



Columbia University

*Department of Economics
Discussion Paper Series*

Managing Homeless Shelters

Brendan O'Flaherty

Discussion Paper No.: 0809-15

*Department of Economics
Columbia University
New York, NY 10027*

April 2009

Managing Homeless Shelters

Brendan O'Flaherty

Department of Economics

Columbia University

New York, New York 10027

This draft, April 26, 2009

Abstract

This is a formal analysis of how homeless shelters should operate: in particular, what quality of accommodations they should provide and how they should help their residents in securing conventional housing. I also examine timing. The results extend to cover optimal police response to street homelessness as well. I draw heavily on the unemployment insurance literature.

Acknowledgments

I am grateful to Jay Bainbridge, Sewin Chan, Lance Freeman, and participants at the Columbia-NYU conference on How to House the Homeless for helpful comments; to Adonis Antoniades for excellent research assistance; and the Russell Sage Foundation for financial support. Support from the Columbia Center on Homelessness Prevention Studies is gratefully acknowledged. Many of the ideas grew out of conversations with Dennis Culhane, Robert Hess, and Maryanne Schretzman. I am solely responsible for errors and opinions.

Keywords: homelessness, homeless shelters, street homelessness, unemployment insurance.

JEL code: I39.

Homeless shelters are one of the most crucial parts of the U.S. safety net. Although the number of people who stay in homeless shelters is fairly small—about 391,000 on a night in January 2007 (US HUD 2008)—those who use shelters are among the worst-off people in the country. Many shelters have minimal eligibility requirements, and so are essentially housing of last resort. When all else fails, shelters are the place where when you go there they have to take you in.

Hence, how shelters operate—what quality of accommodations they offer, for instance, what help they give to people who leave—is important for the functioning of the U.S. safety net. Yet shelter operation has received very little formal analysis, even though both moral hazard (nicer shelters draw more clients, and they stay longer) and adverse selection (residents know more about their out-of-shelter opportunities than administrators do) are present. The purpose of this paper is to begin to address this lack of attention.

Current macroeconomic difficulties have renewed interest in homelessness and spurred governments to do more in this area. The first Obama stimulus package (the American Recovery and Re-investment Act of 2009), for instance, included \$1.5 billion for homelessness prevention and rapid rehousing. This paper seeks in part to illuminate how those funds would best be spent.

I ask several questions: what is the optimal quality of accommodations for a shelter to provide? How should shelters allocate housing subsidies among those who depart? When should subsidies be given? These are basic questions that any shelter must answer, but they have not been studied rigorously. As far as I am aware, no other paper tries to answer the basic question what optimal shelter quality is. The answers are the following.

Optimal shelter quality depends on a tradeoff between insurance and moral hazard (absent moral hazard shelters would be the same quality of housing that residents enjoy during

good times). At the optimum, shelter quality maximizes the reservation housing opportunity—the worst unit (rent and everything else considered) that would just draw an individual from shelter, when shelters are operated on a balanced budget with taxes on conventional housing. This result is a translation of Shimer and Werning (2006) for optimal unemployment insurance. Shelter quality should also deteriorate as a spell lengthens.

Optimal subsidies for leaving shelters are given to everyone who leaves, no matter when they leave. Any sum for subsidies should be spread as widely and as evenly as possible.

If the amount of individual subsidies is fixed and so the sum cannot be spread evenly, then in a homogeneous population, any subsidies that are given should be given immediately on shelter entry. There is no reason to reward people who stay in shelters longer. But if the population is heterogeneous and has private information, some placements may be delayed. O’Flaherty (2009) examines this specific question, but with less generality and more constraints on the instruments available to shelter operators.

The next section of the paper describes in general terms the analogy between shelters and unemployment insurance. Section 2 derives the conditions for optimal shelter quality when quality is not allowed to vary during a shelter spell. Section 3 shows that when quality can vary within a shelter spell, declining quality is usually optimal. Section 4 turns to the question of placement into subsidized housing when placement probability cannot vary with spell length; section 5 discusses this question when placement probability can vary by spell length. Section 6 applies some of the results in the previous sections to street homelessness.

1. Shelters as social insurance

Shelters are like unemployment insurance. People get unemployment insurance when they lose their jobs, and they get a bed in a shelter when they lose their houses. People continue to receive unemployment insurance until they find a new job or reach an administrative limit; people stay in shelters until they find a new place to live or reach an administrative limit.

But there are important differences. First, shelters are in-kind while unemployment benefits are cash. Moreover, living in a shelter precludes living anywhere else (by the laws of physics); receiving cash from unemployment insurance does not preclude receiving cash from other sources (including concealed employment). Thus people who try to enter shelters self-select more than people who try to collect unemployment insurance. Shelters, in this regard, are more like an employer-of-last-resort than unemployment insurance benefits.

Those who design shelter programs thus have one more instrument than those who design unemployment insurance programs: they can set a level of quality for the shelter as well as a cash transfer between the recipient and the shelter authority, and a duration. (Both shelters and unemployment insurance programs can be augmented with additional long-term benefits: subsidized housing in the case of shelters and job training or job placement assistance in the case of unemployment insurance. I postpone until section 4 consideration of these additional benefits.

Shelters also differ from unemployment insurance because the conditions that begin and end receipt are not so clear-cut. Unemployment spells begin with a job loss, and most models of unemployment insurance take this as an exogenous random event, out of the worker's control and perfectly observable. The beginning of homelessness spells is not so clear-cut, especially for families who are doubled-up; nor is it so easy to consider it exogenous. Similarly, unemployment spells end when the worker gets a new job. This is usually modeled as the result

of a good-enough draw from a sequence of job offers, where the worker decides how good is good-enough and sometimes can affect the sequence of draws.

