Targeting with In-kind Transfers: Evidence from Medicaid Home Care

Ethan M.J. Lieber and Lee M. Lockwood*

July 3, 2017

Abstract

Many of the most important government programs make transfers in kind as opposed to in cash. Making transfers in kind has the obvious cost that recipients would prefer cost-equivalent cash transfers. But making transfers in kind can have benefits as well, including better targeting transfers to desired recipients. In this paper, we exploit large-scale randomized experiments run by three state Medicaid programs to investigate this central tradeoff for in-kind provision. We find that in-kind provision of formal home care significantly reduces the value of benefits to recipients while targeting benefits to a small subset of the eligible population that is sicker and that has worse informal care options than the average eligible. Under a wide range of assumptions within a standard model, the targeting benefit exceeds the distortion cost. This highlights an important cost of recent reforms that move toward more flexible, cash-like benefits.

*Lieber: Notre Dame, Ethan.Lieber.2@nd.edu; Lockwood: Northwestern University and NBER, lee.lockwood@northwestern.edu. We are grateful to Norma Coe, Gopi Shah Goda, Seema Jayachandran, Brian Melzer, Matt Notowidigdo, Iuliana Pascu, Mike Powell, Emmanuel Saez, Diane Schanzenbach, Jesse Shapiro, Courtney Van Houtven, and many seminar participants for helpful comments. We thank Jose Carreno, Vishal Kamat, and Zeyu Wang for excellent research assistance. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) through the Michigan Retirement Research Center (Grant #5 RRC08098401-05), funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of the Social Security Administration, any agency of the Federal Government, or the Michigan Retirement Research Center.

1 Introduction

In-kind transfers are a ubiquitous feature of government programs, private contracts, and charitable giving. In the U.S. in 2015, government spending on in-kind health care alone totaled more than 8 percent of GDP (Centers for Medicare and Medicaid Services, 2017). In-kind transfers are also at the heart of a crucial debate about the relative desirability of benefit programs that are more universal and flexible versus more targeted and restrictive.[1]

At the heart of this debate is a key tradeoff inherent to in-kind transfers. In-kind provision has a fundamental cost: Recipients would prefer cost-equivalent cash transfers. But this cost is linked to an important potential benefit: If in-kind provision leads certain people to take up more benefits than others, it can improve the targeting of transfers when information or other constraints preclude more direct targeting (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988). In the context of insurance, for example, if someone's valuation of a particular in-kind benefit is larger in states of the world in which marginal utility is higher, in-kind provision can help concentrate benefits in those states and thereby better insure the risk. In such cases, there is a tradeoff between providing benefits that are more valuable to recipients (for which less restrictive cash-like benefits are best) and providing benefits that better target transfers to higher-marginal utility states (for which more restrictive in-kind benefits might be best). These costs and benefits are crucial determinants of the optimal benefit design.

In this paper, we develop a general framework for analyzing this key tradeoff of in-kind provision, and we apply it to the context of home care. Home care is a set of services such as assistance with eating, dressing, and bathing that helps individuals with chronic health problems live at home instead of in nursing homes. The vast majority of home care is provided informally by family and friends ("informal care"), and the remainder is provided by hired caregivers ("formal care") (Arno et al., 1999). This is an especially important context in which to analyze the consequences of in-kind provision for two main reasons. First, it is one of the largest and fastest-growing components of what is likely the largest and fastest-growing type of in-kind benefits: health benefits. Spending on home care was $88 billion in 2015 and is growing rapidly both in absolute terms and as a share of health spending.[2] Second, many states in the U.S. and countries in Europe have reformed their

[1] For example, in several rich countries, there is an active debate about the desirability of universal basic income programs versus more restrictive and targeted transfers of food, housing (Collinson et al., 2015), medical care (Doty et al., 2010), and other goods.

[2] This statistic is based on data from the national health expenditure accounts. In the U.S., government spending on in-kind health-benefit programs totals more than $1.3 trillion and the recent Affordable Care Act increased such in-kind health benefits substantially through expanded Medicaid eligibility and subsidies for health insurance. Moreover, much of the remainder of the roughly 20 percent of GDP worth of spending on health care comes through in-kind benefits via private insurance contracts.

home care programs to make the benefits more flexible and cash-like (National Conference of State Legislatures, 2007; Da Roit and Le Bihan, 2010). Fifteen state Medicaid programs allow recipients to use benefits to pay informal caregivers or buy equipment for their homes (Doty et al., 2010). In addition, early versions of the bill that became the Affordable Care Act included a long-term care insurance program that would have paid cash benefits.[3]

An important milestone in the movement toward more flexible benefits, and a key source of evidence in our analysis, is the Cash and Counseling demonstrations. These large-scale experiments, implemented around the year 2000 in Arkansas, Florida, and New Jersey, randomized a subset of Medicaid in-kind home care participants to either Medicaid's standard in-kind formal care or a near-cash benefit.[4] Existing research on the Cash and Counseling experiments has found that at the follow-up people randomized to the near-cash benefit were in no worse health and reported much greater satisfaction with their care than those randomized to the in-kind benefit (Foster et al., 2003; Brown et al., 2007; Lepidus Carlson et al., 2007). This evidence has played an important role in the growth of more flexible, cash-like benefits (Doty et al., 2010).

We begin by extending existing research on the Cash and Counseling experiments by estimating the price sensitivity of demand for formal care and the value to recipients of the traditional in-kind benefit.[5] Using the exogenous variation generated by the randomization, we find that a one-dollar reduction in the hourly price of formal care leads to a 1.8-hour-per-week increase in consumption. This implies that a full subsidy would increase consumption by 24.7 hours per week—3.5 times average consumption among those randomized to the near-cash benefit—and that a recipient of the average in-kind transfer would value it at just 28 percent of its cost. This low value is consistent with earlier research on other, more subjective measures of value to recipients using the Cash and Counseling experiments.

But the low value of the in-kind benefit does not imply that more flexible benefits are better. That recipients value the in-kind benefit so much less than its cost raises the likelihood that in-kind provision has important effects on take up, and thereby on the distribution of benefits within the eligible population. We find that program expenditures are much more concentrated among people randomized to the in-kind benefit; the variance in benefits is 10 times greater in the in-kind group. In addition, using nationally representative data, we

[3]This program, known as the CLASS (Community Living Assistance Services and Supports) Act, was eventually repealed due to concerns about its budgetary sustainability.

[4]We use the term "near-cash" because, although recipients were required to spend their benefit on care-related expenses, virtually every recipient received enough informal care that the benefit should be infra-marginal.

[5]The large, exogenous variation in the price of formal care generated by these experiments offers an important advantage over other settings. Non-experimental estimates face a particularly difficult identification problem. Many factors that shift the supply of formal care also shift the supply of *informal* care and thereby shift the demand for formal care. This leads to simultaneity bias.

estimate that between 4 and 16 percent of people eligible for Medicaid home care choose to take up benefits. Those who do take up benefits tend to be sicker and poorer and tend to have fewer potential informal caregivers. To the extent that these states of the world tend to have relatively high marginal utility, these patterns are consistent with in-kind provision improving risk sharing.

The low value that recipients place on the in-kind benefit raises the question of whether a less costly means of targeting could be employed. In principle, insurers could use tags—observable, ideally immutable characteristics—to target cash benefits to high-marginal utility states instead of having the benefit itself serve the additional purpose of targeting (Akerlof, 1978). In practice, we find that even very rich tags, including one based on a detailed medical examination, would leave the vast majority of this risk uninsured. This suggests that in the context of home care, it may not be possible to avoid a stark version of the tradeoff described above: Restrictive in-kind benefits are much less valuable to recipients, but flexible cash-like benefits leave most of the risk uninsured.

The reduced-form analysis provides strong evidence that in-kind provision reduces the value of benefits to recipients and suggestive evidence that in-kind provision improves targeting. This raises the question, what is the net welfare effect of in-kind provision? We use a structural model to quantify these costs and benefits in a unified framework. We make the minimum necessary adjustments to the standard model and test the robustness of our results to a wide range of alternatives. The key inputs are the price sensitivity of demand for formal care and the distribution of demand for formal care, both of which we estimate directly. We find that across a wide range of assumptions, the optimal contract involves a large in-kind component and delivers substantial welfare gains over cash-benefit contracts. The desirability of providing formal care in kind arises from the significant heterogeneity in the demand for formal care, which in the absence of a large formal care subsidy translates into significant heterogeneity in non-care consumption and marginal utility.

Our results provide new insights into long-term care risk, insurance, and policy. There are many known barriers to private, voluntary long-term care insurance, including Medicaid and adverse selection (see Brown and Finkelstein, 2011, for a review). We estimate the importance of two fundamental barriers to any long-term care insurance, private or government, voluntary or mandatory: hard-to-verify heterogeneity and moral hazard. Our findings reveal a fundamental dilemma for insuring long-term care risk. On the one hand, that formal care consumption is highly sensitive its price implies that many recipients would be significantly better off *ex post* with a cost-equivalent cash transfer. It also implies that a large "moral hazard tax" plagues most long-term care insurance contracts—an under-appreciated difficulty facing a market with many other challenges besides—and raises the effective loads to

consumers above existing "supply-side" estimates that do not account for moral hazard (e.g., Brown and Finkelstein, 2007; Friedberg et al., 2014). On the other hand, that even richly-tagged cash benefits leave most of the risk uninsured suggests that providing home care in kind might be an unfortunate but necessary cost of insuring the risk from chronic health problems. Especially when combined with the other potential benefits of in-kind provision of home care, our findings raise concerns about the many recent reforms that make long-term care benefits more flexible and cash-like.[6]

Our paper contributes to the literature on in-kind transfers. Although the theoretical branch of this literature analyzes a variety of potential benefits of in-kind provision (see Currie and Gahvari, 2008, for a review), the empirical literature has mostly focused on estimating the consumption distortion from providing benefits in kind.[7] For example, Moffitt (1989), Whitmore (2002), and Hoynes and Whitmore Schanzenbach (2009) all find that providing food in kind has relatively small effects on recipients' choices relative to cost-equivalent cash transfers because many recipients are inframarginal and resale opportunities allow them to convert their in-kind benefit into cash. Whereas the apparent similarity of food stamps to cash benefits reduces the likelihood that in-kind provision has important effects in that context, in the many contexts in which in-kind transfers are large and resale is difficult—such as the major in-kind health care and education benefit programs—the fundamental tradeoff we analyze could be of great importance.

