinear link between the independent and response variables when fitting CT data. In contrast to conditions at the CD level where HB offices might run out of resources before they run out of good cases to treat (see section 4.1.1), the opposite is likely to obtain at the CT level. The many low shelter use CTs are likely to run out of treatable cases before HB resources run out. Thus we expect HB centers will tend to allocate resources within their catchments proportional to need at the CT level, as indexed by the number of families

entering the shelter. Our a priori reasoning is consistent with the empirical distribution of HomeBase cases. Low use census tracts averaged 0.040 new HB cases opened each month during period of official program eligibility. The mean number of new HB cases per months jumps to 0.34 and 0.96 for moderate and high use CT's, respectively. DHS reinforced the natural tendency of HB to target high-risk families from areas where shelter use is greatest by contracting services to nonprofit organizations located in or close to neighborhoods with the highest volume of family shelter entries. Under conditions where HB centers devote more resources to high use CTs, a constant HB effect for individual families would be expected to translate into a proportional effect when data are aggregated as counts at the CT level, which is how Poisson model coefficients are interpreted. The Poisson model has the additional virtue that in constructing a counterfactual condition, a single set of month fixed effect allocates city-wide shocks to the shelter system proportional to a CT's shelter use. Of course even if HB offices allocate resources roughly proportionate to CT-level need for services, this doesn't necessarily translate into equal effectiveness at the individual family level. We test this assumption when HB effects are estimated separately by CT's level of shelter space use.

The effect of foreclosures on shelter entries over an 18 month period was also estimated at the CT level. In a minor departure from the CD level analysis, all monthly lags from 1 to 18 months are estimated in the CT equations. Because of their much smaller size, monthly LP housing unit filings are relatively rare occurrences at the CT level. There were no LP filings in three-quarters of monthly CT observations and 1 to 5 LP filings accounted for another 23 percent of CT-months. At the other extreme, 52 monthly observations with the largest numbers of LP filings ranged between 59 and 690 filings. Despite the very small number of outlying CT-months (.04 percent of all observations), preliminary analysis indicated that these extremely high LP counts had a substantial depressing effect on estimates of family entries. To remove the distorting effect of very high LP counts, the LP count was truncated at 50 and an indicator variable for these outlying observations (1=observations with 50+

filings, 0 otherwise) was added for each lagged variable. Effects of foreclosure are estimated for the linear model, but not the Poisson model, since each additional foreclosure is assumed to have a fixed additive effect on shelter entries regardless of spatial aggregation or the volume of families entering the shelter system.

We also grouped CTs by the time that HB officially began operating—either November 2004, July 2007, or January 2008. Results of this analysis are available in the working paper version of this paper (Messeri et al. 2011).

4.2.0 Results

HB appeared to decrease entries to the shelter system. This holds at both the CD and CT levels.

4.2.1 CD level

Table 3 shows OLS regression estimates for all variants of the CD coverage model. HB operation appeared to reduce shelter entries by about 4.6 per CD-month. Since the average CD-month during our study period had about 10.7 shelter entries, this reduction is economically as well as statistically significant. The sign of the coefficient on experience indicates that HB offices became less effective (although not significantly so) after they had been open for two months; this suggests that delay probably played a role, though a small one, or that the client mix changed as the program was becoming better known in the community.

Foreclosure initiations appear to increase homelessness. The coefficients indicate that for every 100 LP filings, about five additional families enter shelters. The effect was strongest at 12 months after the initial filing. Since most LP filings do not result in foreclosures, the results are economically as well as statistically significant.

Results for the logarithmic equation (2') were not so strong. Coefficients have the same signs as they do for the linear model, but estimates are less precise. The estimated magnitude of the effect of HB is also somewhat smaller, but is still significantly different from zero when all the coefficients are added. HB produced a reduction of around 1.58 entries in the average fully operating CD-month, rather than over four with the linear specification.

The estimated effect of foreclosures was also smaller and less significant. This suggests that the LP effect was tied directly to displaced tenants, and not a general indication of housing market difficulties.

The linear model appears to fit the data better than the logarithmic model, but the fixed effects of the linear model run into the problems we discussed in the methods section. Consider CD 503, the south shore of Staten Island, a relatively affluent CD. During the pre-program period, less than half a family entered the shelter system from this CD in an average month. Yet the counterfactual with the linear model was that absent HB, around 9 families would have entered in August 2008.

Because of this problem, we present the results of the stratified equations in table 3. HB effects were quite small and insignificant for the stratum with low shelter use CDs, and while they were significant for the stratum with high use CDs, the effect was smaller than the unstratified effect. Specifically, on average the stratified coefficients imply that a fully operating CD averts 1.72 shelter entries a month, while the unstratified coefficients imply more than 4. (The 95 percent confidence intervals do not overlap.) The effect of foreclosures in the stratified equations was also smaller: 100 LP filings led to around three shelter entries, as compared to five with the unstratified equation. This effect was still highly significant, however.

4.2.2 CT level

Table 4 presents estimates of the effect of HB coverage on shelter family entries for the CT-level models. The pooled Poisson and linear models produced divergent estimates of HB effects, while the stratified model estimates lie between the two extremes. The estimated average HB program effect, presented as the incidence ratio, for the pooled Poisson model, after combining the three HB parameters, was .949 (95% C.I.=.897,1.001) or 5.5 families averted for each 100 shelter entries.---about half a family in average operating CD-month. The linear formulation of the coverage model estimated a much stronger HB program effect, possibly too large an effect: 0.13 (95% C.I.=0.106,0.153) family shelter entries averted for each month of official HB operations in a CT. Extrapolated to the CD level, this would imply about 4.2 entries averted in an average operating CD-month. When the linear coverage model is stratified by shelter use level, the HB effect is substantially reduced and is now much more in line with the estimated effect of the Poisson model. The stratified linear model indicates that the HB effect strengthened with increasing CT use for shelter services. Qualitatively similar results obtain for stratified Poisson models.³

5.0 Shelter entries: Service analysis

5.1 Methods

³We also estimated stratified Poisson models. Similar to the stratified linear model these models indicate that the HB effect strengthened with increasing shelter use. The point estimate and 95% confidence intervals for the incidence ratio for low, moderate and high use strata were respectively 1.25 (1.14,1.38), .927 (.871,.986) and .904 (.848,.964) The positive effect for the low use stratum we believe is not real, but an artifact of a fixed effect Poisson model, which drops a cluster, in this case a CT, when the outcome is constant, in this case it drops CTs with zero shelter entries, but retains CTs with only one or two shelter entries during the study period. Dropping the zero entry CTs removed CTs in which HB cases may have potentially averted shelter entries while retaining a matched group of low use CTs in which HB failures could have occurred. Thus a disproportionate loss of successful HB cases may have positively biased treatment effects in the low use CTs, whereas the linear model suggests the HB has no effect when serving low use census tracts.

5.1.1 CD level

An alternative approach is to assume that what affects shelter entries is not the simple presence of an HB office, but the number of families that the HB office serves. Both direct and indirect effects of HB depend on how many families are served. This relationship argues for making the independent variable the level of HB service in a month, not a description of the coverage of the area.

Let HB_{ct} denote the number of families living in CD c whom HB served in month t . Thus we would like to estimate a linear equation like:

$$S_{ct} = \alpha + \beta HB_{ct} + \gamma_c + \delta_t + \sum_s \varphi_s F_{c(t-s)} + \varepsilon_{ct}. \quad (4a)$$

or a quadratic equation like

$$S_{ct} = \alpha + \beta HB_{ct} + \check{\beta} HB_{ct}^2 + \gamma_c + \delta_t + \sum_s \varphi_s F_{c(t-s)} + \varepsilon_{ct}. \quad (4b)$$

However, HB service in a particular CD-month is likely to be endogenous—for instance, a CD-specific event that causes many families to visit HB is likely to cause many other families to enter shelters at the same time. So OLS will not produce unbiased estimates for equation (4).