By contrast, researchers generally do not portray homeless people as searchers for apartments and homeless spells as frictions that result from failures to coordinate moves. Rather, the feeling has been that homeless people have more severe problems than not having found the right apartment, and that shelters are not just convenient places to stay while house-hunting.

There are three reasons why homelessness researchers have something to learn from unemployment insurance research. First, very little is known about short homelessness spells, and many of them may in fact be frictional. Second, for many purposes it does not matter whether moral hazard operates on the end of spells or on the beginning of spells. Homelessness research has paid a lot of attention to moral hazard at the beginning of spells—do shelter conditions draw people into shelters?—even though empirical work (O’Flaherty and Wu, 2006, 2008; Culhane et al. 2007) indicates that moral hazard operates at the end of spells, too. Third, formally, it does not matter whether the things that arrive are wage offers, rent offers, or something more complex—complete packages that represent opportunities for living outside the shelter. The packages would include, *inter alia*, apartment rents, apartment qualities, apartment locations, companionship, jobs and other income sources, opportunities to consume non-housing goods, professional support and medical help—in short, a complete description of the best available option that day for living outside the shelter. All that we require in order to use the unemployment insurance framework is that these outside options vary stochastically from day to day.

2. Optimal shelter quality

Let me illustrate these ideas by applying them to the problem of optimal shelter conditions, assuming that shelter conditions are constant during a family's stay. I will use Shimer and Werning's (2007) analysis of optimal unemployment benefits, and simply translate it into shelter terms.

In the absence of moral hazard and adverse selection problems, optimal shelters would replicate the quality of housing that their residents would consume in better times. Rent would be set at a level that would allow their residents to consume the same bundle of non-housing commodities that they consumed in good times. (Intertemporal optimization requires that the marginal utility of income in good times equal the marginal utility of income in bad times. Intra-temporal optimization requires that the marginal utility of housing, per dollar, equal the marginal utility of non-housing, per dollar, in both states of the world, and that the marginal utility, per dollar, of both goods equal the marginal utility of income. Thus the consumption of housing must be the same in both states of the world, and so too must the consumption of non-housing.)

Optimal shelter quality is less than this because of moral hazard and adverse selection. The unemployment insurance literature has many insights into how to combine these ideas.

Preliminaries

Ever since Shavell and Weiss (1979), the unemployment insurance literature has made a distinction between recipients who live hand-to-mouth and recipients who have access to well-functioning capital markets. For the latter, unemployment benefits serve only an insurance function—they compensate the recipient when something bad happens to her. For the former, unemployment benefits represent liquidity as well as insurance—they allow the recipient to smooth consumption partially over her lifetime in a way that she would do if she could borrow

and save at market interest rates. O'Flaherty (2008) looked at the capital market constraints that people at risk of becoming homeless faced, and concluded that they had some opportunities to borrow and to save, but not unlimited ones. In this paper I will concentrate on models of hand-to-mouth recipients, because people who become homeless may often have exhausted all their opportunities to borrow.

Following Shimer and Werning (2007), I consider a fiscally balanced shelter system. Housed people are assessed a tax or insurance premium τ per unit time. This revenue must cover the cost of the shelter.

Since shelters, unlike unemployment benefits, are not financed by premiums, what does it mean to think about a balanced system? Consider a shelter funded only by a distant magnanimous donor who contributes an arbitrary amount of money. The amount of that donation implicitly determines shelter quality.

Suppose the donation is small and the quality of the shelter is poor. The quality is so poor that the shelter's users and potential users would be willing to set up a self-funded insurance pool to improve the shelter's quality. Since the voluntary establishment of this insurance pool would be a Pareto improvement, the original shelter quality could not have been optimal.

Similarly, suppose the original donation had been very large, and so the shelter's quality is lavish. Then the shelter's users could take the donation, split it among themselves, set up a self-funded insurance pool for a shelter with lower quality, and pocket the rest of the money. If they would be better off with the lower quality shelter and more money in their pockets, then the original lavish shelter could not have been optimal.

Thus a shelter, no matter how it is funded, can be of optimal quality only if it is optimal for some self-funded insurance pool. Thus we are deriving necessary, not sufficient, conditions

for a shelter's quality to optimal. (Since different initial wealth conditions may produce different optimal shelters, optimal shelter quality is not unique. However, given any initial wealth distribution and a balanced budget, we will derive conditions for optimal shelter quality to be unique.)

The set-up

People who have an apartment keep it for T amount of time and then become homeless again. The period T is exogenous so we can concentrate on moral hazard only at the end of homeless spells.

In each period and in each state-of-nature, people consume a quality of housing h and a quantity x of the numeraire good. Sub-utility from this consumption is Cobb-Douglas

$$C = x^\beta h^{1-\beta}, \beta \in [0,1].$$

People maximize the expected present value of utility from consumption

$$E \int_0^\infty e^{-rt} U(C(t)) dt = E \int_0^\infty e^{-rt} (-e^{-\gamma C(t)}) dt$$

where $r > 0$ is the discount rate (both market and subjective), and $\gamma > 0$ is the coefficient of absolute risk aversion. Thus we assume constant absolute risk aversion (CARA) utility in consumption, which is a Cobb-Douglas function of housing and other goods.

A homeless person stays in a shelter of quality S and receives net cash transfers that allow consumption X of other goods, all controlled by the homeless shelter agency. This costs the shelter agency

$$X + pS$$

per day per person. Let b denote the amount the shelter agency spends per person-day. Assume the agency maximizes the utility of residents subject to the break-even constraint. Then

$$X = \beta b$$

$$pS = (1 - \beta)b$$

$$C = \delta(p)b$$

where

$$\delta(p) = \frac{\beta^\beta (1 - \beta)^{1-\beta}}{p^{1-\beta}}$$

and $\delta'(p) < 0$.