Our paper also contributes to the literature on incomplete take up of transfer programs. That many important government programs are taken up by only a small fraction of those eligible for benefits has motivated research aimed at understanding take up decisions (see Currie, 2006, for a review). While in some contexts low take up is undesirable, our analysis suggests that in the context of home care, low take up significantly improves risk sharing. Researchers have investigated the extent to which different personal and household characteristics are associated with take up in contexts such as disability insurance (Low and Pistaferri, 2015; Deshpande and Li, 2017), Medicaid (Cutler and Gruber, 1996), housing assistance (Reeder, 1985), and Supplemental Security Income (Benitez-Silva et al., 2004).

[6] An extensive theoretical literature investigates a variety of potential benefits of in-kind provision in addition to targeting, including increasing the efficiency of the tax system (Munro, 1992), reducing moral hazard in the context of the Samaritan's Dilemma (Bruce and Waldman, 1991), internalizing externalities, indulging paternalistic preferences, shifting prices in a desirable way ("pecuniary effects," Cunha et al., 2011), mitigating asymmetric information problems, solving political economy problems, and redistributing resources within households or families (see Currie and Gahvari, 2008, for a review). Some of these other potential advantages of in-kind provision may be important in the context of home care as well, especially tax system efficiency (since providing informal care appears to reduce market work, e.g., Ettner, 1995) and perhaps the Samaritan's Dilemma (providing care to one's elderly parents may make it more likely that one will rely on means-tested transfers in the future). Any such benefits would increase the relative desirability of in-kind provision of home care.

[7] An exception is Cunha et al. (2011), who find that in-kind provision of food transfers reduced food prices in Mexican villages.

Other work has investigated the targeting effects of ordeals (Atalas et al., 2016), subsidized prices (Cohen and Dupas, 2010), and delegating authority over the distribution of benefits to local leaders (Atalas et al., 2012; Basurto et al., 2017). Our work complements and extends this literature by analyzing not only the targeting effects but also the welfare consequences of different benefit designs and eligibility rules, including both the targeting effects and the value of benefits to recipients. The approach we develop to do so should be useful in many other contexts as well.

2 Theory

This section presents a simple model of in-kind benefits that focuses on the targeting benefit and distortion cost of in-kind provision. It combines features from the welfare analysis of a tax or subsidy, where efficiency costs are central, with features from optimal income taxation and the welfare analysis of social insurance, where redistribution across types is central.

The key feature of in-kind provision is that the size of the transfer an individual receives depends on his or her consumption of the good in question. One can view an in-kind benefit program as providing a cash benefit while at the same time imposing a restriction on recipients that they must consume at least a certain amount of the good in question. As Nichols and Zeckhauser (1982) emphasize, imposing restrictions on recipients can improve the targeting of benefits to desired recipients who cannot otherwise be distinguished from would-be "mimics," if meeting the restriction is more costly for mimics than for desired recipients. Imposing such a restriction relaxes the incentive compatibility constraints on mimics' participation and thereby allows the program to make greater transfers to desired recipients.

An in-kind benefit can be modeled as a (potentially non-linear) price subsidy. Many in-kind benefit programs, such as food stamps, offer individuals up to a fixed quantity of the good at no charge. When resale is not possible, this has the same effect on a participating individual's budget constraint as a non-linear price subsidy of 100 percent on units up to the benefit limit and 0 percent on units above the limit.[8] In this section we focus on the case of a subsidy program with no quantity limit. We do this both for simplicity of exposition and because in many states, including the states that ran the Cash and Counseling experiments, the Medicaid home care program does not appear to have binding benefit limits in practice.

[8]The nature of resale opportunities, if any, is an important determinant of the effects of in-kind benefit programs. In the case of home care benefits, resale is impossible. In the case of food stamps, by contrast, resale markets are an important feature of the environment. Whitmore (2002) presents survey evidence that food stamps trade at about 65 percent of their face value in the resale market.

The results are easily extended to cases with benefit limits.

The key considerations for in-kind provision can be seen in Figure 1. Figure 1 shows the values (in terms of equivalent variations) and efficiency costs of a price subsidy on good X for each of two consumers with different levels of demand for X. The price subsidy is worth less to each consumer than it costs the government or insurance company to provide due to the induced consumption distortion. The size of this distortion is increasing in the compensated own-price elasticity of demand. This is entirely standard. Less-commonly emphasized is the key role of heterogeneity in the level of demand for X in determining both the distributional consequences of subsidizing X and the share of total spending on the subsidy that is costed away by distortions. The same price subsidy is more valuable to people who consume more of the subsidized good, so, relative to a cost-equivalent cash benefit, the subsidy redistributes toward people who consume more X from people who consume less X. In addition, the share of the program's spending on the subsidy that benefits the recipient as opposed to being costed away by distorting consumption is increasing in the recipient's level of demand for the good. As a result, the extent to which consumption distortions reduce the desirability of in-kind provision depends not only on the price sensitivity of demand but also on the distribution of the level of demand across types.

2.1 The benefit program and its budget constraint

Consider the problem of designing a mandatory benefit program for a population of ex-ante identical individuals whose ex-post distribution of types is $F(\theta)$.[9] An individual's type, $\theta_i \sim F(\theta)$, embeds all of the individual's characteristics that are relevant for determining the costs and benefits of alternative benefit designs, including any relevant heterogeneity in preferences and budget constraints. The planner knows the distribution of types, $F(\theta)$, but cannot verify any single individual's type.[10,11]

[9] That the program is mandatory is not essential. It just simplifies the exposition by eliminating take-up decisions.

[10] We focus on the problem of designing benefits for a population of ex-post heterogeneous individuals who cannot be verifiably distinguished. This population could be a sub-population of a broader population, where the sub-population is distinguished from the broader population by the values of some verifiable characteristics. Differences in verifiable characteristics allow the program to "partition" the population into sub-groups and treat these groups differently. Groups of individuals who can be verifiably distinguished can be given different benefits directly; the planner need not resort to imposing restrictions on recipients in order to redistribute across these groups. In the extreme case in which all of the heterogeneity in marginal utility occurs across rather than within groups, pure cash-benefit contracts can achieve the first-best allocation. In the opposite extreme in which all of the heterogeneity in marginal utility occurs within rather than across groups, pure cash-benefit contracts can provide no insurance within the eligible population.

[11] For simplicity, we ignore any second-best considerations that might arise from the interaction between the program and other distortions in the economy. We discuss such considerations in Section 6 and the conclusion.

Consider an idealized in-kind benefit program that potentially combines two elements: a cash benefit, b, and a linear subsidy on good X, σ. The cash benefit and subsidy rate are common across all eligible individuals and are automatic in the sense that there are no take-up decisions; all eligible individuals receive the cash benefit and are subsidized on their purchases of good X. Two special cases of this combined cash-plus-subsidy program are a pure cash-benefit program ($b > 0, \sigma = 0$) and a pure subsidy program with no cash benefit ($b = 0, \sigma > 0$). A pure in-kind benefit program like Medicaid home care has a zero cash component and a full subsidy, ($b = 0, \sigma = 1$).

Average per-eligible spending on the program, B, is divided between funding the cash benefit, b, and the subsidy on X, σ:

$$\int_{\Theta} \left(b + (\sigma p_X(\sigma)) x_X(\sigma;\theta)\right) f(\theta) d\theta = B,$$

where $p_X(\sigma)$ is the subsidy-exclusive price of X (the sellers' price) and $x_X(\sigma;\theta)$ is the consumption of X by type θ as a function of the subsidy rate.

2.2 Analysis of a budget-neutral shift toward in-kind benefits

This section analyzes a marginal shift in benefits toward in-kind benefits. This shift involves marginally increasing the subsidy rate, σ, and at the same time decreasing the cash benefit in order to maintain the same program budget.

For simplicity, suppose that the supply of every good is perfectly elastic. In this case, an increase in the subsidy reduces buyers' after-subsidy price of X one-for-one (no incidence on supply), $p_X(\sigma) = (1 - \sigma)p_X^0$, and has no effect on the prices of goods other than X, $p_i(\sigma) = p_i^0$ for $i \neq X$, where p_i^0 is the price of good i without any benefit program.