Fortunately, administrative decisions provide us with several instruments for HB service, and so we estimate equation (4) by instrumental variables. The first instrument is obvious: H_{ct} , the variable indicating whether HomeBase was officially operating. CD-months when HomeBase was officially operating should see more HB cases than CD-months when it was not. The next instrument is the distance from the CD centroid in question to the HB office that was nearest that month. The distance is set to zero for CDs in which an operating HomeBase office is located.

We also use a series of instrumental variables to reflect administrative differences between groups of CDs and differences over time. For each of the three cohorts defined by starting date (the big

six, the 2007 cohort, the 2008 cohort) we have a dummy variable that turns on only for the CDs in this cohort, and only after the start of official HomeBase operations for that cohort. In addition, for each cohort and each fiscal year after HomeBase operation begins for that cohort, we have a dummy variable that turns on for CDs in that cohort in that fiscal year. We use fiscal year dummies because most DHS contracts are on a fiscal year basis and many important policy changes coincide with new contracts, including changes in funding levels that influence the overall number of HB cases the HB centers are able to open each year.⁴

Thus, for instance, for a CD-month when HomeBase is not operating, all instruments will be zero, except the distance to the nearest operating HomeBase office.⁵ For a CD-month in fiscal year 2009 for a CD in the 2007 cohort, the positive dummy variables will be H_{ct} , the 2007 cohort dummy, and the FY-2009-interacted-with-2007-cohort dummy.

These instruments are highly correlated with HB service; the R^2 on the first-stage CD-level regression for HB_{ct} is .797. The distance variable and H_{ct} are significant and their estimated effects have the right sign.

The other question about the instruments is whether they satisfy the necessary exclusion restriction. To be valid instruments, they must affect shelter entries only through the services that HB provided. Unconditionally the location of HB offices and the period of formal HB operations in a CD are correlated with number of shelter entries. Thus the big six CDs that were first to receive HB services were chosen because of their high rate of shelter entry. The expansion of HB services and the opening of new HB centers to serve the remaining CD occurred during the final 18 months of the study period when

⁴ Of course, we omit dummies as appropriate.

⁵ In contrast to the *coverage* model in which observations prior to the start of HB operations in November 2004 provide useful information, observations for the service model are restricted to the period of HB operations.

the rate of shelter entry was higher. However the inclusion of the month- and CD- fixed effects should remove the correlation with shelter entries mediated by the level of HB services. There remains the possibility that DHS may have re-allocated resources in response to transitory events in a particular month and CD, but we have no evidence for such a fine-tuned response. DHS created HB coverage through contracts with nonprofit organizations, and nonprofits had to gear up and acquire staff and space to implement the contracts. These are complex and time-consuming processes. HB was implemented in only three waves, moreover, so there was little scope for adjusting start-up dates. Hence it is likely that the effect of CD-and-month-specific events on HB coverage was at most small. To the extent it was present, however, it biases down our estimates of the effect of HB service on shelter entries.

Because of the size of our sample and because all of the big six CDs are in the large stratum, we were unable to produce meaningful stratified results.

5.1.2 CT level

The CT service model is estimated using the linear form and instrumental variables applied to the CD-level data. However, we do not estimate a quadratic form of the service model, as the observed number of HB cases opened in any month at the CT level is seldom more than 3 cases. In this study, we did not attempt to investigate the potential conditioning effect of this and other CD level attributes on CT level effects.

5.2 Results

5.2.1 CD level

Table 5 shows the results of the IV regressions, both the linear and quadratic versions. Both equations imply that HB averted shelter entries, but the size of the effect is considerably larger, on

average, in the quadratic equation than in the linear. In the linear specification, it appears that shelter entries fell by about 10.3 for every 100 families that HB serves. This value is consistent with the CT coverage results with the Poisson specification. The significant coefficient on the quadratic term indicates that the linear equation might be misspecified. The quadratic results imply that the marginal effectiveness of an HB office declined considerably as the number of families it serves increased. Taken literally, the quadratic results imply that after about 30 cases, more cases were counter-productive (but very few CD-months had more than 30 cases). Among CD-months with HB officially operating, the average number of cases from within the CD was 9.7 and the average number of squared cases was 296.99. This implies an average reduction in shelter entries of 2.6.

Why was HB more effective with small caseloads than with large? We do not know. Answering this question is clearly important for future decisions about the size of HB.

One possibility is simple congestion. With staff fixed, participants in CD-months with higher caseloads received less attention and outcomes may have deteriorated as a result.

Selection is another possible reason for the differential effectiveness. The size of the coefficients on P_{ct} suggests that selection is an important part of the story. These coefficients indicate that unofficial operation of HB in a CD reduced shelter entries by about two a month. Receiving services before the start of official operation (“pure unofficial operation”) means that residents of that CD had had to travel to another CD to apply for HB services, and then had been served—DHS may have permitted and even instructed HB centers to open cases for families living in ineligible CDs. During an average CD-month of pure unofficial operation, 0.58 families were served (many CD-months of pure unofficial operation had zero families because unofficial operation started when the first family in the CD received services). The astounding effectiveness of unofficial operation was probably an important part of the diminishing marginal effectiveness result.

Why is pure unofficial operation so effective? Families who travel outside their neighborhoods, may be special, especially those who can convince HB workers to bend the rules. Either the families themselves or the HB workers may recognize that they are in imminent danger of homelessness and that specific assistance can resolve that danger. The same may be true for HB offices that are small and not well-known. Convenient, well-publicized HB offices may draw many families that HB cannot help avoid homelessness.

In both regressions, the sum of foreclosure coefficients is positive and significant—slightly larger, in fact, than it was in the linear OLS regressions. The service results also imply a larger impact of foreclosures on shelter entries than the coverage results imply. In the long run, a hundred LP filings result in 8.2 more shelter entries. In the first-stage regression, moreover, foreclosures also increased HB cases. This suggests that our decision to use IV was not misguided: if foreclosures increased both shelter entries and HB cases, other CD-and-month-specific shocks probably did the same. (In fact, if equation (4a) is fit by OLS, the coefficient on HB cases is a positive and significant: HB appears to have increased shelter entries.) This estimate is higher than the estimate from the coverage equation, and most other estimates of foreclosure effects in this paper do not support such a large impact.

5.2.2 CT level

Table 5 also presents IV estimates for the CT service equation. The service model indicates a reduction of 12 shelter entries for every 100 HB cases opened. This is in line with the estimate obtained from the Poisson coverage model.

5.3 Reconciliation of coverage and service results

Table 6 and Figure 4 consolidate all of the service and coverage results and transform the

original coefficients into a consistent treatment effect metric: entries averted per 100 HB cases. We use historical accounting in this table; that is, the weighted average across the actual number of months in each condition (unofficial operation, official operation, and experienced operation) that occurred in during the study period. The results consistently show that HB reduced shelter entries, but are not consistent on the size of the reduction.

Generally, CD models imply larger effects than CT models. For coverage models with only one set of monthly fixed effects, however, the CD and CT results have the same pattern. For an unstratified linear model both data sets imply a very large reduction in entries, about 68 entries per hundred HB cases for the CD model and about 65 for the CT model. For the unstratified logarithmic or Poisson model, both data sets imply smaller reductions: about 25 entries per hundred cases for CD estimates and 12 per hundred for CT estimates. The difference between the CD and CT logarithmic estimates, however, is not statistically significant. For both data sets, the stratified linear model produces very similar results that are much closer to the unstratified logarithmic or Poisson models: 19 per hundred cases for CD estimates and 21 entries per hundred cases for CT estimates.

Since we have reason to believe that the unstratified linear results are too high, we are left with a range of 10 to 20 shelter entries averted per hundred HB cases.

With the service models, the CD and the CT linear models produce about the same estimate: 10 to 12 averted shelter entries per hundred HB cases. The CD quadratic estimated over twice this level, or 26 cases averted.

We should not be surprised that the CT linear and CD linear service estimates are about the same. Our instruments are at the CD level, not the CT level; they use no CT-specific information. Thus we are looking at the instrumental variation in the number of CT HomeBase cases that is driven by CD-level phenomena. By construction, month-to-month variation in HB cases is constrained to be the same

in each CT of a CD. Month-to-month variation in shelter entries in a CD is the sum of month-to-month variation in shelter entries in the component CTs, and so if month fixed effects are working about the same way on the CD and CT levels, the estimated effect of HB cases on shelter entries should be the same in the CT linear and CD linear equations. The CT linear and CD linear estimates will also differ because the number of CTs forming a CD differs.