Opportunities for living outside the shelter have a Poisson arrival rate λ . To be concrete, an opportunity is a 3-tuple $(\tilde{p}, \tilde{s}, \tilde{y})$: an opportunity to live in an apartment of quality \tilde{s} at rent \tilde{p} per period and have an income \tilde{y} per period. Let \tilde{w} denote the per-period sub-utility from such a 3-tuple:

$$\delta(p)(\tilde{w} - \tau) = (\tilde{y} - \tilde{p}\tilde{s})^\beta \tilde{s}^{1-\beta}$$

(Obviously we could use an arbitrarily higher dimensional description of each opportunity.) The income-equivalent \tilde{w} are drawn independently from a cumulative distribution function F . A homeless person who receives an outside opportunity decides whether to accept it or reject it. If she accepts it, she leaves the shelter immediately and lives in the new apartment for T periods. If she rejects it, she stays homeless. Recall is impossible (and with CARA utility, never optimal).

Let \bar{w} denote a homeless person's "reservation utility": the minimum value of outside opportunity that will induce her to leave the shelter. Let

$$\alpha = \int_0^T e^{-rs} ds$$

denote the present value of receiving a dollar for the next T periods. Let

$$B = b + \tau$$

denote the difference between the agency's contribution to a homeless person and to a housed person.

Shimer and Werning derive (proposition 2) an equation for the reservation wage.

Translating this into our setting implies

$$(3)U(\delta\bar{w}) = U(\delta B) + \alpha\lambda \int_{\bar{w}}^{\infty} (U(\delta w) - U(\delta\bar{w})) dF(w).$$

Since the left-hand side of (3) is an increasing function of the reservation opportunity and the right-hand side is a decreasing function, comparative statics are simple.

Raising B —which means either increasing shelter quality or increasing the tax or both—shifts the right-hand side up, and so raises the reservation opportunity; people stay in shelters longer. Raising T , the duration of housing, also shifts the right-hand side up: since they will stay in the new house longer, search is more valuable, and so people stay in shelters longer.

On the other hand, raising $\delta(p)$ (by reducing p) reduces the reservation opportunity. Since δ is always multiplied by γ , the coefficient of absolute risk aversion, in (3), raising δ is the same as making searchers more risk averse. They become more willing to accept the opportunity they have before them rather than reject it and take a chance on something better. So people leave shelters sooner.

We can also examine the effect of changes in opportunities. An improvement in opportunities (a shift to a new more attractive distribution of opportunities) moves the right-hand side of (3) up, and so raises the reservation opportunity. Whether this results in longer shelter stays, however, is impossible to say a priori: people wait for better opportunities, which slows exits, but better opportunities come more often, which speeds exits.

We can demonstrate three results that always hold in this framework, no matter what the value of $B \geq 0$ is.

First, the reservation opportunity is always strictly positive and so the amount of homelessness is always strictly positive. No matter how bad shelters are, someone will always stay there.

To see this, consider $\bar{w} = 0$. By direct substitution, $U(\delta 0) = -1$. Similarly $B \geq 0$ and so $U(\delta B) \geq -1$. Finally

$$\alpha\lambda \int_0^\infty (U(\delta w) - U(0)) dF(w) > 0.$$

Thus the right-hand side of (3) is greater than the left-hand side at $\bar{w} = 0$. Hence the reservation opportunity cannot be zero and so it must be positive.

Second, there is always a reservation opportunity. Let $\bar{w} \uparrow \infty$. Then the left-hand side of (3) approaches zero and

$$\alpha\lambda \int_{\bar{w}}^\infty (U(\delta w) - U(\delta \bar{w})) dF(w)$$

approaches zero, too. Since $U(\delta B)$ is fixed and negative, the right-hand side is smaller than the left-hand side. By continuity, for some $\bar{w} \in [0, \infty]$, the left- and right-hand sides are equal.

Third, consumption is not smoothed completely. Sub-utility while homeless is less than sub-utility while housed. Since

$$U(B) = U(b + \tau) > U(b),$$

and from (3),

$$U(\bar{w} - \tau) > U(B),$$

we have

$$U(\bar{w} - \tau) > U(b).$$

Since the actual utility while housed is at least as great as $U(\bar{w} - \tau)$, utility while housed is at least as great as utility $U(b)$, utility while homeless, and greater with probability one.

Optimality

The homeless shelter agency maximizes a homeless person's expected utility (results would not change if the agency valued housed people's expected utility, too, since homeless people contemplate expected utility while housed, and vice versa.) The shelter agency chooses b and τ but its choice is constrained by the break-even constraint:

$$(4) Db = \alpha\tau$$

where

$$D = \frac{1}{\lambda(1 - F(\bar{w}))}$$

is the duration of the average homeless spell.

From the construction of the reservation opportunity, the lifetime expected utility of a homeless person is

$$U(\delta(\bar{w} - \tau))/r.$$

Thus the optimal homeless shelter plan has b and τ that maximize the post-tax income-equivalent of the reservation opportunity $(\bar{w} - \tau)$.

The question for optimality therefore is whether a balanced budget improvement in homeless shelter conditions will increase or decrease the post-tax income-equivalent reservation opportunity $(\bar{w} - \tau)$. If the improvement increases the reservation opportunity, the original shelter conditions were worse than optimal; if it decreases the reservation opportunity, the original conditions were better than optimal.

Every increase in benefits raises the reservation opportunity \bar{w} . But it also requires a higher tax rate τ to pay for it. Shelter operators should make only those improvements in shelter quality that raise the income-equivalent reservation opportunity more than they raise the tax rate.

Comparative statics are fairly straightforward. An increase in the duration of housing T or the arrival rate of offers λ increases the proportion of time that the average person pays taxes. Holding the tax rate constant, this increases the revenue available to improve shelter quality. Cities where spells of homelessness are shorter or intervals between spells longer should have higher quality shelters.