Marginally increasing the subsidy rate while at the same time decreasing the cash benefit in order to maintain the same program budget implies the following change in the cash benefit:

$$\begin{aligned}
\frac{\partial b(\sigma, B)}{\partial \sigma} &= -\int_{\Theta} \left[x_X(\sigma;\theta) p_X^0 + (\sigma p_X^0) \frac{dx_X(\sigma;\theta)}{d\sigma} \right] f(\theta) d\theta \\
&= -\left[E_{\Theta}\left(x_X(\sigma;\theta) p_X^0 \right) + E_{\Theta}\left((\sigma p_X^0) \frac{dx_X(\sigma;\theta)}{d\sigma} \right) \right].
\end{aligned}$$

The cash benefit must fall by the increase in average per-eligible spending on the in-kind benefit (subsidy). Average spending on the subsidy is the sum of two terms: (i) the mechanical increase in spending on the subsidy due to the increase in the subsidy rate, holding fixed each type's consumption of X, $E_{\Theta}\left(x_X(\sigma;\theta)p_X^0\right)$ ("mechanical effect"); and (ii) the increase

in spending on the subsidy due to the induced change in consumption of X in response to the shift in program benefits, $E_\Theta\left((\sigma p_X^0)\frac{dx_X(\sigma;\theta)}{d\sigma}\right)$ ("behavioral effect").[12]

2.2.1 The net ex-post value for each type of a shift toward in-kind provision

For each type, the net ex-post value of the change in the program is the benefit of the increase in the subsidy on X (i.e., the benefit from the reduction in the after-subsidy price of X) less the cost of the reduction in the cash benefit. A marginal increase in the subsidy rate on X reduces the after-subsidy price of X by p_X^0. The value (in units of income) of this reduction in the price of X to an individual of type θ is, by the envelope theorem (Roy's identity),

$$\frac{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial p_X(\sigma)}\frac{dp_X(\sigma)}{d\sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}} = x_X(\sigma;\theta)p_X^0,$$

where $v(p(\sigma), m(\sigma, B); \theta)$ is the indirect utility function of type θ and $m(\sigma, B) = m^0 + b(\sigma, B)$ is benefit-inclusive income. This benefit from a lower after-subsidy price of X must be weighed against the reduction in the cash benefit required to hold fixed total spending on the program. Combining these two elements gives the net value (in units of income) of a budget-neutral marginal shift toward in-kind benefits of

$$\begin{aligned}\frac{dV(\sigma;\theta)}{d\sigma} \equiv \frac{\frac{dv(p(\sigma),m(\sigma,B);\theta)}{d\sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}} &= \frac{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial p_X}\frac{dp_X(\sigma)}{d\sigma} + \frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}\frac{\partial b(\sigma,B)}{\partial \sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}} \\ &= x_X(\sigma;\theta)p_X^0 - \int_\Theta \left[x_X(\sigma;\theta)p_X^0 + (\sigma p_X^0)\frac{dx_X(\sigma;\theta)}{d\sigma}\right] f(\theta)d\theta \\ &= \left[x_X(\sigma;\theta)p_X^0 - E_\Theta\left(x_X(\sigma;\theta)p_X^0\right)\right] - (\sigma p_X^0)E_\Theta\left(\frac{dx_X(\sigma;\theta)}{d\sigma}\right). \quad (1)\end{aligned}$$

The marginal net value for an individual of type θ of a budget-neutral marginal shift in benefits toward in-kind benefits is the net benefit of the resulting redistribution to his type (redistribution benefit), $\left[x_X(\sigma;\theta)p_X^0 - E_\Theta\left(x_X(\sigma;\theta)p_X^0\right)\right]$, which is greater for types with greater levels of demand for X, less the average marginal distortion cost from the induced change in consumption of X (distortion cost), $(\sigma p_X^0)E_\Theta\left(\frac{dx_X(\sigma;\theta)}{d\sigma}\right)$.

Equation 1 shows that the shift toward in-kind provision has two key effects. It redistributes toward people with above-average demand for the good, and it distorts consumption of the

[12]The "behavioral effect" can be positive or negative, though in most cases it will be positive. It embeds the income effects from the reduction in cash benefits, which tend to reduce the consumption of X (provided X is normal), and substitution and income effects from the reduction in the after-subsidy price of X, which tend to increase consumption of X. A shift toward in-kind provision increases average consumption of X unless income effects of demand for X are much larger among those who lose from the shift than among those who gain.

good. The extent to which a particular type gains from a marginal shift toward greater in-kind provision is increasing in that type's level of consumption of the good and decreasing in the average sensitivity of the demand for the good in the population. Absent heterogeneity in the consumption of X, all types lose from the shift toward in-kind provision. More generally, a marginal shift toward in-kind provision makes *all* types worse off (a Pareto worsening) if heterogeneity in the consumption of X is small enough relative to the distortion caused by the subsidy that any redistribution gain enjoyed by high-X types is more than offset by the reduction in the cash benefit due to the distortion.

2.2.2 The net ex-ante value of a shift toward in-kind provision

Ex-ante expected utility is

$$\max_{\sigma} EU(\sigma) = \int_{\Theta} v(p(\sigma), m(\sigma, B); \theta) f(\theta) d\theta.$$

The total derivative of expected utility with respect to the in-kind component σ (adjusting the cash component b in order to hold fixed total program spending) is:

$$\begin{aligned}\frac{dEU(\sigma)}{d\sigma} &= \int_{\Theta} \frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma} f(\theta) d\theta = \int_{\Theta} \lambda(\sigma; \theta) \frac{dV(\sigma; \theta)}{d\sigma} f(\theta) d\theta = E_{\Theta}\left(\lambda(\sigma; \theta) \frac{dV(\sigma; \theta)}{d\sigma}\right) \\ &= Cov_{\Theta}\left[\lambda(\sigma; \theta),\ x_X(\sigma; \theta) p_X^0\right] - (\sigma p_X^0) E_{\Theta}\left(\lambda(\sigma; \theta)\right) E_{\Theta}\left(\frac{dx_X(\sigma; \theta)}{d\sigma}\right), \quad (2)\end{aligned}$$

where $\lambda(\sigma; \theta)$ is the marginal utility of income.

Equation 2 shows the key roles of heterogeneity in the level of demand for X and the sensitivity of the demand for X to the composition of benefits in determining the welfare consequences of in-kind provision. The extent of heterogeneity in the demand for X and the extent to which this heterogeneity is correlated with marginal utility together determine the targeting benefit of in-kind provision. The greater is the covariance across types in marginal utility and the demand for X, the greater is the targeting benefit of in-kind provision. The sensitivity of the demand for X to the composition of benefits determines the distortion cost of in-kind provision. The greater is the sensitivity of the demand for X to the composition of benefits, the greater is the distortion cost of in-kind provision.[13]

[13] Appendix A analyzes the optimal mix of in-kind and cash benefits. Absent heterogeneity in the demand for X, the optimal policy is a pure cash benefit with no subsidy on X, ($b = B, \sigma = 0$). Absent any consumption distortion, the in-kind benefit simply redistributes resources across different types (as defined by their level of demand for X), at no efficiency cost. In this case, the optimal policy eliminates the covariance between marginal utility and the demand for X. If the demand for X is at least somewhat elastic, by contrast, the optimal policy trades off the insurance benefit of increasing in-kind provision against the distortion cost. In most cases it will stop short of eliminating the covariance between marginal utility and the demand for

3 Long-Term Care Risk and Financing and the Design of Home Care Benefits

Chronic health problems are the source of one of the most important risks people face over the life cycle. Most people will at some point develop functional limitations and require significant, time-intensive assistance with activities such as bathing, eating, dressing, and managing their household. Roughly 15 percent of Americans over age 50 have at least one functional limitation and at least one person helping them as a result.[14] People deal with their functional limitations in a wide variety of ways. Among people receiving help, the vast majority (87 percent) live in the community; only 13 percent live in nursing facilities, mostly nursing homes. Among those living in the community, most (72 percent) receive only informal care from family and friends, 10 percent receive a mix of formal and informal care, and only 4 percent rely exclusively on formal home care. (We will refer to formal home care as "formal care" and to informal home care as "informal care.") Of all hours care, 64 percent are informal care, 10 percent are formal home care, and 27 percent are from care-giving facilities. The costs of such care are quite large. In 2008, total spending on formal long-term care was $203 billion, with one-third going to home and community based services (Centers for Medicare and Medicaid Services, 2010). The cost of informal care is estimated to be no less than 60 percent of, and is often thought to exceed, the cost of formal care (e.g., Arno et al., 1999).

Although the costs of dealing with chronic health problems are highly uncertain and sometimes quite large, much of these costs are uninsured. In many countries, universal health insurance programs (like Medicare in the U.S.) do not cover long-term care. Yet markets for private long-term care insurance are uniformly small. In the U.S., just 10 percent of people 65 and older own private long-term care insurance, and coverage rates are similarly low in other countries. The large costs of long-term care together with the low rates of insurance leads to a situation in which means-tested government programs play a major role in financing long-term care costs. In the U.S., for example, the means-tested Medicaid program covers at least part of the costs of 70 percent of nursing home residents at any point in time (Kaiser Commission on Medicaid and the Uninsured, 2013).

Medicaid also plays a major role in financing home care. In most states, Medicaid provides home care primarily through two programs: the Medicaid Title XIX PCS optional State plan and the Medicaid 1915(c) HCBS waiver program. Eligibility is based primarily upon financial means tests and an assessment of the individual's "need" for home care based on

X, since at the margin there would be only a distortion cost and no targeting benefit.

[14] This and the other statistics about care-giving arrangements in this paragraph come from Barczyk and Kredler (2016).

her health.[15] Medicaid home care programs have grown rapidly in recent years. In 1999, 1.9 million people received home care through Medicaid; by 2013, that number had increased to nearly 3 million. In addition to the growing number of participants, the fraction of Medicaid long-term care dollars that go to home care has risen from 18 percent in 1995 to 51 percent in 2014 (Kaiser Commission on Medicaid and the Uninsured, 2016).

Traditionally, government and private long-term care insurance programs have provided in-kind formal care benefits (including subsidies on formal care) to people who meet health- and sometimes financial-related eligibility criteria.[16] But in recognition of the importance of informal care and other means of dealing with chronic health problems, many programs have shifted toward more flexible, cash-like benefits. These programs tend to allow people to spend their benefits on a much wider array of options than just the traditional formal care, often including paying family and friends for informal care, and at the extreme including anything that a recipient wishes to purchase, whether care-related or not. Germany, France, Italy, Austria, Sweden, and the Netherlands all have long-term care programs that either pay benefits in cash or allow recipients to choose between cash and in-kind benefits (Da Roit and Le Bihan, 2010). In the U.S., many state Medicaid home care programs provide flexible, cash-like benefits, including 15 states with "Cash and Counseling" programs (Doty et al., 2010). Early versions of the Affordable Care Act included a voluntary, long-term care insurance program in which individuals meeting health-related eligibility criteria would have received a daily cash benefit.