Thus the approximate equivalence of the CT linear and CD linear estimates gives us confidence in the service equations. The CD quadratic equation nests the CD linear equation, and lets us reject the hypothesis of no nonlinear effects. The nonlinear effects we find on the CD level have no obvious counterpart on the CT level. (Congestion, for instance, arises from an imbalance between the resources of a CD level office and the total number of families seeking service, not matter how they are divided among CTs.)

6.0 Issues and possible overstatements

The effect of HomeBase on shelter entries might be overstated or understated for two reasons: because we may not have properly accounted for the effect of HomeBase on non-participants and because HomeBase may merely perturb the time pattern of shelter entries. In this section we discuss these two issues. Our estimates in the previous two sections, especially the coverage equations, appear to be net of both effects, but a great deal still needs to be learned.

6.1 “Musical chairs,” contagion, and other effects on non-participants

We do not know whether the shelter entries that HB averted would have been HB participants or other families; all that we can estimate is the net number of entries averted. But that is the relevant number for evaluating HB, not the number of entries averted among participating families.

HB may either increase or decrease entries among non-participating families. Increases would arise from “musical chairs;” decreases from contagion. A priori, there is no way of predicting whether musical chairs or contagion is more powerful.

Do spillovers to non-participants bias our estimate of the net effect of HB? We need to consider each of the four classes of estimate.

For CT coverage estimates, there is no bias. All CTs within a CD are either operating in a given month or not. Spillovers between CTs will cancel each other out in the aggregate. For CT service estimates, there would be bias if we used OLS or different instruments, but our instruments for CT service are all on the CD level. Instrumented CT service levels move in tandem within a CD, and so no bias arises.

For CD coverage estimates, a bias could arise if HB activity in an operating CD affected shelter entries in a neighboring non-operating CD. If musical chairs were the stronger effect, our estimates of entries averted would be too high; if contagion were the stronger effect, they would be too low.

To check for these geographic spillovers, we formed pairs of adjacent CDs, and looked at entries from each pair for each month. (Since there are an odd number of CDs and Staten Island with three CDs is relatively isolated, we consolidated all of Staten Island as “CD-pair.”)⁶If HomeBase service in a CD

⁶ Specifically, outside of Staten Island, we first constructed an adjacency table describing which pairs of CDs were adjacent. We consider the East River (but not the Harlem River or Newtown Creek) to be impenetrable (and so CDs on opposite sides of the East River are not adjacent, but CDs on opposite sides of the Harlem River are). (The Harlem River is treated differently from the East River because it is bridged more often.) Similarly we considered Central, Van Cortland, and Flushing Meadow Parks to be impenetrable, but not the Bronx Zoo or Forest Park.

We then arranged the boroughs in the order Manhattan, Bronx, Brooklyn, Queens; within boroughs CDs are numbered. We then did iterations of the following process:

affects non-participants in adjacent CDs adversely through musical chairs, then the estimated HomeBase effect estimated from CD-pairs will be smaller than the effect estimated from single CDs; if contagion is stronger, the CD-pair estimate will be bigger than the single CD estimate. Table 7 shows the result. The HB coverage effect is larger when estimated on double-CDs, but the difference is not significant. Thus, we find no evidence of sizable spillovers between adjacent CDs.

There still could be larger, more diffuse spillovers—HomeBase activity in Queens may affect non-participants in Staten Island. We have no way of checking for such spillovers. Since we have found no evidence of spillovers between adjacent CDs, we do not think net spillovers between non-adjacent CDs are likely to be large.

For CD service estimates, bias could arise from spillovers between CDs that were operating at a high level and CDs operating at a low level, as well as spillovers between operating and non-operating CDs. The former effect is a form of attenuation bias that bias our estimated coefficients down no matter whether contagion or musical chairs were operating. Because our preferred CD service equation is nonlinear, we were unable to form CD-pairs to check for this bias, but the CD coverage exercise suggests that it may not be large.

Thus the estimates in the previous sections appear to approximate the net HomeBase effect. The HomeBase effect on participants only may be different, and could be found with a randomized

-
- a. Check to see whether any unassigned CD has only one potential unassigned partner; if so, assign the CD and that partner.
 - b. When no unassigned CD has only one potential partner, assign the unassigned CD that comes first in order to the unassigned adjacent CD that comes first in order.
 - c. Start over at a.

Clearly this algorithm is not unique.

controlled experiment. Such an experiment, however, could not find the effect on everyone, which is the relevant measure. By combining the results from a randomized experiment with the results from this paper, one might be able to learn whether HomeBase affected non-participants.

6.2 Postponement and other inter-temporal effects

HomeBase has effects beyond the month in which service starts, and these effects may alter the interpretation of our results, particularly the service equations. It is helpful to begin with a precise understanding of inter-temporal linkages.

Let $V(t)$ denote the probability that a family receiving HomeBase services remains out of the shelters for t months after these services begin, $t = 1, 2, \dots$. Define $V(0) = 1$. Let $v(t) = V(t-1) - V(t)$ denote the probability that a family enters shelter in month t , conditional on receiving HomeBase services. Let $U(t)$ be the probability that this same family would remain out of shelter for t months if it did not receive HomeBase services—the counterfactual. Set $U(0) = 1$ and let $u(t) = U(t-1) - U(t)$ denote the corresponding probability of entering shelter for the first time in month t . Let $q(t) = u(t) - v(t)$, $t = 1, 2, \dots$

In the long run, the expected number of shelter entries that one HomeBase case averts is the limit of $U(T) - V(T)$ as T goes to infinity. Define

$$E(\infty) = \lim_{T \rightarrow \infty} [U(T) - V(T)] = \lim_{T \rightarrow \infty} \sum_{t=1}^T q(t)$$

This is what we would like to know. The naïve estimate of HomeBase effects is $q(1)$, shelter entries averted in the month service begins.

A priori, there is no reason to think that $E(\infty)$ is either bigger or smaller than $q(1)$. Some families whom HomeBase serves may enter shelters in later months; they might have entered in later months

without service. Some HomeBase families without HomeBase might have entered shelters in later months; because of HomeBase they never enter. The expression

$$E(\infty) - q(1) = \lim_{T \rightarrow \infty} \sum_{t=2}^T q(t)$$

can be either positive or negative.

Let $HB(t)$ denote the number of HomeBase cases begun in month t of calendar time. If HomeBase starts operating in month 0, then the net reduction in entries through month T due to HomeBase is

$$E^*(T) = \sum_{t=0}^{T-1} HB(t)q(T-t)$$

The most direct approach is to try to estimate the sequence $(q(t))$. Since the CT service equation is plausibly linear, we estimate shelter entries from a CT as a function of current HomeBase cases and lags of HomeBase cases, with lags extending up to six months. Table 5 shows the result. The sum of the contemporaneous effect and six lags results is substantially larger: -.30 compared with the contemporaneous estimate of -.12. We are not sure what to make of the large increase in the lag effects for the instrumented number of cases. However it surely argues against a postponement hypothesis.

We undertook a second analysis that is parallel to the one applied at the CD level to the geographical question. We re-estimated both the coverage and services model by now grouping observations into longer time periods, 2, 3 and 6 months. The estimates in table 8 show only modest declines in the HB effect when the duration of the observation period is lengthened up to six months,

Thus we cannot reject the hypothesis that the sum of $q(t)$ for t running from 2 to 6 is zero. This suggests that ignoring lags does not lead us seriously astray in the service equations. If the sum of $u(t)-v(t)$ for t running from 7 to infinity is zero, too, then $q(1) = E(\infty)$. This result is not “no postponement”; it is “no net postponement”, or “equal postponement regardless of HomeBase services.” With a quadratic term and less data, lags are not practical with CDs.