Improvements in shelter technology that raise utility for given expenditure should also raise both shelter quality and tax rates, since they raise the marginal rate at which reservation opportunity increases as a function of expenditure and revenue.

Shelter operators can calculate an implicit tax rate from accounting and flow data, but they do not observe the income-equivalent reservation opportunity \bar{w} . The shelter exit rate is observable, however, and is $\lambda[1 - F(\bar{w})]$. For very small changes in reservation opportunity, any continuous distribution can be approximated by a uniform distribution. A small improvement in shelter quality that causes a big decrease in the shelter exit rate for a very small tax rate increase should probably be adopted; and conversely. Finer estimation is not impossible.

Shelter operators should not be afraid of making their shelters attractive.

Uniqueness

For completeness and for further reference, we would like to know when optimal shelter quality is unique.

Formally, we call the problem of finding optimal shelter quality the “uniform benefit problem” to distinguish it from settings where benefits are allowed to vary over time, and where a proper subset of recipients might receive subsidies to become housed. We use an equivalent but slightly modified form of (3):

$$(3') U(w - \tau) = U(b) + \alpha\lambda \int_w^\infty [U(t - \tau) - U(w - \tau)] dF(t).$$

We define $w(\tau)$ as the (unique) value of w that solves (3') and (4). (Uniqueness follows because the left-hand side of (3') is increasing in w and the right-hand side is decreasing.)

Define

$$z(\tau) = w(\tau) - \tau.$$

The uniform benefit problem is to choose τ (and from (4), b) to maximize $z(\tau)$.

From total differentiation of (3') with (4) substituted, we obtain

$$\frac{dz}{d\tau} = \alpha\lambda \frac{n(\tau)}{d(\tau)}$$

where

$$n(\tau) = U'(b(\tau)) [1 - F(w(\tau)) - \tau f(w(\tau))] - \int_{w(\tau)}^\infty U'(t - \tau) dF(t)$$

and

$$d(\tau) = U'(z(\tau)) [1 + \alpha\lambda (1 - F(w(\tau)))] + \alpha\lambda U'(b(\tau)) \tau f(w(\tau)).$$

Since it is easy to see that $d(\tau)\tau > 0$ for all τ , the first order condition for a maximum is

$$n(\tau) = 0.$$

The first term in $n(\tau)$ is the budgetary effect of increasing τ : the term in square brackets shows how much benefits can rise when taxes go up, and multiplying by $U'(b(\tau))$ translates the benefit increase into utility units. The second term gives the expected utility loss from higher taxes when housed. Since this loss is positive, the expression in square brackets is positive at the optimum: optimal taxes are always on the good side of the Laffer curve.

In order for the optimum to be unique, we assume the following second-order condition:

Condition R: (1) $n(\tau)$ is a strictly decreasing function of τ .

$$(2) n(0) > 0.$$

Higher taxes and higher benefits cause strictly decreasing marginal improvements in social well-being. The second part of condition R states that the optimal shelter quality level is positive. This assumption is not strictly necessary, but it simplifies exposition to ignore boundary solutions.

3. Time-varying shelter conditions

The unemployment insurance literature also helps us understand how shelter conditions might vary over the course of a homeless spell, if they can vary. The previous section, of course, assumed that they could not vary at all—every day had to be the same. But it may be possible for the homeless shelter agency to treat new entrants differently from people who have been in the shelter for several months.

As with the optimal constant benefit problem, the unemployment insurance literature makes a distinction between recipients who live hand-to-mouth and recipients who have access to perfect capital markets, and I will pay attention only to hand-to-mouth recipients. The basic papers for this group are Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997). The problem is to maximize the ex ante expected lifetime utility of an unemployed person, subject to an arbitrary constraint on expected agency expenditure.

The unemployment insurance literature finds that in a homogeneous population a declining profile of benefits is optimal; people who have been unemployed longer should get less money per week. This result translates immediately into shelter terms: long-term shelter residents should have lower quality accommodations and consume less non-housing goods than newer shelter residents.

The intuition is the trade-off between incentives and insurance. If recipients were not risk averse, then the optimal strategy would be to give them a lump sum at the beginning of the spell, or a wonderful place to stay for a day, and nothing thereafter. This would achieve the desired level of ex ante utility, and would not distort incentives to find a new job or a new apartment. But it would provide very poor insurance against the bad luck of having a very long spell of unemployment or of homelessness. The compromise between incentives and insurance is a declining sequence of benefits.

Shavell and Weiss provide a quick demonstration that constant benefits are not optimal. Start with a constant stream of benefits and then increase benefits at time t by a small amount and decrease benefits by a second small amount at time $(t+1)$. Choose the two small amounts so that the expected present value of expenditures is constant. Since originally the marginal utility of consumption was the same in both periods, the direct effect on lifetime utility is approximately zero, holding the probability of getting a job constant at each period. But by reducing expected future benefits at time t , the change induces a lower reservation wage at time t , and a higher probability of accepting a job or an apartment then. Reservation wages and opportunities after t are not affected. So the probability of having found a job by any period after t falls. This reduces expected expenditure while keeping lifetime utility the same. So the original constant sequence of benefits could not have been optimal.

Thus benefits should fall during an unemployment spell. Since many shelters have maximum spell durations—an extreme variety of falling benefits—this result is not counterintuitive to shelter operators. But smoothly falling benefits provide better insurance than a hard cutoff.

Most varieties of heterogeneity do not greatly change this result. Suppose that as the length of a homeless spell increases, the quality of opportunities available to people deteriorates, either because they lost contacts or because they are negatively stereotyped because of their shelter experience. Then conditions should fall even faster, because optimal constant conditions are lower for people who have worse outside opportunities. On the other hand, if shelters impart skills that improve the quality of opportunities available, shelter conditions might improve during a homeless spell.