The shift toward more flexible, cash-like benefits is motivated by the perception, supported by a range of evidence, that many people with chronic health problems would much prefer a cash transfer to a cost-equivalent in-kind transfer of formal care. An important milestone in the debate about more- vs. less- flexible benefits was the Cash and Counseling demonstrations. These were large-scale experiments run by three state Medicaid programs around the year 2000. The experiments randomized participants between the traditional in-kind formal home care benefit and an alternative more flexible, cash-like benefit that allowed recipients to spend the benefit on a wide range of care-related goods and services, including paying relatives or friends for providing care.

[15] Medicaid is financed jointly by the federal and state governments, so Medicaid policies vary somewhat across states. For the elderly, the means tests for Medicaid home care are often less restrictive than those for general Medicaid coverage. The majority of states provide coverage for individuals with incomes up to 300 percent of the monthly Supplemental Security Income amount (LeBlanc et al., 2001). The amount of Medicaid home care for which an individual qualifies is determined by a medical exam. The applicant's health care provider must submit a care plan that details the services deemed appropriate based on the applicant's health status. Summaries of Medicaid-provided home care services are available in LeBlanc et al. (2001) and Kaiser Commission on Medicaid and the Uninsured (2011).

[16] Many programs, for example, offer benefits to people who have limitations with at least two activities of daily living.

The main goal of the Cash and Counseling experiments was to test whether recipients could effectively manage their cash benefits and receive "enough" care. The results were almost uniformly positive. Members of the cash-benefit treatment group reported greater satisfaction with their care (Foster et al., 2003) and had similar health outcomes; if anything, their health outcomes appeared to be slightly better than those of the traditional in-kind benefit group (Lepidus Carlson et al., 2007). In the official final report on the experiments, Brown et al. (2007) conclude that the cash transfer has overwhelmingly positive effects on recipients. The large body of work evaluating the Cash and Counseling experiments demonstrates quite conclusively that more flexible, cash-like benefits can be much more valuable to recipients than traditional in-kind benefits. This points to an important cost of providing home care benefits in kind: The consumption distortion from in-kind provision appears to be large in this context.

Although much is known about the consumption distortion from in-kind provision of Medicaid home care benefits, little is known about the potential benefits of in-kind provision, whether for Medicaid home care or for other programs more generally (Currie and Gahvari, 2008). In-kind provision can have a variety of effects in addition to distorting consumption, including changing the distribution of benefits among those eligible for benefits. This could significantly increase targeting efficiency, especially in the many contexts in which state verification is particularly costly. Yet little is known about the likely magnitude of such benefits. This is the gap we aim to fill in this paper.

In the context of long-term care, the diversity of ways in which people cope with chronic health problems and the large role of informal care may raise important challenges for insuring this risk. Hard-to-verify heterogeneity in the costs of coping with chronic health problems—which may arise from hard-to-verify differences in health status or in would-be care-givers' opportunity costs of providing informal care—seem likely to cause heterogeneity in marginal utility and in the demand for formal care. In a second-best setting in which this heterogeneity is not verifiable, insurance programs cannot make different cash transfers to people on the basis of their coping costs. Any redistribution from people with lower to higher coping costs must instead come through changes in prices. A potentially worthwhile distortion might be to subsidize formal care. Everyone in the same market faces the same price of formal care, which makes formal care consumption a potentially good "meter" of the total costs of bad health (of which formal care is only a part). A formal care subsidy therefore tends to redistribute to people with larger coping costs at the cost of distorting consumption. As discussed in Section 2, two of the main determinants of the desirability of such a subsidy are the distribution of formal care consumption and the sensitivity of formal care demand to the composition of benefits. We now turn to estimating these objects.

4 Value to Recipients of Medicaid In-Kind Home Care

Estimations of the demand for formal care confront an especially challenging simultaneity problem. As always, estimating the slope of demand requires an instrument that shifts supply but not demand. The key challenge in the case of formal care is that many factors that shift supply are also likely to shift demand by changing the opportunity cost of informal care. For example, consider using minimum wage laws (or their changes over time) as instruments for the price of formal care. Many formal home care workers earn roughly the minimum wage, so changes in the minimum wage are likely to shift the supply of formal care. But at the same time, changes in the minimum wage are also likely to change the opportunity cost of informal care-giving by changing the wage or employment prospects of some potential informal care-givers. Changes in the supply of informal care likely shift the demand for formal care since formal care and informal care are closely-related goods. The close links between formal care and informal care make it especially difficult to find valid instruments for the supply of formal care.

We circumvent this identification issue by taking advantage of the randomization in the Cash and Counseling experiments. People randomized to the cash benefit face the market price of formal care, whereas people randomized to the traditional Medicaid program face a price of zero.[17] Moreover, the price variation from the experiment is not only exogenous to the demand for formal care, it also spans the full range of prices most relevant for policy, from the market price down to zero.

The Cash and Counseling experiments were conducted by the Medicaid programs of Arkansas, Florida, and New Jersey. Participants were enrolled beginning in 1998 in Arkansas, in 1999 in New Jersey, and in 2000 in Florida. In New Jersey and Florida, only people who had been receiving Medicaid home care were eligible for the experiment, whereas in Arkansas anyone who was eligible for Medicaid home care could enroll. Although the exact implementations of the experiment differed somewhat between the three states, each shared the same basic features. After completing a baseline survey, participants were randomly assigned to one of two equal-sized groups. Members of the control group remained subject to the rules of the traditional, in-kind Medicaid home care program. Members of the treatment group were given cash budgets that could be spent on a wide range of personal care goods and services, including assistive devices, home modifications, and, most important, informal care from family or friends. Although the benefit had to be spent on personal care services, for many recipients it was likely to be equivalent to a pure cash transfer, since the vast majority of

[17] Appendix B.2 presents evidence that the marginal value of Medicaid formal care is zero for most recipients of the traditional in-kind benefit. It also tests the robustness of the estimated price sensitivity to alternative assumptions about the marginal value of Medicaid formal care.

recipients had been receiving enough informal care at baseline to more than exhaust the budget they received from Medicaid.[18]

We use data from the nine-month follow-up survey for respondents who were at least 65 years old and who had non-missing data on their demographics (age, gender, race, and educational attainment) and self-rated health. Our sample contains 860 participants from Arkansas (428 in the cash group), 482 in Florida (245 in the cash group), and 604 in New Jersey (312 in the cash group). Table 1 provides summary statistics. At baseline, average formal care consumption ranged from 8 (Arkansas) to 16 (New Jersey) hours per week, and on average participants had two informal caregivers. The remaining rows of Table 1 show that the treatment and control groups are balanced on demographics such as gender, age, race, and education. Although non-negligible fractions of the treatment and control group attrited from the experiment before the nine-month follow-up survey (20 and 35 percent of treatment and control group members, respectively), the attrition did not lead to any apparent imbalance between the treatment and control groups in terms of observables. Of the 30 balance tests, none of the differences between treatment and control groups are statistically distinguishable from zero at the 5 percent level and only two are significant at the 10 percent level. These are fewer significant differences than would be expected to arise by chance without any differential attrition.

A simple measure of the sensitivity of the demand for formal care to the composition of benefits is the difference in average formal care consumption by people randomized into cash versus in-kind benefits. Table 3 makes this comparison. On average, people randomized to in-kind benefits consumed over twice as much formal care as people randomized to cash benefits: 14.8 versus 7.1 hours per week. Similarly large differences occured in each state. Figure 2 shows the empirical cumulative distribution functions of formal care consumption by the in-kind and cash groups. The distributions have two key features. First, formal care consumption by the in-kind group is greater throughout the distribution. Second, many people consume no formal care, especially in the cash group. Whereas about one-fifth of the in-kind group consumes no formal care, over half of the cash group consumes no formal care. Receiving cash as opposed to in-kind benefits leads many people to a corner solution in which they consume no formal care. The prevalence of people consuming no care means that the observed mean differences between the two groups tend to understate the price sensitivity

[18] Recipients had to submit receipts documenting that they spent their benefits on personal care services, though they were also allowed to spend up to 10 percent of their allowance on services that could not be readily invoiced (e.g., payments to a neighbor for mowing the lawn). The cash allowances received by members of the treatment group were slightly smaller than those that would have enabled them to purchase the care they would have received under the traditional Medicaid program. This is mainly because of a requirement that the experimental cash treatment be budget-neutral. This requirement meant that the costs of paying the counselors who helped treatment group members manage their care came out of the cash allowances.

of demand. Overall, these simple comparisons suggest that the demand for formal care is quite sensitive to the composition of benefits.

The slope of the demand curve is not simply the ratio of the observed difference in average quantities to the difference in prices for two reasons. First, as just discussed, many individuals consume zero hours of care. We account for this by treating an individual's observed hours of care, q_i, as the outcome of a censored, latent demand for care, $q_i = \max\{0,\ q_i^*\}$. Second, some of the participants in the cash group reverted to traditional Medicaid home care, and some of the participants in both groups left Medicaid home care altogether. The differences in Table 3 therefore correspond to intent-to-treat parameters rather than direct measures of price sensitivity. We handle this issue by instrumenting for the price each participant faces with her randomized assignment.[19] We estimate the system

$$q_i^* = \alpha + \beta p_i + X_i\gamma + \varepsilon_i$$

$$q_i = \max\{0,\ q_i^*\}$$

$$p_i = \mu_0 + \mu_1 Cash_i + X_i\mu_2 + \nu_i,$$

where p_i is the price of formal care, $Cash_i$ is an indicator of whether the participant was randomized to the cash treatment, and X_i includes indicators for gender, education level, race, self-rated health, five-year age bins, and state. We begin by assuming (ε_i, ν_i) are jointly normal and estimate this system using an instrumental variables Tobit specification.

The first-stage results are presented in Table 4. On average, being assigned to the cash group increases the price of formal care by \$8.84, 64 percent of the average market price of \$13.73. The point estimate is fairly precise and the first-stage relationship is strong; the F-statistic exceeds 1,000, well above the levels at which weak instruments become a concern (Stock and Yogo, 2002). As one would expect given the random assignment, adding control variables has little effect on the estimated relationship between treatment assignment and the price of care (column (2)).