The interpretation of the coverage equations is simpler. The coverage effects without the “experienced” variable are unbiased estimates of the average value of $E^*(T)$ during the study period. With the experienced variable included, the estimates with $R=0$ give the average value of $E^*(1)$ and $E^*(2)$, and those with $R=1$ give the average value of $E^*(T)$, $T \geq 2$. Thus the coverage coefficients are estimates of the HomeBase effects that are uncontaminated by postponement problems; but they cannot be confidently extrapolated to the future without a more serious structural model like those in the service equations. However, the approximate equivalence of the estimated HomeBase effect in the coverage and services equations is also weak evidence for “no net postponement.”

A final weak piece of evidence for “small net postponement” comes from the study of residuals from the coverage equation for CDs. If short-term net postponement occurs, then months with unusually large numbers of shelter entries averted should be followed by months with unusually low numbers of entries averted, as some of the families served in the first month enter in the second. So the serial correlation between residuals should be more negative (less positive) when HomeBase is operating than when it is not operating.

To test for such an effect, we use our fullest model of entries as a function of HB coverage, equation (2). We divide the original sample into two subsamples. In one subsample of CD-months, HB services are “on”, and in the other they are “off.” We compute three sets of residuals from equation (2) estimated over our samples (the original, the “on” subsample, and the “off” subsample), and then

estimate a linear regression of each residual set on its one-period lag. Note that a t-test using the estimated coefficient on the lagged residual in any of these three equations is a test for residual autocorrelation at lag one (assuming all regressors are strictly exogenous). However, in our case, we are primarily interested in whether the estimable serial correlation is statistically significantly different between our two subsamples.

To answer this, we conduct a Chow Test. The Chow statistic is computed using the sum of squared residuals from each lagged equation, the number of observations in each subsample (N_1 and N_2), and the total number of parameters (k); in our case, the test statistic equals 5.66. The test statistic follows the F distribution with k and $N_1 + N_2 - 2k$ degrees of freedom, in general, and 3 and 3121 in our case. The null hypothesis for this test is that the regression fit is equal across our three proposed samples, or in other words that the serial correlation is the same for the full sample and both subsamples.

When HB is officially operating, serial correlation between residuals is 0.894 (95% C.I.=0.865, 0.923), and when HB is not officially operating, the serial correlation is 0.825 (95% C.I.=.800, .850). The critical value is 1.091 at the 5% significance level (with a p-value of 0.0007), so we reject the null hypothesis, concluding that there is strong evidence that the expected values in the two groups differ. The residuals are more-positively correlated when HB operates, not more negatively correlated. This is not consistent with the postponement story. But the effect, even if significant, is small.

Thus our coverage results are unbiased estimates of the average effect of HomeBase operation, and several pieces of evidence suggest that net postponement is either small or nonexistent

7.0 Exits and spell length

7.1 Theory

To find the full effect of HB on New York City's family shelter population, we must look at exits as well as entries. HB could affect exits in three different ways. Two of these ways would make spells longer, and one would make spells shorter.

Selection might make spells longer. If HB were more successful in averting homelessness for families with less serious problems than for families with more serious problems, and if homeless spells are longer for families with more serious problems, then HB will be more successful in averting spells that would have been short than in averting spells that would have been long. The average spell that starts when HB is operating would therefore be longer than the average spell that starts when HB is not operating.

Spillovers might also make spells longer. For instance, if HB participants stay in apartments they would otherwise vacate, fewer apartments will be available for shelter residents to move into. When HB is operating, then, spells may end less often.

Postponement could make spells shorter. Suppose some families leave shelters when an exogenous favorable event occurs—winning the lottery, finding a good job, getting married—and the hazard of the good event does not depend on whether the family is in a shelter or not, and rises over time. Then if HB delays shelter entry, it reduces the expected interval between shelter entry and the favorable event. The average shelter spell for families who entered after HB started would be shorter.

Notice that each of these three effects tells us to look at a somewhat different set of exit hazards. Postponement tells us to look at all days in spells that began after HB had been operating a few months. Selection tells us to look at those days, as well as all days in all the other spells that began when HB was operating, too. The spillover story, by contrast, tells us to look at days when HB was operating, not the spells in which they are embedded.

To keep the analysis simple, we will concentrate on selection and postponement, and not test directly for spillovers. To the extent that spillovers are geographically diffuse, moreover, our dataset may not let us test for them at all (we have no way of knowing whether a family that originally came from Brooklyn would not have left the shelter system and moved to the Bronx if HB had not been operating).

During the period we study, the majority (about two-thirds) of eligible families left the shelter system by receiving a subsidized apartment through DHS. This is called “placement.” Families’ circumstances and decisions have some bearing on when placement occurs, but DHS resources, rules, and queues are very important. We would not expect selection to have a large effect on the timing of placement, and postponement should have none, since time-in-shelter is what matters for DHS queues.

7.2 Methods

We use Cox proportional hazard methods. Because we want to consider placement and non-placement exits separately, we use a competing risk specification. The goal is to find whether families who entered the system when HB was operating had a smaller or larger hazard of non-placement exit than other families.

Specifically, let $\lambda_j(d, f, m, c)$ denote the hazard that in calendar month m , family f , which came from community district c , will leave the shelter system after d days, and the exit will be type j (placement or non-placement). Our basic equation is

$$\lambda_j(d, f, m, c) = \lambda_0^j(d) \exp\{\beta X_{fc} + \gamma_c + \delta_m\} + L[\tilde{\beta} X_{fc} + \tilde{\gamma}_c + \tilde{\delta}_m]\} + \varepsilon_{jfc dm}$$

(5)

Here $\lambda_0^j(d)$ is the baseline hazard for type j exits; the Cox method does not estimate this directly. The vector X_{fc} is a vector of characteristics of family f and community district c . In particular

$$\beta X_{fc} = \beta_1 A_f + \beta_2 K_f + \beta_3 \overline{H_{fc}} + \beta_4 \overline{P_{fc}} + \beta_5 \overline{R_{fc}}.$$

Here A_f is the (demeaned) number of adults in family f , K_f is the (demeaned) number of children in family f , $\overline{H_{fc}}$ is a dummy equal to one if and only if $H_{ct}=1$ for the month in which family f entered the shelter system, $\overline{P_{fc}}$ is a dummy equal to one if and only if $P_{ct}=1$ for the month in which family f entered the shelter system, and $\overline{R_{fc}}$ is a dummy equal to one if and only if $R_{ct}=1$ for the month in which family f entered the shelter system.

Continuing with (5), γ_c and δ_m are dummies for the CD and the current month respectively, and L is a dummy variable equal to one if and only if the exit is a placement. Thus the independent variables are fully interacted with placement type.

The coefficients we are most interested in are $\beta_3, \beta_4, \beta_5$. These coefficients indicate how the non-placement hazard changes for families who entered the shelter system when and where HB was operating. We are also interested in $(\beta_3 + \widetilde{\beta}_3), (\beta_4 + \widetilde{\beta}_4),$ and $(\beta_5 + \widetilde{\beta}_5)$. These sums indicate how the placement hazard changes for families who entered the shelter system when and where HB was operating.

7.3 Results

Figure 5 shows the baseline survival probabilities for placement and non-placement exits for families who entered when HB was officially operating and those who entered when it was not. (This is for a family with the mean number of adults and children.) HB appears to make little difference, especially for placements.

Table 9 provides the results from estimating equation (4). It confirms the general picture that HB makes no significant difference to exits. HB families seem to leave slightly faster, but the difference is tiny and not statistically significant. Family size has large and significant effects: larger families stay longer, though the effect is stronger for non-placements than for placements. In fact families with more children are placed sooner than families with fewer children; the effect is small but significant.

Selection and postponement may both be operating, but if they are, they are cancelling each other out. There is reason to suspect, however, that neither is operating. If operating in the hypothesized direction, postponement would increase the non-placement hazard rate for CDs that have been operating several months, relative to CDs that have been operating less time. So postponement implies that the effect of experience on non-placement exits should be positive and large. Selection has no implication for a change in effect after several months. So if both postponement and selection are strong and offsetting, the coefficient on experience for non-placements should be positive and significant. But it is not: the point estimate is .004 (95% C.I. = -.083, .091). Moreover, we have independent evidence in section 6.2 that net postponement is unlikely to be large.