Similarly, suppose there are two types of people, type 1 and type 2, and type 1 people have better outside opportunities. Type 1 people draw from a distribution of opportunities F_1 and type 2 people draw from a distribution F_2 . The distribution F_1 stochastically dominates F_2 . If the homeless shelter agency could distinguish between types, it would set better conditions for type 1 people. But assume that the types are indistinguishable to outside observers.

Then the agency must set the same conditions for both types, although conditions can vary with duration. If type 1 people leave the shelter faster under these conditions, then over time the proportion of type 2 people in the shelter population grows, and the optimal benefit should be closer to the type 2 stand-alone benefit than to the type 1 stand-alone benefit. This would be another reason for benefits to fall with duration. (Shimer and Werning (2006) show this rigorously, but for the case of recipients with access to capital markets and CARA utility. In this case, constant benefits are constant under homogeneity.) On the other hand, if type 2 people leave the shelter faster (because their lower reservation opportunity makes up for their worse opportunities), benefits might rise with duration.

Hagedorn et al. (2007) have argued that a different approach to heterogeneity is optimal: instead of trying to find a one-size-fits-all benefit sequence, it is better to offer a menu of

different sequences and let people self-select. They show that the people with the best outside opportunities will always end up choosing a declining Shavell-Weiss sequence, but that other types may have benefits increasing over time. The intuition is that late benefits are relatively more attractive to people with poor opportunities than to people with good opportunities, because people with good opportunities are more likely to leave before they can enjoy the late benefits. Including a sequence with relatively increasing benefits on the menu is a way to separate the types, and once the types have been separated, to design sequences appropriate to each (subject to the incentive-compatibility constraints that keep the types separated).

One problem with this result is that the cost to a recipient of re-starting a shelter spell is comparatively small—she need only spend one night on the street or in a hotel. That does not mean that realizing the mutual gains from declining benefits is impossible. Instead, the implication is that reset rules—how long you must be gone to have a spell really over, what happens if you return to the shelter before that—have to be very carefully crafted, and benefit rules have to be constrained by reset rules. The problem of optimal benefit profiles with reset rules is an area where considerable shelter-relevant research can be done.

4. Selective post-shelter subsidies

Many shelter systems, New York City being the best known, offer housing subsidies as an inducement for families or individuals to leave shelters. The federal economic stimulus legislation funds more of these subsidies, but gives little guidance on how they should be deployed. These subsidies, however, are not an entitlement; most shelter-leavers do not receive them. At first blush, such selective subsidies seem to be sub-optimal. Risk averse individuals prefer a subsidy for certain to a random subsidy of the same expected cost. Subsidies induce

those who receive them to leave shelters too soon—absent the subsidy they would wait for a better opportunity—and they induce those who don’t receive them to stay in shelters too long—waiting for a subsidy gives them a reason to stay in the shelter longer.

But these intuitions are definitive only in a first-best world with no other distortions present. Shelters are a distortion: they drive a wedge between the true social cost of passing up a housing opportunity and private cost of doing so. If individuals were not risk averse, or if they could be insured some other way, there should be no shelters. If we are talking about shelters at all, we are in a second-best world where first-best intuitions do not necessarily apply.

However, in this case the intuitions from the first-best case are not misguided. Selective subsidies are not an optimal way to encourage people to leave shelters, under mild regularity conditions.

To prove this, we need to expand the notation we used in the discussion of the uniform benefit problem (section 2). A policy is a 3-tuple

$$(\tau, \pi, \sigma) \in \mathbb{R}_+ \times [0,1] \times \mathbb{R}_+$$

of a tax rate, a probability of receiving a housing subsidy, and a housing subsidy amount. A housing subsidy is the same as a selective reduction in the housing tax. Before housing opportunities are drawn, the proportion π of shelter residents are drawn to receive subsidies: if these residents find apartments, they will receive subsidies of σ per period. Our restriction that σ be non-negative entails no loss of generality: a tax of τ with a subsidy of $(-\sigma) < 0$ with probability π is the same as a tax of $(\tau - \sigma)$ and a positive subsidy of σ with probability $(1 - \pi)$. Similarly, we rule out $\pi = 1$ without loss of generality to avoid writing the same policy two ways: $(\tau, 1, \sigma)$ is the same as $(\tau + \sigma, 0, 0)$.

Let w denote the (gross of tax) reservation opportunity for residents without a subsidy; then $(w - \sigma)$ is the reservation opportunity for residents with a subsidy. (You can think of everyone as drawing from the same distribution, but with some draws being higher because they come with a subsidy.) The net-of-tax reservation opportunity is given by

$$(5) \quad U(w - \tau) = U(b) + \alpha\lambda\left[\pi \int_{w-\sigma}^{\infty} [U(t - \tau + \sigma) - U(w - \tau)] dF(t) + (1 - \pi) \int_w^{\infty} [U(t - \tau) - U(w - \tau)] dF(t)\right].$$

This is analogous to equation (3') in section 2. The budget constraint becomes

$$(6) \quad b = \alpha\lambda\left[(\tau - \sigma)\pi(1 - F(w - \sigma)) + \tau(1 - \pi)(1 - F(w))\right].$$

Denote by $w(\tau, \pi, \sigma)$ the unique value of w that solves (5) and (6). An optimal selective benefit policy is one that maximizes

$$Z(\tau, \pi, \sigma) = w(\tau, \pi, \sigma) - \tau.$$

In approaching the problem of finding an optimal selective subsidy policy, we need to take an approach slightly different from the one we took for the uniform benefit problem. Let

$$s = w - \tau.$$

Consider the right-hand side of (5) as a function of s , given the policy 3-tuple:

$$(7) \quad R(s|\tau, \pi, \sigma) = U(b) + \alpha\lambda\left[\pi \int_{s+\tau-\sigma}^{\infty} [U(t - \tau + \sigma) - U(s)] dF(t) + (1 - \pi) \int_{s+\tau}^{\infty} [U(t - \tau) - U(s)] dF(t)\right].$$

Then $Z(\tau, \pi, \sigma)$ is the unique s that solves

$$U(s) = R(s|\tau, \pi, \sigma)$$

and

$$b = \alpha\lambda\left[(\tau - \sigma)\pi(1 - F(s + \tau - \sigma)) + \tau(1 - \pi)(1 - F(s + \tau))\right] \triangleq b(s|\tau, \pi, \sigma).$$

We need the following regularity condition:

Condition S:(1) $\partial R / \partial \tau$ is a decreasing function of τ .