The instrumental variables estimate of β is presented in Table 5. The estimate implies that a one dollar increase in the hourly price of formal care reduces consumption by 1.8 hours per week. Evaluated at the sample means, this implies an elasticity near -1.2. When the demographic controls are added to the specification (column (2)), the estimate is virtually unchanged.

This price sensitivity suggests that formal care subsidies could potentially significantly in-

[19] For individuals who leave Medicaid home care or are in the cash group, the price they face is the market price in their state. For individuals who are in the traditional Medicaid home care program, the price they face is zero.

crease consumption of formal care, as many people who otherwise without subsidies would have relied more heavily on informal care substitute toward formal care. This also suggests that subsidizing formal care may have large moral hazard costs. For someone consuming the average amount of formal care among people randomized to traditional in-kind benefits (14 hours per week), the ex-post deadweight loss from in-kind provision is about 72 percent of Medicaid's spending on the individual's care.[20] For someone consuming 25 hours of formal care per week, the deadweight loss from in-kind provision is about 49 percent of Medicaid's spending on the individual's care. More generally, and as discussed in Section 2, the ex-post deadweight loss as a fraction of Medicaid spending is decreasing in the level of demand for care.

In Appendix B.2, we analyze the consequences for the estimated price sensitivity of making a wide range of alternative assumptions about the distribution of the error terms and limits on how much Medicaid home care recipients of the traditional in-kind benefit can consume. The conclusion that the demand for formal care is quite sensitive to its price is robust to a wide range of alternative assumptions.

Also in Appendix B.2, we discuss the generalizability of the estimates of the demand for formal care based on the Cash and Counseling experiments to other policies and populations of interest. There are two key issues. First, selection into the experiment likely means that the average participant in Cash and Counseling was more price sensitive than the average recipient of Medicaid home care and potentially either more or less price sensitive than the average person in the population of people eligible for Medicaid home care. Second, the nature of the experiment—especially its unexpected nature and uncertain duration—seems likely to lead to smaller quantity responses than one would expect from permanent changes in policies. Although these are important issues for generalizability, the benefits of the large, exogenous variation in the price of formal care make the Cash and Counseling experiments an important source of evidence on the demand for formal care. In light of the possible issues with generalizability, however, in the policy analysis (Section 5) we test the robustness of our results to a wide range of values of the price sensitivity.

5 Targeting Effects of In-Kind Provision of Home Care

In this section, we characterize the extent of the risk associated with formal care and assess whether in-kind provision of home care targets transfers. We do this using data from the Cash

[20] With a price sensitivity of -1.8, someone consuming 14 hours of care per week has an equivalent variation of formal care benefits (assuming no income effects) of $54 per week, or 28 percent of Medicaid's $192 of spending on that care (14 hours per week at an average price of $13.73 per hour).

and Counseling experiments as well as the 1999 wave of the National Long Term Care Survey (NLTCS). The NLTCS is a nationally representative sample of Medicare enrollees, including people in care-giving institutions such as nursing homes. It is designed to provide information on the health and functional status of Americans who are at least 65 years old. We focus on people who are living in the community with at least two ADLs because this is the population of interest for home care benefits programs. The data include information on formal care consumption, income, physical limitations, and a host of demographic characteristics.

Heterogeneity in formal care consumption implies, holding other resources fixed, heterogeneity in non-care consumption, which in many models implies heterogeneity in the marginal utility of wealth. Figure 3 shows the empirical CDF and PDF of formal care hours from the NLTCS. There are two important features of this distribution. First, many people do not consume any formal care—approximately 62 percent of the sample. Second, there is a long right tail of people who consume large amounts of care. Conditional on consuming any care, median consumption is 14 hours per week, whereas the 95th percentile is 168 hours per week. If purchased at market prices, the latter amount of care would cost more than $122,000 per year.[21]

In-kind transfers can target benefits on both the extensive and intensive margins. To the extent that taking up benefits is costly, people with relatively low demand may not choose to join the program. If that occurs, then there are fewer individuals over which to spread a program's budget and those who do sign up receive greater transfers. Take-up rates are therefore an important metric of the extent to which providing benefits in kind targets transfers to a subset of individuals. We use the NLTCS to estimate the fraction of people eligible for Medicaid home care who actually take up benefits. Among those eligible, we find that between 4 percent and 16 percent actually take up benefits.[22] The low take-up rate

[21]If we do not limit the sample to those with at least two ADLs, we continue to find a large amount of dispersion in formal care use. Eighty-six percent of individuals do not consume any care. Conditional on consuming care, mean consumption is 8 hours per week and the 95th percentile is 141 hours per week. Again, the upper-tail risk is considerable. To purchase that level of care in the market would cost more than $102,000 per year.

[22]Estimating take-up rates for Medicaid home care, and means-tested social programs more generally, is notoriously difficult (U.S. Department of Health and Human Services, 1992; Currie, 2006). Eligibility rules are complex, vary from state-to-state, and often depend upon household characteristics that are unobservable to the researcher. We use the NLTCS to estimate the fraction of the elderly who are eligible for benefits, based on the eligibility criteria from Schneider et al. (1999). We combine that estimate with the size of the sixty-five and older population and administrative estimates of the number of Medicaid home care users from LeBlanc et al. (2001). The main source of uncertainty in our estimated take-up rate is the incompleteness of the information on household assets in the NLTCS. In all cases, a person must have at least two activities of daily living limitations. The upper end of our range, 16 percent take-up, should be interpreted as an upper bound on the take-up rate because we imposed (much) more restrictive income and asset requirements than the actual limits in the vast majority of states. In particular, we imposed that the household earned no more than 100 percent of the SSI benefit and had no cars (car value is one of the primary inputs to the asset tests). Our least restrictive eligibility threshold uses the income limits from Schneider et al. (1999) and imposes that the household have fewer than two cars.

means that, holding fixed total spending on the program, benefits per individual who takes up are between 6 and 24 times greater than they would be under a pure cash program with 100 percent take up.

To incorporate targeting on the intensive margin, we compare the in-kind and near-cash groups in the Arkansas Cash and Counseling experiment. We are limited to this single state because it is the only one for which we have a measure of expenditures on recipients in both the in-kind and near-cash treatments. To the extent that in-kind provision targets benefits more than the near-cash benefit does, there should be greater dispersion in benefits received in the in-kind than the near-cash group. Figure 4a shows the kernel density plots of expenditures on members of the in-kind and near-cash groups.[23] The in-kind group has substantially more mass than the near-cash group at both low and high benefit levels. Less than 14 percent of the near-cash group received less than $28 of expenditures per week (the approximate price of two hours of formal care) while more than 35 percent of the in-kind group received less than this amount. At the other tail, 24 percent of the in-kind group received more than the 90th percentile of expenditures received by the near-cash group. The in-kind group's distribution has 10 times the variance of the cash group's distribution. Figure 4b shows the difference in benefits between the in-kind group and either the near-cash or a pure cash benefit at different quantiles of the distribution.[24] These distributions suggest that in-kind provision of home care targets benefits to a greater degree than does the near-cash transfer in the Cash and Counseling Demonstrations.[25]

Whether this targeting of benefits is desirable depends on who is being targeted and the planner's objective function. Table 6 shows summary statistics from the NLTCS on the characteristics of the population that is eligible for Medicaid home care divided into those who do and those who do not actually take up the in-kind benefits. Individuals who take up the program are 13 percentage points more likely to have at least 4 ADLs, 12 percentage points more likely to say their health is fair or poor (instead of good or excellent), 19 percentage points more likely to be unmarried, and 20 percentage points more likely to be living alone. Even after conditioning on the observables used by Medicaid to determine eligibility, those who actually take up Medicaid home care benefits appear to be in worse health and have less good informal care options than eligibles who do not take up benefits.

[23]For the near-cash group, we can not directly observe the dollar transfer made to the participants. However, because we have individuals care plans, we can estimate benefits by multiplying hours of care by the cost to Medicaid of providing those hours. For the in-kind group, we multiply hours of care received by Medicaid's cost per hour.

[24]Note that a pure cash transfer is a hypothetical, untagged cash transfer to all Arkansas Cash and Counseling participants. Keeping the budget fixed, this implies a transfer of approximately $122 per week to each participant.

[25]It is worth noting that the near-cash transfer also targets resources more than a pure cash transfer would.

Appendix C provides additional evidence from the Arkansas Cash and Counseling experiment on who is targeted by in-kind provision of home care. Despite the fact that the tagging of both the near-cash and in-kind benefits minimizes the differences in targeting on *observables* across the two groups, we find suggestive evidence consistent with our findings from the NLTCS. Those who are sicker or who have less good informal care options tend to be targeted by the in-kind benefit. Although we can not observe the marginal utility of income directly, it seems likely that the sicker and those with fewer informal care options are likely to be those with greater marginal utility and thus that in-kind provision helps to insure risk.

If the states targeted by in-kind formal care are those in which individuals have higher marginal utility, then in-kind provision provides insurance value. However, given the large cost associated with in-kind provision, there could be alternative ways to target resources that are less costly. A natural candidate is a tagged cash transfer. As we saw in Table 6, there are a number of characteristics that predict take-up of Medicaid home care that could potentially be used as tags. However, the targeting properties of potential tags depend on the degree to which observable characteristics explain the variation in the risk.