Notice that neither selection nor postponement would affect the length of shelter spells in a Markovian world where current condition was a sufficient statistic for all predictions. See O'Flaherty (2012) for a discussion of this proposition. This result is weak support for a Markovian model of shelter transitions.

8.0 Conclusion

8.1 HB worked

Our best estimates are that HomeBase reduced shelter entries by between one and two for every ten families it served. It did not change homeless spell lengths. Our evidence suggests that, on

net, non-participants did not enter shelters more often as a result of HB, and that either “small net postponement” or “no net postponement” occurred. While we have been able to estimate net effects, especially historical net effects, we have not been able to estimate gross effects—particularly gross postponement. Experiments would be very helpful in estimating many of these gross effects.

Our findings lead us to believe that HomeBase effects are heterogeneous. We are reasonably confident that HB worked better in neighborhoods that historically generated many shelter entries than in neighborhoods that historically generated few. This is intuitive—you can’t avert shelter entries that would not have occurred anyway. Even in neighborhoods where potential shelter entries were plentiful, however, HB had decreasing returns to scale: it was less effective when it had to contend with many cases. Whether this deterioration is due to congestion or to selection is an open and important question.

We would like to know why HB effects were heterogeneous, but we cannot tell. Outcome variation might stem from variation in what service providers do, variation in the characteristics and motivations that HB participants bring to the offices, or variation in neighborhood and historical environments. Because HB operations have evolved and economic and housing conditions have dramatically changed since the start of HB operations, we are also cautious about extrapolating effect sizes estimated from the study period to today’s HB program.

8.2 Was HB cost effective?

Despite uncertainties about the true effect size, these findings suggest that HB compares favorably with other strategies to prevent homelessness and probably saved the city money. A reduction of even ten entries per hundred cases, the lower bound of our range of reasonable estimates of the HB effect, is large in the homelessness literature. Since homeless families on average stay in shelter for close to a year, these results imply a reduction in point-in-time (PIT) homelessness of at least

almost ten (and maybe as much as twenty) per hundred HB cases. In contrast, subsidized housing probably reduces PIT homelessness by about 3-7 per hundred households served. (See Ellen and O’Flaherty 2010 for a review of the literature and calculation.) Of course, these results may not apply beyond New York, which has a constitutional right-to-shelter, and a large family shelter system, in which families stay for long periods.

A rough comparison suggests that HB probably saved the city money. The budget cost of HB in this period was clearly below the estimated \$30-60 million in shelter costs that HB saved (this uses \$30,000 per shelter stay, which is the standard figure).

A more thorough cost-benefit analysis, of course, would consider the benefits to the participants themselves of both the direct assistance that HB provided (even if it did not avert homelessness) and the benefits both to participants and the rest of society to averting homelessness. HomeBase may have promoted residential stability, for instance, quite apart from any effect on homelessness, and a large body of evidence (Haveman et al. 1991, Astone and McLanahan 1994, Aaronson 2000, Mohanty and Raut 2009) suggests that residential instability hurts children’s cognitive development).

By analogy, people do not generally evaluate fire prevention on whether it saves municipalities money; they evaluate it on the value of lives saved and property preserved. We do not see why an evaluation of homelessness prevention should be limited to impacts on the city budget. (Shelters and police services are not evaluated on that basis.) The real unanswered question that is needed for a serious cost-benefit analysis of prevention is what the cost of a homeless spell is—to the homeless family, to third parties, and to the city government. Without such a number, cost-benefit analysis of homeless prevention is a feckless exercise.

8.3 Foreclosure

We believe our study is also the first to model the effect of the foreclosure process on homelessness. In New York's largely rental market, foreclosures have a strong association with shelter entries. The association between the foreclosure process and shelter entries is always statistically significant, but the point estimates typically range between 3 and 5 shelter entries over an 18-month period per 100 filings. Given the large number of LP filings during this period--115,000--these estimates suggest that 3,400 to 5,700 shelter entries during this period may be associated with the foreclosure process. Our data do not allow us to draw a strong causal inference. We cannot state with a high degree of confidence whether this association is a direct consequence of the foreclosure process or whether it is merely an indicator for broader economic factors that are tied to variation in housing stability between neighborhoods and over time. A better understanding of this connection is an important topic for future research.

8.4 Directions for future research on homelessness prevention

This study was designed to answer the question of whether HB worked during the early years of its operation. Left for future research are the equally important questions of the how and why it worked.

The first research question should be how and where prevention programs find potential clients, and what information those clients have? HB began with a community-based approach. Using data from DHS showing the spatial distribution of last known addresses of shelter applicants, staff at HB centers developed and implemented targeted outreach strategies to publicize services in high-shelter-use neighborhoods, and established relationships with a referral network of local service providers and government agencies that come into contact with families with housing emergencies. Although this study is based on HB cases found through community outreach, DHS has supplemented the community-

based approach to case finding with a “diversion approach:” HB staff recruit clients in New York’s central shelter intake center as families are applying for shelter services.

A second question is what families benefit most from HB? During its early years, DHS left considerable discretion to HB center staff in selecting eligible clients. During the first four years of operation, 9000 families were turned away for HB offices—40% of all applicants. Only a quarter of these families were rejected based on bright-line rules: high income(5%) or residence in an ineligible CD (20%). The leading reasons for denying eligibility involved staff judgment: the applicant’s housing problems were better served by other programs (37%), the applicant had an insufficient housing crisis(19%), and non-compliance with the application process (19%). As our findings show, HB case manager discretion apparently extended to waiving the residency requirement. We suspect that a careful study of the application procedure may also turn up evidence that there was some flexibility even when it came to the income criterion.

DHS is now introducing a structured screening instrument to better target HB services based upon results of Shinn and Greer’s 2011 study of risk factors for shelter entry among family applicants for HB. The theory is that by constraining case manager discretion in selecting eligible cases, a structured screening instrument will yield a higher concentration of families at high risk of shelter entry, and as a consequence will result in increased allocation of HB services targeted at families most likely to enter shelter. It is an empirical question as to whether substitution of structure for professional judgment results in more effective HB operations. Is there valuable soft information in the intake process that is unobservable to outside econometricians? How much of this soft information do potential participants have? How much do intake workers have? What are the incentives to reveal this information and to use it in decision-making?

Then there is the question about the sources of effect heterogeneity. An important source of variation is the mix and delivery of the basket of HB prevention services. A distinctive feature of the HB program is the discretion HB case managers have in selecting from a menu of prevention services that are matched to the characteristics of housing emergency confronting each client. We had no information on the mix of services received by clients or the nature of housing emergencies and only minimal characteristics of family characteristics. These are all pieces of information that would be required for a refined analysis designed to explain variation of an HB effect between HB centers and between neighborhoods served by each center. Is the soft information from intake used to make decisions about what services to provide? How do the service decisions affect the incentives to reveal information? What are the incentives of service providers?

A prominent theme running through evolving HB policies and procedures is an effort to limit or imposing structure on the broad discretion given to case managers and center in deciding to whom and how to delivery HB services. These changes raise interesting questions about the optimal balance between structure and flexibility—between hard information and soft-- in prevention programs.

To answer the above questions, will require a combination of observation studies that peer into the black box and describe how HB and other community prevention programs go about outreach and case finding, determining client eligibility, and matching services to client needs. More detailed description of how community prevention programs work may then provide the basis for more refined experiments and quasi experiments to better understand why they work.

Besides studying variation in program structure and operations, future research should investigate how prevention programs operate in different neighborhoods, different historical contexts, and different milieus of local housing and homeless policies. How easy it is to resolve housing

emergencies may depend on neighborhood characteristics, average income, ethnicity, and the quality and price of the housing.

Finally local housing and homeless policy may matter in determining prevention program effectiveness. Homelessness can be prevented only among people who would otherwise be homeless, and local housing policies and the structure of homeless programs together exert considerable influence on who potentially homeless people are. In the limit, homelessness prevention would not “work” at all in a city with no shelters (and Draconian street policies). More expensive and attractive shelters may make homelessness prevention a more attractive strategy from the point of view of a municipal budget, but a less attractive strategy from the point of view of family well-being.