(2) $\partial R / \partial \sigma$ is a decreasing function of σ .

(3) $\frac{\partial R(Z(0, \pi, \sigma))}{\partial \tau} > 0$.

This condition is quite similar to condition R. Notice that it is less stringent than standard second-order conditions for 3-dimensional optimization: for instance, I place no restrictions on cross-partials or derivatives with respect to π .

The basic result of this section is that under these conditions, selective subsidy schemes are never optimal. Let τ^* denote the tax rate that solves the uniform benefit problem.

Proposition 1. If condition S holds, then a policy (τ, π, σ) is optimal if and only if $\tau = \tau^$ and $\sigma = 0$.*

Proof is in the appendix.

Thus in general, when the homeless shelter agency has a full set of instruments available to it, selective subsidies are not optimal. It is better to subsidize all shelter exits a little than to subsidize a few exits a lot.

The theorem also lets us see what sort of constraints on shelter operation might make selective subsidies look attractive. For instance, selective subsidies might be optimal if the homeless shelter agency is constrained to give subsidies of at least a certain size, or if the benefits of a subsidy are non-convex (\$180 a month will get you nothing but \$200 a month will get you a room), or if subsidies of a certain size are cheaper to the homeless shelter agency, perhaps because that is what the federal government provides.

Selective subsidies, then, are optimal only if the shelter agency faces some additional constraints or additional non-convexities. Specifying those constraints or non-convexities is the first step in developing any plan for selective subsidies.

5. Timing of selective subsidies

Since receiving a selective subsidy is a good thing, the optimal timing of selective subsidies is similar to the optimal timing of benefits. This problem, of course, presents itself only when selective subsidies should be given and when the shelter agency can reasonably observe how long an individual has stayed in the shelter. If the conditions of proposition 1 hold, or if individuals can re-start their shelter clocks almost costlessly, then it makes no sense to look for the optimal timing of selective subsidies.

I examine the optimal problem at length in O’Flaherty (2009), which draws heavily on Shavell and Weiss (1979) and Hagedorn et al. (2007). When the shelter population is homogeneous, the solution is simple, obvious, and totally contrary to current practice: to the extent that subsidized placements occur, they should occur immediately on shelter entry.

The setting is one in which shelter quality and the value of the placement subsidy are fixed. The shelter agency chooses probabilities of placement after different durations of shelter stay in order to minimize its costs subject to a minimum ex ante utility constraint for homeless individuals.

The intuition is that any policy other than immediate placement generates moral hazard problems: it rewards people who stay in shelters longer with a higher probability of receiving a subsidy. The standard objection to an immediate placement policy—that it will make shelters too attractive and draw too many people in—has no bite in this setting, because the probability of (immediate) placement can be adjusted to make shelter entry as attractive or unattractive as desired.

Suppose, for instance, that in an initial state a shelter with immediate placements is “too attractive”: ex ante expected utility is too high and too many individuals are entering the shelter. The shelter system can be made less attractive either by delaying placement or by reducing the probability of immediate placement. Delaying placement makes people stay in the shelter longer and so costs money for the shelter agency. Reducing the probability of immediate placement saves the shelter agency money. Reducing the placement probability is obviously the better adjustment. With a homogeneous population, spending money in order to make people miserable because you will subsequently spend money to make them too happy is not a cost-effective way of doing things.

With a heterogeneous population, some placements can be delayed. The optimal strategy is to present shelter entrants with a menu of different ‘contracts’ and let them choose the contract that they think is best for themselves. One contract offers immediate placement at some probability, with no chance for later placement. Another contract offers a smaller probability of immediate placement, but a positive probability of later placement. Individuals who think they have a better chance of leaving the shelter on their own (“good searchers”) will choose the immediate placement contract—if the menu is designed right—and individuals who expect more difficulty finding an apartment on their own (“poor searchers”) will choose the delayed placement contract.

The advantage to the shelter agency of this arrangement over a more standard uniform placement probability profile is that it reduces the ex ante utility of good searchers and their associated cost. A uniform contract would need to have a high enough placement probability that poor searchers could achieve the minimum utility ex ante; good searchers then would be able to achieve much higher ex ante utility. But with a menu of contracts, the good searcher

contract needs to be only as attractive to good searchers as the poor searcher contract is to them (good searchers). Since good searchers are not likely as poor searchers to stay in the shelter long enough to benefit from late delayed placements, the poor searcher contract can be made unattractive to them, and so the shelter agency can save money (and induce fewer entries) by reducing the probability of immediate placement for good searchers.

This separating menu of contracts works only if shelter entrants have some private information about their chances of leaving the shelter on their own. We do not know whether they have any such useful information. A useful research project would be to ask shelter entrants how long they thought they would stay in the shelter, and see whether their predictions helped predict actual shelter stay, once the usual observable variables (possibly including staff predictions) were controlled for.

6. Street homelessness

This approach is applicable to street homelessness as well as shelter homelessness. Many homeless people (about 281,000 in January 2007 (U.S. HUD 2008) do not use shelters. Instead they sleep on subways, in transportation terminals, under highways, and in many other places that were not designed for human habitation.