We assess the degree to which observable characteristics are able to predict formal care use. We focus primarily on the NLTCS population rather than the Cash and Counseling participants because the former more closely represents the population risk. We follow the standard approach from the predictive modeling literature: randomly split the sample into two subsamples, train the model on one subsample (the training sample), use the trained model to make predictions for the other subsample (the test sample), and assess predictive power in the test sample. We implement two separate modeling approaches: OLS and machine learning. The latter is likely to be near the frontier of methods with which an insurer could predict risks and form tags and so helps to establish an upper bound on how well formal care use can be predicted. The variables we use to predict formal care use are age, indicators for the number of ADLs, self-rated health, gender, being married, having children, and linear and quadratic terms in income. Because those on Medicaid face a lower price for formal care, we adjust observed hours of formal care for differences in prices.[26]

Table 7 reports the results from this exercise. In column (1), only age, the number of ADLs, self-rated health, and gender are used ("health" controls). The first row presents the R^2 based on predicting outcomes in the same data used to train the model. This in-sample R^2 corresponds to the R^2 reported in most regression analyses and will tend to overstate the

[26]For individuals who face a price of zero, i.e., those on Medicaid, we adjust their observed quantities using our estimated price sensitivity to predict their consumption when facing the market price. In particular, we multiply the market price for care by 1.8 and subtract that quantity from their observed hours of formal care. To be comparable with those who currently face the market price for care, we censor hours of care at zero for anyone whose adjusted hours of care are negative.

true amount of variation explained due to over-fitting.[27] Indeed, it is more than twice as large as the R^2 obtained in the test data. In the final row of Table 7 we present results for the machine learning model.[28] Although the machine learning algorithm does perform better than simple OLS, it is not a dramatic improvement. Columns (2) and (3) add variables that the models can use to predict formal care. In each case, predictive power increases, but the vast majority of variance in formal care is not accounted for by these observables.

An alternative, costlier means of targeting is case-by-case examinations such as those commonly used by disability insurance programs. Medicaid home care programs require participants to have a medical exam with a physician or nurses who will create a care plan meant to reflect the recipients "need for care. Intersetingly, "need" is meant to include the availability of paid and unpaid caregivers Dale et al. (2004). In principle, it should be as close to a summary measure of demand upon which an insurer could potentially contract as is feasible. In column (4) of Table 7, we use data from the Arkansas Cash and Counseling experiment to assess how much variation in formal care use can be accounted for by an individuals care plan.[29] Among those using the in-kind benefit at the nine month follow up survey, care plan hours account for less than 7 percent of the variation in individuals formal care use.[30] Taken together, it does not appear that much of the variation in formal care use is predictable with the types of observables that insurers would have access to. As a result, it is not likely that even richly-tagged pure-cash transfers would prove as useful in targeting benefits to high marginal utility individuals as in-kind provision.

6 Welfare Analysis

In this section we adapt an approach from the literature on optimal taxation to evaluate the welfare consequences of alternative home care benefit programs.

[27] See Mullainathan and Spiess (2017) for a review of this issue and the recent rise of machine learning in economics.

[28] See James et al. (2013) for a discussion of machine learning techniques. These methods are becoming increasingly widespread in prediction problems in economics (Mullainathan and Spiess, 2017). We implement random forest models that use five-fold cross-validation to tune the model. Dimensions that are optimized include the number of trees, tree depth, minimum leaf size, and number of variables to (randomly) sample at each split. We have also experimented with two different aggregation methods, bagging and boosting.

[29] Arkansas is the only state for which we have information on care plans.

[30] Care plans are updated every six months in Arkansas. We use two measures of care plan hours in the analysis. The first is care plan hours at baseline; the second is care plan hours twelve months after baseline. Because care use is measured at nine months, there is some question about differences in measured care plan hours and the individuals true care plan hours at that time. The lack of variation in care plans over the twelve month period—the correlation between an individuals care plan at baseline and follow up is 0.86—allays this concern to some degree.

6.1 Model

Individuals draw their type from a distribution of types, $\theta \sim F(\theta)$. Then they choose their formal care consumption and non-care consumption to maximize utility subject to a budget constraint that depends on the policy in operation. The budget constraint is

$$A + pF = m,$$

where F is formal care consumption, A is non-care consumption (i.e., "all other goods," the numeraire), p is the after-subsidy price of formal care, and m is benefit-inclusive nominal income. The utility function and the corresponding demand for care are

$$U(A, F; \theta) = u\left(A - \frac{(\max\{\alpha, 0\} - F)^2}{2\beta}\right),$$

$$F(p, m; \theta) = \max\left\{0,\ \min\left\{m/p,\ \alpha - \beta p\right\}\right\}.$$

α is the quantity of care at which the individual is satiated, i.e., the amount of care the individual would consume when facing a price of zero. β determines the utility cost of consuming levels of care other than the satiation level and thereby determines the sensitivity of the demand for formal care to the composition of benefits. $F(p, m; \theta)$ is the Marshallian demand function for formal care.

This utility function is motivated by key evidence from our setting. It produces a simple function for the demand for formal care that is consistent with the sensitivity of formal care consumption to its price and that people become satiated at finite levels of formal care consumption.[31] This utility function also has several appealing features. It nests as a special case the widely-used model in which health spending is equivalent to a wealth shock.[32] It implies that the demand for formal care is linear in its price within the range of prices in which the individual is not at a corner. It has an intuitive interpretation: Utility is decreasing in any unmet, residual health needs, $(\alpha - F)$, the size of which is decreasing in

[31] The most direct evidence of satiation is that among Cash and Counseling participants with information on their care plan hours, 43 percent consumed less care than they were entitled to based on their care plan. Intuitively, satiation might arise from a demand for privacy or space, since home care involves close contact with caregivers in one's home. This utility function is also consistent with the fact that most people who need assistance do not consume any formal care. This implies that there is no Inada condition on formal care consumption and that formal care is not too complementary with other goods that people consume.

[32] As β approaches 0, formal care consumption approaches α ($F(p, m; \theta) \to \alpha$, ignoring corner solutions), and the indirect utility function approaches $v(p, m; \theta) = u(m - p\alpha)$. For $\beta > 0$, demand for formal care is sensitive to its price and the indirect utility function is $v(p, m; \theta) = \begin{cases} u\left(m - \frac{\max\{\alpha,0\}^2}{2\beta}\right), & \text{if } \alpha < \beta p; \\ u\left(m - p(\alpha - \beta p) - \frac{\beta p^2}{2}\right), & \text{if } \alpha \geq \beta p. \end{cases}$ This differs from the benchmark case in which health spending is a wealth shock by just a slight adjustment, which is necessary to accommodate the observed price sensitivity of demand for formal care.

formal care consumption, F, and increasing in the level of demand for formal care, α. This captures the idea that certain health problems are costly for people to cope with on their own. Marginal utility of income depends on the demand for formal care mainly through the budget constraint: Greater spending on formal care means lower non-care consumption.

6.2 Baseline parameter values

The key parameters of the model are the sensitivity of formal care demand to its price, β, and the distribution of the level of demand for formal care in the population of people eligible for home care benefits, $F(\alpha)$. Everyone has the same price sensitivity of demand for care, β, equal to our main estimate from the Cash and Counseling experiment. We use this β to convert the joint distribution of formal care consumption and formal care prices observed in the NLTCS into a distribution of the level of demand for formal care, $F(\alpha)$. For the main analysis, which takes as given standard eligibility criteria for home care benefits, our sample is everyone aged 65 and older with at least two activities of daily living limitations. For the tags analysis, we estimate separate $F(\alpha)$ distributions for each sub-group of this population as defined by their tagged characteristics (e.g., for people with different numbers of activities of daily living limitations). Estimating $F(\alpha)$ would be entirely straightforward were it not for people who consume no care when facing a positive price. For the 62 percent of the population of interest who consume no formal care, however, revealed-preference analysis only bounds the level of their demand: their marginal value at zero hours of care is no greater than the price. But because we will be analyzing policies that reduce the prices people face, it is important to know at which price each individual would begin purchasing care. We handle this fundamental unobservability issue by extrapolating the observed distribution among people who consume a strictly positive amount of care backward to "fill in" the unobservable α values of people who consume no formal care when facing a positive price. Details of this calculation are reported in Appendix D.

Figure 5 presents our main estimate of the distribution of the level of demand for formal care, $F(\alpha)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits a long right tail (the mean far exceeds the median) and that most of the mass is at low values.

The remaining parameters take standard values. We follow most of the literature on health spending risks and use a constant relative risk aversion utility function, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ (e.g., Brown and Finkelstein, 2008; De Nardi et al., 2010; Ameriks et al., 2011). In our model, the argument c is "net consumption," non-care consumption net of any residual coping costs, $c = A - \frac{(\alpha-F)^2}{2\beta}$. We follow Brown and Finkelstein (2008) and many others in taking as a

baseline value a coefficient of relative risk aversion, γ, of three. Income before transfers, m, is \$15,000 per year. The distribution of before-subsidy prices of formal care is the empirical distribution observed in the NLTCS. People who cannot achieve net consumption of at least $\bar{c}$ =\$5,000 per year receive transfers that enable them to enjoy net consumption of \$5,000 per year (a consumption floor). This is meant to approximate the important means-tested programs, Medicaid and Supplemental Security Income.

With these parameters, the risk within the set of people traditionally eligible for home care benefits (with two or more ADL limitations) is substantial. In order to make the individual as well off as she is with the first-best policy under an alternative pure-cash benefit program, the cash benefit would have to be about \$9,377—137 percent—greater than the average cost of the first-best program.

6.3 Welfare effects of in-kind provision

In this section we calculate the welfare effects of varying degrees of in-kind provision, taking as given total spending on program benefits and standard eligibility criteria for home care benefits. Following standard practice for Medicaid home care and private long-term care insurance, we focus on programs that limit eligibility to people with two or more activities of daily living limitations. We consider policies under which total program spending equals the spending on a pure in-kind benefit program, a 100 percent subsidy with no cash benefit. Policies with smaller subsidy rates have larger cash benefits.

Figure 6 summarizes the key results. It shows the equivalent variation of the mixed in-kind and cash benefit policy as a function of the in-kind component, the subsidy rate σ. The optimal subsidy rate is 88 percent, close to a pure in-kind program (under which the after-subsidy price is zero). The optimal subsidy increases welfare substantially relative to a pure-cash benefit program. In order to make the individual as well off as she is with the optimal policy under an alternative pure-cash benefit program, the cash benefit would have to be about 80 percent greater than the average cost of the in-kind program. Figure 6 also shows, however, that the optimal subsidy is significantly less valuable than the hypothetical first-best policy. The optimal in-kind subsidy achieves 59 percent of the incremental value over a pure-cash benefit that the first-best policy does.