Community-based prevention programs may prove to be an important policy tool in reducing family homelessness. This study is a first step in a larger program of research on how they work in New York City and how they might be adopted elsewhere.

References

Aaronson, D., 2000, "A note on the benefits of homeownership," *Journal of Urban Economics* 47(3): 356-69.

Apicello, Jocelyn, 2008, "Applying the population and high-risk framework to preventing homelessness: A review of the literature and practice," working paper, Columbia Center on Homelessness Prevention Studies.

---, William McAllister, and Brendan O'Flaherty, 2012, "Homelessness: Prevention," article 00352
In: Susan J. Smith, MarjaElsinga, Lorna Fox O'Mahony, OngSeowEng, Susan Wachter, editors.
International Encyclopedia of Housing and Home. Oxford: Elsevier.

Astone, N.M., and S.D. McLanahan, 1994, "Family structure, residential mobility, and school dropout: A research note," *Demography* 31(4): 575-84.

Buckley, Cara, 2010, "To test program, some are denied aid," *New York Times*, December 8,
accessed at <http://www.nytimes.com/2010/12/09/nyregion/09placebo.html?pagewanted=all>.

Burt, Martha, C. Pearson, and A.E. Montgomery, 2005, *Strategies for Preventing Homelessness*.
Washington, DC: U.S. Department of Housing and Urban Development.

Cragg, Michael, and Brendan O'Flaherty, 1999, "Do homeless shelter conditions determine shelter population? The case of the Dinkins deluge," *Journal of Urban Economics* 46: 377-415.

Ellen, Ingrid Gould and Brendan O'Flaherty, 2010, Introduction, in Ellen and O'Flaherty, eds.,
How to House the Homeless. New York: Russell Sage.

Furman Center for Real Estate and Urban Policy, 2010, Foreclosed Properties in New York City: A
Look at the Last 15 Years. Accessed on January 21. 2011 at

http://furmancenter.org/files/publications/Furman_Center_Fact_Sheet_on_REO_Properties.pdf.

Haveman, Robert, B. Wolfe, and J. Spaulding, 1991, "Childhood events and circumstances
influencing high school completion," *Demography* 28(1): 133-57.

Mohanty, Lisa L., and Lakshmi K. Raut, 2009, "Home ownership and school outcomes of
children: Evidence from the PSID Child Development Supplement," *American Journal of Economics and
Sociology* 68(2): 465-89.

O'Flaherty, Brendan, 2012, "Individual homelessness: Entries, exits, and policy," *Journal of
Housing Economics* 21(2): 77-100.

--- and Ting Wu, 2006, "Fewer subsidized exits and a recession: How New York City's family
homeless shelter population became immense," *Journal of Housing Economics* 15 (2): 99-125.

Table 1: Summary Statistics (1/2003 to 11/2008)

	CD-Level(N=59)	CT-Level (N=1,889)
Total Monthly Observations	4,189(100%)	134,119(100%)
CD/CT-Months of Formal HB Operations	1,063(25.4%)	35,117(26.2%)
CD/CT-Months of Informal HB Operations	1,166(27.8%)	37,316(27.8%)
CD/CT-Months of Experienced HB Operations	946(22.6%)	31,339(23.4%)
Total Shelter Entries		45,088
Average Shelter Entries/Monthly Observation	10.7	.336
Total HB Cases Opened		10,978
Average HB cases Opened /Monthly Observation ^a	3.8	.118
Total Units in Foreclosure proceedings		115,320
Average Foreclosures started/Monthly Observation	27.53(45.37)	.86(4.88)
Average Distance to Closest HB Center (in Miles)		3.05(1.91)

^aAveraged over period of HB operations starting in November, 2004

Table 2: Annual Trends in Shelter Entries and HomeBase Cases

	2003	2004	2005	2006	2007	2008 ^a
All CDs						
Shelter Entries	6,459	6,726	5,951	8,116	8,190	9,646
HB Cases Opened	^b	194 ^c	2,440	3,042	2,111	3,191
(Shelter Entries)-to-(HB Cases) ratio ^d		1.32	.60	.65	2.09 ^e	3.02
Units Starting Foreclosure Process	16,163	13,298	12,798	17,763	27,421	27,877
Restricted to "Big Six" CDs						
Shelter Entries	1,499	1,589	1,382	1,896	1,969	2,325
HB Cases Opened		180	2,312	2,897	1,404	822
(Shelter Entries)-to-(HB Cases) ratio		1.32	.60	.65	1.40 ^f	2.83
Units Starting Foreclosure Process	4,998	2,934	3,208	4,436	6,717	5,924

^aStudy Period ends November, 2008

^b HB not Operating

^c Limited HB operations begin in "Big Six" CDs in November 2004

^d Restricted to CDs during official HB operations

^e For first six month of 2007 shelter-HB ratio=.86; after expansion of HB program starting in July 2007 Shelter-to-HB ratio=3.73

^fFor first half of 2007 shelter-HB ratio=.86; second half of 2007 shelter-HB-ratio=2.80

Table 3: Effects of HB Coverage on Monthly Shelter Entries, CD Level Results (OLS Estimates)

	All CD Linear Model	Low-Use CD Linear Model	High-Use CD Linear Model	Loglinear Model
HB Coverage Variables				
Official Operations	-2.5** (-3.67, -1.32)	-0.198 (-1.05, .65)	-2.10† (-4.40, .19)	-0.06 (-.16, .04)
Unofficial Operations	-2.17** (-2.77, -1.57)	-0.319† (-.70, .06)	-0.93 (-2.54, .68)	-0.055* (-.11, -.00)
Experienced Operations	0.07 (-1.09, 1.24)	-0.193 (-1.04, .65)	0.62 (-1.65, 2.88)	-0.012 (-.11, .09)
Sum of HB Coverage Coefficients	-4.59** (-5.61, -3.58)	-0.71† (-1.54, .12)	-2.42* (-4.52, -.31)	-0.126** (-.21, -.04)
Sum of foreclosure coefficients	.053** (.040, .066)	0.038** (.024, .052)	.027** (.007, .047)	.001* (.000, .003)
Within R ²	0.34	0.21	0.59	0.22
<i>N</i>	3127	2014	1113	3127

All regressions include month and CD fixed effects.

†p<.1

*p<.05

**p<.01

Table 4: Effects of HB Coverage on Monthly Shelter Entries, CT Level Results(Point Estimates & 95% Confidence Intervals)

	All CT Linear Model	Low Use CT Linear Model	Moderate Use CT Linear Model	High Use CT Linear Model	Poisson Model ^a
HB Coverage Variables					
Official Operations	-0.056** (-0.071,-0.007)	-0.018* (-0.035,-0.002)	-0.019 (-0.068,0.029)	-0.009 (-0.176,0.158)	0.948 (0.881,1.018)
Unofficial Operations	-0.071** (-0.085,-0.056)	0.007† (-0.001,0.015)	-0.034* (-0.062,-0.006)	-.120† (-0.244,0.004)	0.973 (0.929,1.019)
Experienced Operations	-0.003 (-0.030,0.25)	0.002 (-0.015,0.019)	0.002 (-0.046,0.050)	0.015 (-0.150,0.018)	1.03 (0.958,1.106)
Sum of HB Coverage Coefficients	-0.130** (-0.153,-.106)	-0.010 (-0.025,0.006)	-0.051* (-.093,-.010)	-0.114 (-0.270,0.043)	0.949† (0.897,1.095)
Sum of Foreclosure Coefficients	0.038** (0.030,0.046)	0.013** (0.005,0.021)	0.026** (0.014,0.038)	-0.004 (-0.040,0.032)	
Within R ²	.026	0.007	0.035	0.095	
N	100,117	52,735	37,206	10,176	131,279

All regressions include month and CT fixed effects.