Extensive debates are carried on, often emotionally, between those who advocate police action of some form to roust the homeless, and those who counsel forbearance and sometimes assistance. The argument for rousting is that street homelessness has costly externalities for the rest of society (and sometimes for the homeless person himself); the argument for forbearance is insurance—society should not add to the woes of those who are already down on their luck.

If transportation terminals and other places are considered as shelters, the reasoning in section 2 can easily be applied to this controversy. The costs of operating these “shelters” are the full social costs, including externalities and unintended consequences for the homeless themselves. Since the technology for producing benefits in these “shelters” is extremely inefficient, the optimal quality is probably quite low, and police harassment may be necessary in some cases to produce optimally low quality. But not in all cases: situations that are inherently uncomfortable, that produce low external costs or that are only rarely used may require no police response, and possible assistance instead. Ellickson (1996) similarly argues for a differentiated response to street homelessness. Thus the theory of optimal shelter quality is also a theory of optimal street enforcement.

Street homelessness also raises questions about placement into subsidized housing. Housing First, an extremely popular program in the U.S. and Canada, operates by placing mentally ill street homeless people in subsidized apartments. Analysis of Housing First and similar programs is beyond the scope of this paper, but should be amenable to closely related methods.

7. Conclusion

Housing First shows the two major areas where further theoretical research is needed.

The first is more work on heterogeneity. We need to know more about how shelter management is affected when homeless people differ—by drawing from different distributions of post-shelter opportunities, for instance or experiencing shelter life differently. In some discussions, particularly O’Flaherty (2009), where the shelter agency had fewer

instruments, and in parts of this paper, I have considered heterogeneity, but not in the description of optimal shelter quality or of placement strategy.

The second problem is to extend the analysis to multiple shelters. If “the street” is counted as a shelter, then every city has multiple shelters, and the operators of one do not necessarily cooperate with the operators of others. In O’Flaherty (2009) I showed how heterogeneity across shelters can produce results better than any number of homogeneous shelters could produce. How shelters should operate when there are many of them is a challenging and important question.

References

Culhane, Dennis P. and Stephen Metraux, Jung Min Park, Maryanne Schretzman, and Jesse Valente, 2007, "Testing a typology of family homelessness based on patterns of public shelter utilization in four U.S. jurisdictions: Implications for policy and program planning," *Housing Policy Debate* 18 (1): 1-28.

Ellickson, Robert, 1996, "Controlling misconduct in city spaces: Of panhandlers, skid rows, and public-space zoning," *Yale Law Journal* 105: 1165-1248.

Hagedorn, Markus, Ashok Kaul, and Tim Mennel, 2007, "An adverse selection model of optimal unemployment insurance," working paper, Zentrum für Europäische Wirtschaftsforschung (ZEW), Mannheim.

Hopenhayn, H. A., and J. P. Nicolini, 1997, "Optimal unemployment insurance," *Journal of Political Economy* 105 (2): 412-38.

O'Flaherty, Brendan, 2008, "Homelessness as bad luck: Implications for research and policy." Paper presented at conference on How to House the Homeless, November 21, New York.

---, 2009, "When should homeless families get subsidized apartments? A theoretical inquiry," *Journal of Housing Economics*, forthcoming.

--- 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.

--- and ---, 2008, "Homeless shelters for single adults: Why does their population change," *Social Service Review*, 82(3): 511-50

Shavell, S. and L. Weiss, 1979, "The optimal payment of unemployment insurance benefits over time," *Journal of Political Economy* 87 (6): 1347-62.

Shimer, Robert and Ivan Werning, 2006, "On the optimal timing of benefits with heterogeneous workers and human capital depreciation," working paper, MIT.

--- and ---, 2007, "Reservation wages and unemployment insurance," *Quarterly Journal of Economics* 122 (3): 1145-85.

U.S. Department of Housing and Urban Development, 2008, *The 2007 Annual Homelessness Assessment Report*. Washington, D.C.: the Department.

Appendix : Proof of proposition 1.

The basic idea is to show that for any z , $R(z|\tau, \pi, \sigma)$ is maximized only if $\pi=0$. Given $\pi=0$, the selective subsidy problem is the same as the uniform benefit problem, and so is solved the same way, with $\tau=\tau^*$. We proceed by a series of steps.

Step one. Fix s , substitute the budget constraint into (7) in the text, and take derivatives with respect to the three policy variables:

$$\begin{aligned} \frac{1}{\alpha\lambda} \frac{\partial R}{\partial \tau} &= U'(b) [\pi(1 - F(s + \tau - \sigma) - (\tau - \sigma)f(s + \tau - \sigma)) \\ &\quad + (1 - \pi)(1 - F(s + \tau) - \tau f(s + \tau))] - \pi \int_{s+\tau-\sigma}^{\infty} U'(t - \tau + \sigma) dF(t) \\ &\quad - (1 - \pi) \int_{s+\tau}^{\infty} U'(t - \tau) dF(t). \end{aligned}$$

$$\frac{1}{\alpha\lambda} \frac{\partial R}{\partial \pi} = U'(b) [(\tau - \sigma)(1 - F(s + \tau - \sigma)) - \tau(1 - F(s + \tau))] +$$

$$\int_{s+\tau-\sigma}^{\infty} [U(t - \tau + \sigma) - U(s)] dF(t) - \int_{s+\tau}^{\infty} [U(t - \tau) - U(s)] dF(t).$$

$$\begin{aligned} \frac{1}{\alpha\lambda} \frac{\partial R}{\partial \sigma} &= \pi \left\{ U'(b) [-(1 - F(s + \tau - \sigma)) + (\tau - \sigma)f(s + \tau - \sigma)] \right. \\ &\quad \left. + \int_{s+\tau-\sigma}^{\infty} U'(t - \tau + \sigma) dF(t) \right\}. \end{aligned}$$

Let $\sigma=0$. Then it is easy to see that

$$\frac{\partial R}{\partial \pi} = \frac{\partial R}{\partial \tau} = 0 \text{ if and only if } \tau = \tau^*.$$

Moreover

$$(A1) \quad \frac{\partial R(\tau, \pi, \sigma)}{\partial \sigma} = -\pi \frac{\partial R(\tau, \pi, \sigma)}{\partial \tau}.$$

Hence $(\tau^*, \pi, 0)$ is a local maximum. The remainder of the proof will show that it is also a global maximum.