Table 8 shows a variety of outcomes in several versions of the model. The purpose of this table is to provide intuition for and assess the robustness of the key results. The first column of the table shows results for the baseline specification just discussed. The key tradeoff involved in increasing the in-kind component of the benefit can be seen clearly by comparing the average level of and dispersion in non-care consumption under the optimal

subsidy program and under the cost-equivalent pure-cash benefit program. The optimal subsidy reduces average non-care consumption due to the consumption distortion (and to a lesser extent due to foregone transfers from the consumption floor), but it also greatly reduces the dispersion of non-care consumption, as measured by the standard deviation. Under the pure-cash program the standard deviation of annual non-care consumption is 4.5 times greater than under the optimal program, $5,610 vs. $1,237.

Additional rows of the table unpack these results further. They show that formal care consumption is significantly greater under the optimal subsidy than in the absence of any program, by a factor of 2.4. This translates into a large distortion cost; the total ex-post equivalent variation of the optimal program summed over all recipients is only 48 percent of the total cost of the program. Part of this is due to the optimal program displacing transfers from the consumption floor, but much of it is due to the consumption distortion from the formal care subsidy. This amounts to a significant implicit tax on insurance, equivalent to a tax of almost 100 percent of benefits. The reason that subsidizing formal care is optimal despite the large distortion is that the in-kind subsidy redistributes toward people with greater marginal utility. The correlation between an individual's marginal utility in the absence of any program and his ex-post equivalent variation of benefits under the optimal program is 0.84. The net benefit from in-kind provision comes from making large transfers to the relatively few people with high demand for care (and so low non-care consumption). This can be seen in the bottom row of the table, which shows that under the optimal subsidy program only 16 percent of people value the subsidy program as much as the cost-equivalent pure-cash benefit program ex post. This may help explain why many countries and U.S. states have made home care benefits more cash-like. Making benefits more cash-like helps most individuals ex post, often significantly. A key finding of this paper, however, is that the greater ex-post value of more cash-like benefits comes at the expense of much less redistribution toward people with high demand for formal care, which may worsen insurance.

The other columns of the table test the robustness of the results to making different assumptions about the key ingredients of the model. The price sensitivity of demand for formal care must be quite large—over 10 times larger than we estimate based on evidence from the Cash and Counseling experiment—in order to overturn the conclusion that the optimal subsidy is large. Even if the distribution of partially-identified α values among people who consume no care at positive prices is in the "worst-case" configuration (i.e., each α_i equal to the maximum value consistent with observed behavior), the optimal subsidy rate is still 86 percent. The utility function must exhibit strong state dependence of just the right kind—greatly decreasing the marginal utility of people with high demand for formal care in just the right way—in order to overcome the fact that, holding other resources constant, people with greater formal care consumption must have lower non-care consumption. Although the

right tail of the distribution of demand for formal care is an important determinant of the targeting benefit and so the optimal subsidy, the optimal subsidy remains large even when the right tail of the distribution is chopped off or when all of the α values are scaled down. If people who consume more than 50 hours per week of care are dropped from the analysis, the optimal subsidy is 59 percent. If all of the α values are cut in half, the optimal subsidy is 75 percent. Finally, a combination of relatively low risk aversion together with a relatively generous consumption floor can overturn the optimality of a large subsidy on formal care, although this reflects the undesirability of *any* insurance—including a first-best contract—in situations in which means-tested programs are sufficiently attractive rather than any undesirability of in-kind benefits per se. Appendix D discusses these robustness tests in more detail.

The alternative specifications also provide information about the key factors driving the results. As expected, the net benefit of subsidizing formal care is decreasing in the price sensitivity of demand for formal care. When demand for formal care is completely inelastic ($\beta = 0$), a 100 percent subsidy achieves the first best.[33] The targeting benefit of in-kind provision is increasing in risk aversion and decreasing in the generosity of alternative insurance arrangements, such as any consumption floor or means-tested programs. The targeting benefit of subsidizing formal care is increasing in the extent to which there is state-dependent utility in which marginal utility is greater for people with greater demand for formal care (above and beyond the effects operating through the budget constraint or residual coping costs). If such state-dependence is strong enough, it is optimal to more than fully subsidize formal care (columns 8 and 11).

Although formal care subsidies significantly increase risk sharing, they (optimally) leave some risk uninsured due to the distortion they cause. Both the incompleteness of the insurance and the distortion from the subsidy mean that formal care subsidies fall short of achieving the first best. In the baseline specification, the optimal subsidy achieves about 59 percent of the incremental value over pure-cash benefits of the first-best policy. The shortfall is a measure of the potential gain from using a richer set of policies. A natural enrichment is to condition benefits on verifiable characteristics—i.e., to use tags,—a possibility to which we now turn.

[33] One caveat about this result is that it is based on a model in which formal care is borderline inferior (no income effects). This result need not hold in a more general model with income effects of demand for formal care. It is also important to note that the assumption that formal care is borderline inferior tends to work against the value of in-kind provision by increasing the consumption distortion. The greater are income effects of demand for formal care, the more that the (negative) income effects from subsidizing formal care (due to the consumption distortion) offset the inefficient over-consumption of formal care due to the substitution effect.

6.4 Welfare effects of tags

This section extends the analysis to the case in which different groups of people, defined by their verifiable characteristics, can be offered different benefits. We estimate the gains from catering benefits to different groups of people defined by whether they live alone and the number of activities of daily living limitations they have (2–4, 5, and 6), the two strongest predictors of formal care consumption uncovered in Section 5.[34] The procedure is the same as that in the last section, except that we estimate different α distributions for different groups of people and allow the program to offer different benefits to people in different groups. Figures of the α distributions of each group are reported in Appendix D.

Table 9 shows the results. The planner takes advantage of the flexibility offered by tagging to redistribute across groups by offering high-demand groups more generous benefits. In the case of a tag based on whether someone lives alone, for example, the planner transfers on average over two-and-a-half times more to people who live alone than to people who live with others. Despite this, however, the ex-ante welfare gain from using these tags is quite small. The incremental welfare gain from optimally tagging a pure-cash benefit based on whether someone lives alone is \$227, just 4 percent of the gain from an optimal untagged mixed benefit. The incremental welfare gain from optimally tagging a mixed cash-in-kind benefit based on whether someone lives alone is even smaller—just \$44. The incremental gains from tagging benefits based on the number of activities of daily living limitations someone has are smaller still. The fundamental reason for tags' ineffectiveness in insuring this risk is that much of the heterogeneity in demand for formal care occurs within rather than across groups that can be distinguished on the basis of their verifiable characteristics. This comes through in the correlation between someone's marginal utility in the absence of any program and the optimal tagged pure-cash benefits. This correlation is just 0.20 when whether someone lives alone is used as a tag and just 0.05 when someone's number of activities of daily living limitations is used as a tag.

These results suggest that in the context of homecare, the potential for tags to target transfers toward high-marginal utility types is limited. These results are consistent with those of Mankiw and Weinzierl (2010) on the effects of using height as a tag for optimal income taxation.[35] Although different combinations of observable characteristics could potentially improve on those we have analyzed here, both the small gains from tags based on two of the

[34] We are limited in the number of groups into which we can split the population by the size of the NLTCS sample. We chose the groups to maximize the across-group heterogeneity in the demand for formal care.

[35] In both cases, the optimal tagged transfers are large; the optimal "lives-alone subsidy" is \$4,790 and the optimal "height tax" on someone earning \$50,000 is \$4,500. But the welfare gains from tagging are a small fraction of aggregate income—about 1.5 percent for a "lives-alone subsidy" and about 0.2 percent for a "height tax."

strongest predictors of formal care consumption and the limited extent to which the full set of observable characteristics in the NLTCS predict formal care consumption (as discussed in Section 5) suggest that the scope for tags is limited and reinforce the conclusion that in-kind benefits have an important role to play in terms of targeting benefits to high-marginal utility types.

6.5 Discussion of results

This analysis is subject to several caveats. It assumes that people's decisions about consumption are rational. This ignores possible paternalistic rationales for in-kind transfers, which could be important in the case of home care given the cognitive health problems from which some recipients suffer. Such considerations, which were one of the main motivations for the Cash and Counseling experiments, seem likely to increase the value of in-kind as opposed to cash transfers in this context. It assumes that take-up decisions are fully rational and abstracts from any costs of taking up. This is done to focus on the core tradeoff at the heart of in-kind provision, but it is important to note that many in-kind programs have low take-up rates, whether from low knowledge about or high costs of taking up the programs. It assumes that all ex-post heterogeneity is the outcome of an exogenous process. This feature, which is shared by the vast majority of the large literature on optimal taxation, rules out ex-ante moral hazard (effects of policies on the distribution of ex-post types), which tends to increase the net value of insurance or redistribution. It focuses only on home care and does not explicitly model substitution across other types of care. This was done for simplicity given that there appears to be little substitution across different types of long-term care (Grabowski and Gruber, 2007; Kemper, 1988). Finally, it focuses on singles in order to avoid the many complexities involved in modeling couples, including any financial risk sharing and utility consequences of different caregiving arrangements. An analysis of couples and extended families is an interesting topic for future work.

Acknowledging these caveats, taken as a whole the results suggest that in-kind provision of formal care benefits likely increases welfare despite the large distortion it causes. Although the distortion cost of this particular means of targeting is large, our results suggest that the main alternative means of targeting (using tags) is unlikely to be very effective in this context. These conclusions are robust to a wide range of assumptions. The fundamental reasons for this robustness are the large extent of hard-to-verify heterogeneity in the demand for formal care and the rapid rate at which marginal utility diminishes in the level of consumption under standard utility functions (Kaplow, 2011).