^a Coefficients for Poisson model are incidence rate ratios

†p<.1

*p<.05

**p<.01

Table 5: Effect of HB Services on Shelter Entries: Instrumental Variable Regressions

	CD Level Models		CT Level Service Model	CT Level, 6-month lag Model
HB families served	-0.103*	-0.576*	-0.12**	-0.30**
	(-.14, -.06)	(-.80, -.35)	(-0.15, -0.09)	(-0.49, -0.039)
(HB families served) ²	---	0.00991*		
	---	(.005, .014)		
Sum of Foreclosure Coefficients	0651**	.0820***	0.045**	0.060**
	(0.050, 0.080)	(0.064, .010)	(0.037, 0.053)	(0.050, 0.071)
R ²	0.86	0.85	0.005	0.046
N	2725	2725	100,117	88,783

Instruments include distance to nearest HB center, official operations, Initial Big Six contract period, fiscal year
 Other covariates include month and CT/CD fixed effects and foreclosure variables

†p<.1

*p<.05

**p<.01

Table 6: Estimates of the Historical Effect of HomeBase on Shelter Entries Averted per 100 HB Cases

	Point	Lower	Upper	Point	Lower	Upper
Coverage						
Linear	60.5	47.8	73.2	65.4	54.6	75.2
Stratified	19.2	4.5	34.2	20.8	7.2	34.4
Log/Poisson	25.1	3.5	46.8	11.8	-1.7	25.4
Service						
Linear	10.3	6.0	14.0	12.0	9.0	15.0
Quadratic	26.1	16.9	35.2	-	-	-

Note: Point estimates are calculated by estimating the number of entries averted for each equation during unofficial and official HB operations adjusted for the experience factor. The total is divided by the total number of HB cases opened and multiplying the ratio by 100. 95% confidence intervals are calculated using the Statalincom procedure

Table 7: Effect of HB Estimated with Larger Units, Linear Specification

	Single month
Capacity (OLS)	
Single CD	-4.62* (-5.57, -3.67)
Double CD	-5.6* (-6.77, -4.42)

All regressions have month and CD fixed effects and foreclosures. The capacity effect is the sum of the coefficients on unofficial and official operation.

*Significant at 1 percent level.

Table 8: HB Coverage and Service Effects at CD level for 1-, 2-, 3-, and 6-month Grouping of Observations

	One month	Two month	Three month	Six month
Capacity Equation Poisson	0.952 (.899,1.008)	0.952 (.898,1.010)	0.954 (.899,1.012)	0.949 (.891,1.012)
Services equation	-0.21 (-.24,-.17)	-0.18 (-.21,-.14)	-0.19 (-.23,-.15)	-0.18 (-.21,-.14)

Table 9: Effects of HomeBase Coverage on Exit Rates, Cox Proportional Hazard Model

	Non-placement effect β	β	Placement effect $(\beta + \beta)$
Official operation	0.024 (-.067, .116)	-0.057 (-.202, .088)	-0.033
Unofficial operation	-0.005 (-.093, .083)	0.024 (-.076, .123)	0.019
Experienced	0.003 (-.084, .091)	0.018 (-.124, .160)	0.022
Sum	0.023	-0.015	0.008
Log pseudolikelihood			-426,316.65
<i>N</i>			90,222