Step two. From condition S(1), if $\tau < \tau^*$,

$$\frac{\partial R}{\partial \tau} > 0,$$

and so by (A1)

$$\frac{\partial R}{\partial \sigma} < 0.$$

Hence by condition S(2), there is no maximum of R with $\sigma > 0$, $\tau < \tau^*$.

Step three. Suppose (τ, π, σ) is optimal with $\sigma > 0$. Then

$$\frac{\partial R(\tau, \pi, \sigma)}{\partial \pi} = 0.$$

Consider the last two terms of this derivative:

$$\begin{aligned} & \int_{s+\tau-\sigma}^{\infty} [U(t-\tau+\sigma) - U(s)] dF(t) - \int_{s+\tau}^{\infty} [U(t-\tau) - U(s)] dF(t) \\ &= \int_{s+\tau-\sigma}^{s+\tau} [U(t-\tau+\sigma) - U(s)] dF(t) - \int_{s+\tau}^{\infty} [U(t-\tau+\sigma) - U(s)] dF(t) > 0, \end{aligned}$$

since $t - \tau + \sigma \geq s$ whenever $t \geq s + \tau - \sigma$. Because $U'(b) > 0$, and the derivative is assumed zero, we have

$$(A2) \quad (\tau - \sigma)(1 - F(s + \tau - \sigma)) - \tau(1 - F(s + \tau)) < 0.$$

Step four. Observe that in general

$$(A3) \quad \frac{1}{\alpha\lambda} \frac{\partial R}{\partial \tau} = -\frac{1}{\alpha\lambda} \frac{\partial R}{\partial \sigma} + (1 - \pi)\{U'(b)[1 - F(s + \tau) - \tau f(s + \tau)] - \int_{s+\tau}^{\infty} U'(t - \tau) dF(t)\} \triangleq \\ = -\frac{1}{\alpha\lambda} \frac{\partial R}{\partial \sigma} + (1 - \pi)X(\tau, \pi, \sigma).$$

Here we have defined the expression in curly brackets as $X(\tau, \pi, \sigma)$.

Step five. Suppose $\tau > \tau^*$. Then

$$(A4) \quad \frac{1}{\alpha\lambda} \frac{\partial R(\tau, \pi, 0)}{\partial \tau} = X(\tau, \pi, 0) < 0$$

by condition S(1).

Step six. Suppose (τ, π, σ) is optimal with $\sigma > 0$ and $\tau > \tau^*$. Consider

$$b(s|\tau, \pi, \sigma) = \alpha\lambda[\pi(\tau - \sigma)(1 - F(s + \tau - \sigma)) + (1 - \pi)\tau(1 - F(s + \tau))] \\ = \alpha\lambda[\pi\{(\tau - \sigma)(1 - F(s + \tau - \sigma)) - \tau(1 - F(s + \tau))\} + \tau(1 - F(s + \tau))] \\ < \alpha\lambda\tau(1 - F(s + \tau)) = b(s|\tau, \pi, 0).$$

The inequality follows from (A2). Hence by concavity

$$(A5) \quad U'(b(s|\tau, \pi, \sigma)) < U'(b(s|\tau, \pi, 0)).$$

Step seven. Since we have assumed that (τ, π, σ) is optimal,

$$\frac{\partial R(\tau, \pi, \sigma)}{\partial \sigma} = 0.$$

Thus from (A3)

$$\frac{1}{\alpha\lambda} \frac{\partial R}{\partial \tau} = (1 - \pi)X(\tau, \pi, \sigma),$$

which we label (A6).

Since $\tau > \tau^*$,

$$\frac{1}{\alpha\lambda} \frac{\partial R(\tau, \pi, 0)}{\partial \tau} = U'(b(s|\tau, \pi, 0))[1 - F(s + \tau) - \tau f(s + \tau)] - \int_{s=\tau}^{\infty} U'(t - \tau) dF(t) < 0.$$

Now consider whether

$$1 - F(s + \tau) - \tau f(s + \tau) \leq 0.$$

Suppose so. Since

$$- \int_{s+\tau}^{\infty} U'(t - \tau) dF(t) < 0,$$

$X(\tau, \pi, \sigma) < 0$, and from (A3),

$$\frac{\partial R(\tau, \pi, \sigma)}{\partial \tau} < 0.$$

This contradicts the presumed optimality.

So suppose

$$1 - F(s + \tau) - \tau f(s + \tau) > 0.$$

Then from (A5)

$$\begin{aligned} 0 &> U'(b(s|\tau, \pi, 0))[1 - F(s + \tau) - \tau f(s + \tau)] - \int_{s+\tau}^{\infty} U'(t - \tau) dF(t) \\ &> U'(b(s|\tau, \pi, \sigma))[1 - F(s + \tau) - \tau f(s + \tau)] - \int_{s+\tau}^{\infty} U'(t - \tau) dF(t) \\ &= X(\tau, \pi, \sigma). \end{aligned}$$

Hence from (A6)

$$\frac{\partial R}{\partial \tau} < 0.$$

Thus (τ, π, σ) is not optimal. Since this is a contradiction, we have proved that (τ, π, σ) is optimal only if $\sigma=0$, and that the local maximum at $(\tau^*, \pi, 0)$ is a global maximum.