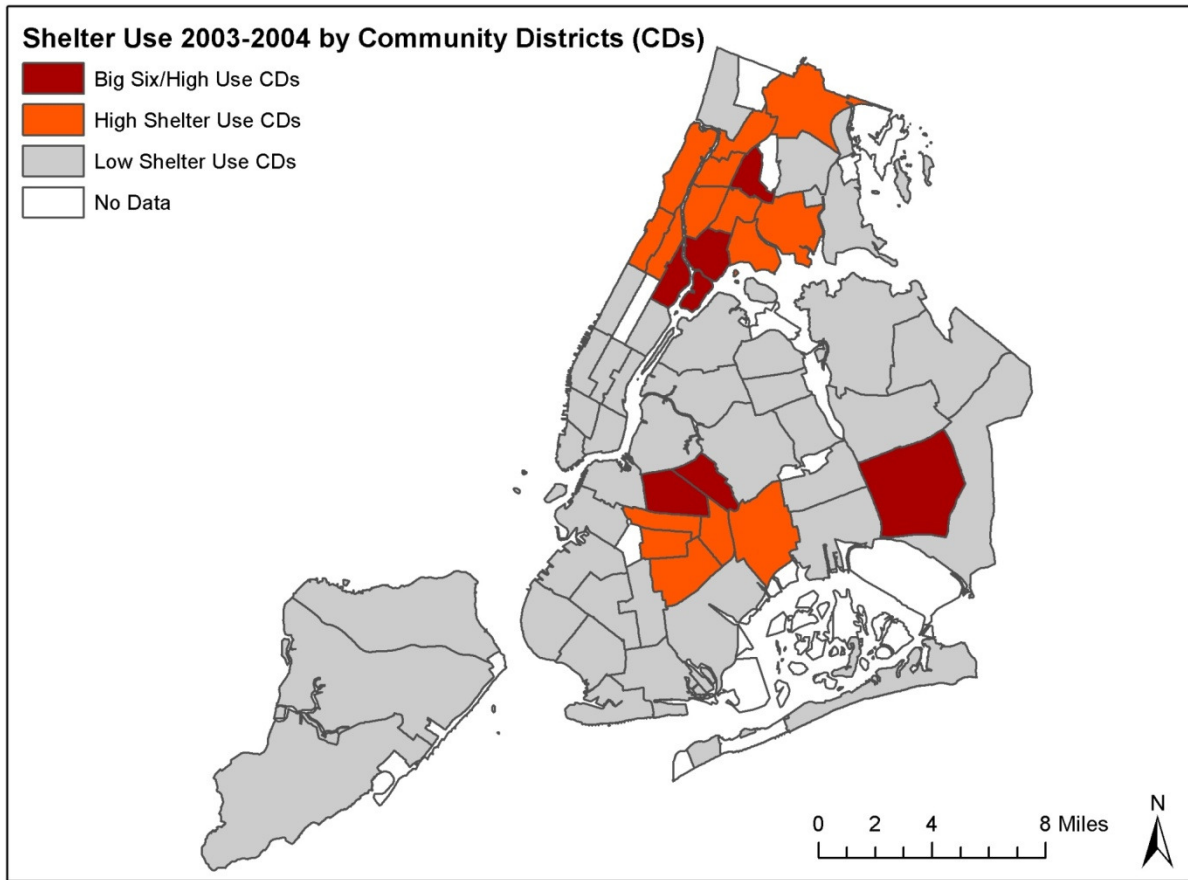


Figure 1: High and low shelter use NYC Community Districts, 2003-2004

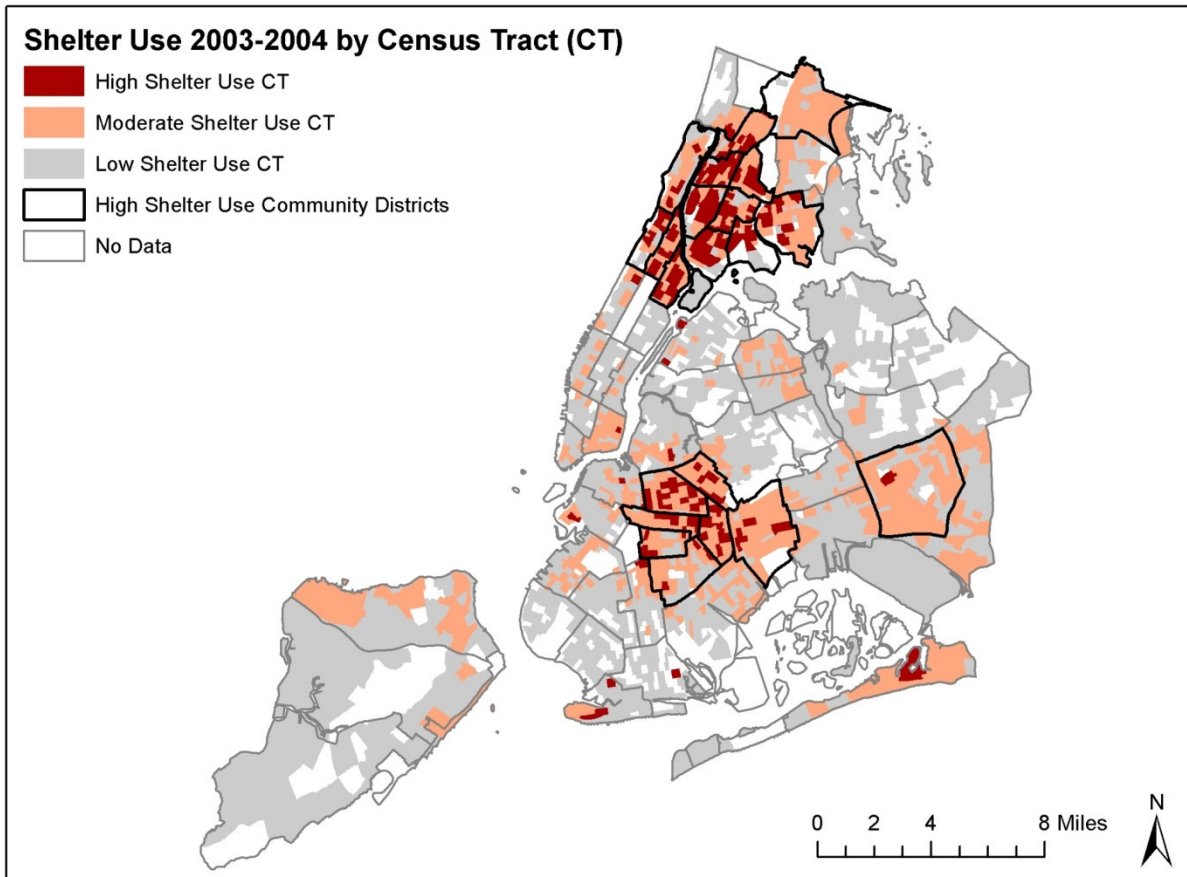


Figure 2: Low, moderate and high use NYC Census Tracts, 2003-2004

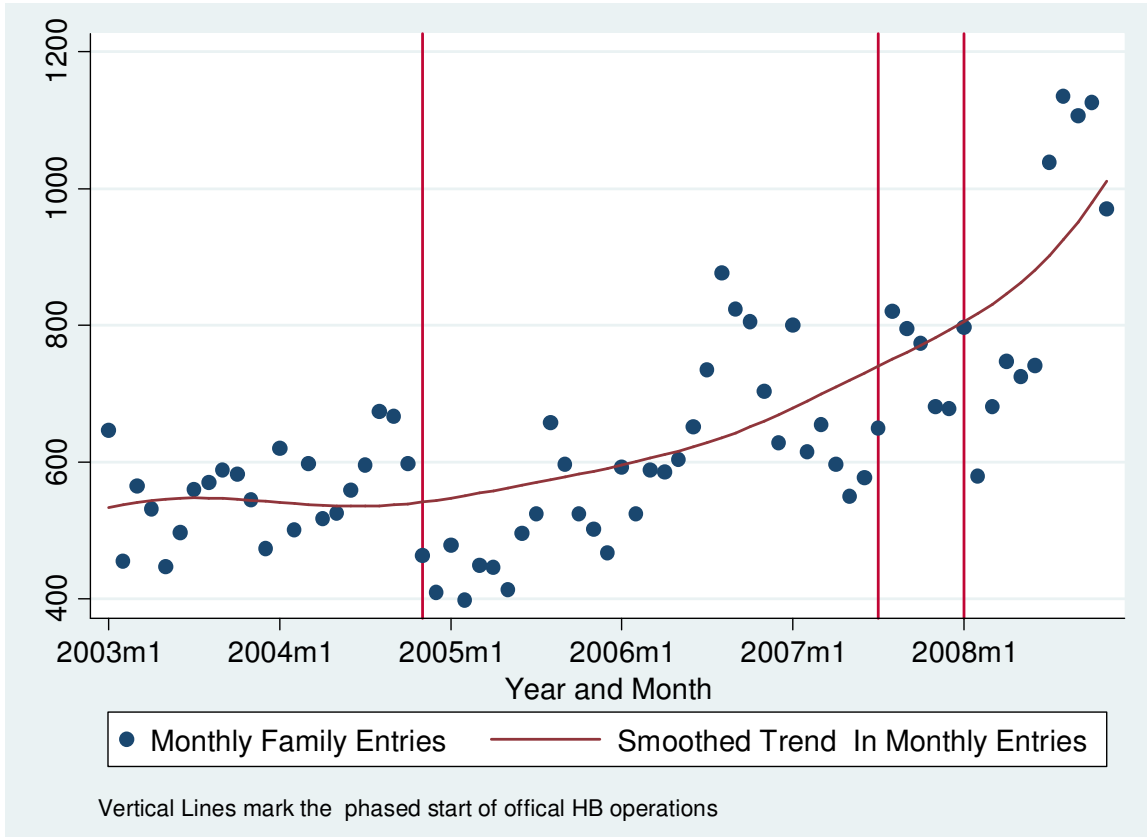



Figure 3: Monthly family entries into the New York City shelter system (2003-2008)



Cov.=Coverage Model
 Stat.=Stratified Model

Figure 4: Estimates of the historical effect of HomeBase on shelter entries averted per 100 HB cases. Markers indicate point estimates and vertical line are 95% confidence intervals.(See Table 6 for values of upper and lower bounds)

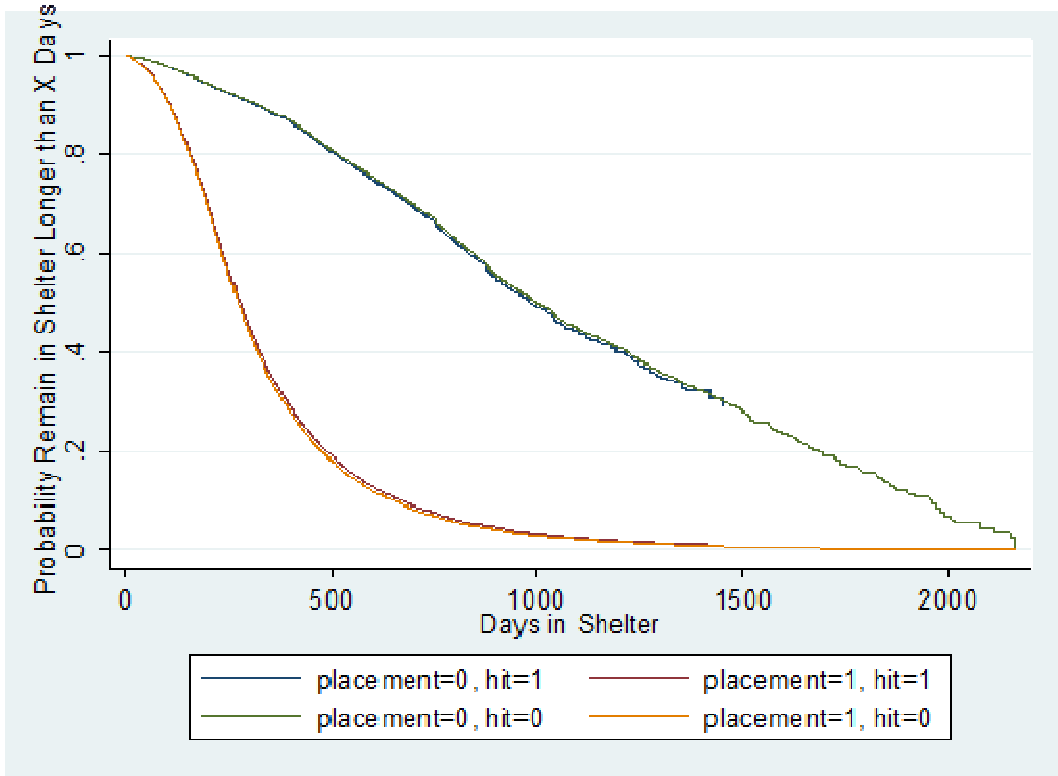


Figure 5: Survival curves for shelter duration. Comparison between shelter spells that ended and don't end in a housing placement(placement=1), and residents in CDs that were (hit=1) and were not eligible(hit=0) for HomeBases services at time of shelter entry.