7 Conclusion

We analyze a central tradeoff inherent to in-kind provision—in-kind provision can improve the targeting of benefits at the cost of being less valuable to recipients—in the context of home care. Despite the ubiquity of in-kind transfers, little is known about the magnitude of these key costs and benefits of in-kind provision. We find that in the context of home care, in-kind provision appears to increase welfare despite imposing large distortion costs, since the targeting benefits are even larger. The key factor driving this result is the significant, hard-to-verify heterogeneity in the demand for formal care—whether from hard-to-verify differences in underlying health or in the costs of coping with a given set of health problems—which implies significant heterogeneity in non-care consumption and so, in many models, in marginal utility.

Two main caveats are important to keep in mind in interpreting our results. First, the magnitude of the targeting benefit depends crucially on the utility function, particularly any (state-)dependence of utility on the unverifiable heterogeneity that motivates in-kind provision in the first place. This caveat, common to all questions about the optimal design of insurance programs and other policies that redistribute across people, is extremely important: Improving our understanding of the mapping between spending on market goods and utility is a high priority for future work. Second, we focus only on the potential targeting benefit of providing formal care in kind. We do not consider other potential benefits, including improving tax system efficiency, alleviating the Samaritan's dilemma, and paternalistic benefits, all of which might be important in this context. These other potential benefits seem likely to increase the net value of in-kind provision of formal care, which would reinforce our conclusion about the desirability of in-kind provision of formal care.

Our results have important policy implications. Several recent policy reforms and proposals have made or propose to make benefits that have traditionally been restrictive in-kind benefits more flexible and cash-like. This is true not only in the case of home care but in many other areas as well. The European Commission, for example, recently suggested that policymakers "always ask the question, 'Why not cash?"', U.N. Secretary-General Ban Ki-moon has argued that "cash-based programming should be the preferred and default method of support", and the desirability of a universal basic income is the subject of an active debate in many countries.[36] A major impetus for these proposals is the view that recipients would much prefer cost-equivalent cash transfers, a view that is consistent with our analysis of the particular case of Medicaid home care benefits. Yet a frequently-overlooked consequence of

[36]Switzerland held a referendum on whether to have a universal basic income on June 5, 2016. (It was rejected with a 77 percent majority.) Pilot programs are planned in Finland and Canada. GiveDirectly is providing a universal basic income to dozens of villages and thousands of people in Kenya.

such reform proposals is that, in addition to reducing consumption distortions, such proposals would also tend to systematically change the distribution of benefits received by different people. To the extent that achieving a good targeting of benefits in any particular context is difficult or infeasible without in-kind provision, as our analysis suggests is the case in the context of home care, any gain from reducing distortions must be weighed against any reduction in targeting efficiency that would result.

The issue of optimal benefit design in government programs is a central one, as many of the most important government programs involve in-kind benefits, including public schooling, food stamps, public housing, and Medicare and Medicaid medical benefits. Although home care shares much in common with other types of health care, the desirability of in-kind provision is necessarily context-specific. It is therefore important to evaluate the costs and benefits of alternative benefit designs on a case-by-case basis, and our hope is that the approach we have developed in this paper will prove fruitful in the analysis of other policies as well.

References

Akerlof, G. A. (1978). The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *American Economic Review 68*(1), 8–19.

Ameriks, J., A. Caplin, S. Laufer, and S. V. Nieuwerburgh (2011, 04). The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives. *Journal of Finance 66*(2), 519–561.

Arno, P. S., C. Levine, and M. M. Memmott (1999). The economic value of informal caregiving. *Health Affairs 18*(2), 182–188.

Atalas, V., A. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2016). Self-targeting: Evidence from a field experiment in indonesia. *Journal of Political Economic 124*(2), 371–427.

Atalas, V., A. Banerjee, R. Hanna, B. A. Olken, and J. Tobias (2012). Targeting the poor: Evidence from a field experiment in indonesia. *American Economic Review 102*(4), 1206–1240.

Barczyk, D. and M. Kredler (2016). Evaluating long-term care policy options, taking the family seriously. *Review of Economic Studies*, forthcoming.

Basurto, P. M., P. Dupas, and J. Robinson (2017). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural malawi. *NBER Working Paper No. 23383*.

Benitez-Silva, H., M. Buchinsky, and J. Rust (2004). How large are the classification errors in the social security disability award process? *NBER Working Paper No. 10219*.

Blackorby, C. and D. Donaldson (1988). Cash versus kind, self-selection, and efficient transfers. *American Economic Review 78*(4), 691–700.

Brown, J. R. and A. Finkelstein (2007, November). Why is the market for long-term care insurance so small? *Journal of Public Economics 91*(10), 1967–1991.

Brown, J. R. and A. Finkelstein (2008). The interaction of public and private insurance: Medicaid and the long-term care insurance market. *American Economic Review 98*(3), 1083–1102.

Brown, J. R. and A. Finkelstein (2011). Insuring long-term care in the United States. *Journal of Economic Perspectives 25*(4), 119–42.

Brown, R. S., B. L. Carlson, S. Dale, L. Foster, B. Phillips, and J. Schore (2007). *Cash and Counseling: Improving the Lives of Medicaid Beneficiaries who Need Personal Care Or Home-and Community-based Services: Final Report*. Mathematica Policy Research.

Bruce, N. and M. Waldman (1991). Transfers in kind: Why they can be efficient and nonpaternalistic. *American Economic Review 81*(5), 1345–1351.

Centers for Medicare and Medicaid Services (2010). National health expenditures by type of service and source of funds.

Centers for Medicare and Medicaid Services (2017, June). National health expenditure accounts.

Cohen, J. and P. Dupas (2010). Free distribution or cost-sharing? evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics CXXV*(1), 1–45.

Collinson, R. A., I. Gould Ellen, and J. Ludwig (2015). Low-income housing policy. *Kreisman Working Paper Series in Housing Law and Policy*.

Cunha, J. M., G. De Giorgi, and S. Jayachandran (2011). The price effects of cash versus in-kind transfers. Technical report, National Bureau of Economic Research.

Currie, J. (2006). *Public Policy and the Income Distribution*, Chapter The Take-Up of Social Benefits, pp. 80–148. Russell Sage Foundation.

Currie, J. and F. Gahvari (2008, June). Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature 46*(2), 333–383.

Cutler, D. and J. Gruber (1996). Does public insurance crowd out private insurance? *Quarterly Journal of Economics 111*(2), 391–430.

Da Roit, B. and B. Le Bihan (2010). Similar and yet so different: Cash-for-care in six european countries' long-term care policies. *Milbank Quarterly 88*(3), 286–309.

Dale, S., R. Brown, and B. Phillips (2004, June). Does arkansas' cash and counseling affect service use and public costs? Technical report, Mathematica Policy Research.

De Nardi, M., E. French, and J. B. Jones (2010, 02). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75.

Deshpande, M. and Y. Li (2017). Who is screened out? application costs and the targeting of disability programs. *NBER Working Paper No. 23472*.

Doty, P., K. J. Mahoney, and M. Sciegaj (2010). New state strategies to meet long-term care needs. *Health Affairs 29*(1), 49–56.

Ettner, S. (1995, February). The impact of "parent care" on female labor supply decisions. *Demography 32*(1), 63–80.

Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association 11*(s1), 221–258.

Foster, L., R. Brown, B. Phillips, J. Schore, and B. Lepidus Carlson (2003, March). Improving the quality of medicaid personal assistance through consumer direction. *Health Affairs Web Exclusive*.

Friedberg, L., W. Hao, W. Sun, A. Webb, and Z. Li (2014, October). New evidence on the risk of requiring long-term care. Working Papers, Center for Retirement Research at Boston College wp2014-12, Center for Retirement Research.

Grabowski, D. C. and J. Gruber (2007). Moral hazard in nursing home use. *Journal of Health Economics 26*(3), 560–577.

Hoynes, H. W. and D. Whitmore Schanzenbach (2009). Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Applied Economics 1*(4), 109–139.

James, G., D. Witten, T. Hastie, and R. Tibshirani (2013). *An Introduction to Statistical Learning*. Springer.

Kaiser Commission on Medicaid and the Uninsured (2011). Medicaid home and community-based service programs: Data update. Technical report, Henry J. Kaiser Family Foundation.

Kaiser Commission on Medicaid and the Uninsured (2013, September). Five key facts about the delivery and financing of long-term services and supports.

Kaiser Commission on Medicaid and the Uninsured (2016). Medicaid home and community-based services programs: 2013 data update. Technical report, Henry J. Kaiser Family Foundation.

Kaplow, L. (2011). *The theory of taxation and public economics*. Princeton University Press.

Kemper, P. (1988). The evaluation of the national long term care demonstration: Overview of the findings. *Health Services Research 23*(1), 161.

LeBlanc, A. J., C. Tonner, and C. Harrington (2001). State Medicaid programs offering personal care services. *Health Care Financing Review 22*(4), 155–174.

Lepidus Carlson, B., L. Foster, S. B. Dale, and R. Brown (2007). Effects of cash and counseling on personal care and well-being. *Health Services Research 42*(1), 467–487.

Low, H. and L. Pistaferri (2015, October). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105*(10), 2986–3029.

Mankiw, N. G. and M. Weinzierl (2010, February). The optimal taxation of height: A case study of utilitarian income redistribution. *American Economic Journal: Economic Policy 2*(1), 155–76.

Moffitt, R. (1989). Estimating the value of an in-kind transfer: The case of food stamps. *Econometrica 57*(2), 385–409.

Mullainathan, S. and J. Spiess (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives 31*(2), 87–106.

Munro, A. (1992). Self-selection and optimal in-kind transfers. *The Economic Journal 102*(414), 1184–1196.

National Conference of State Legislatures (2007). The state of long-term care. Technical report.

Nichols, A. L. and R. J. Zeckhauser (1982). Targeting transfers through restrictions on recipients. *American Economic Review 72*(2), 372–377.

Reeder, W. J. (1985). The benefits and costs of the section 8 existing housing program. *Journal of Public Economics 26*(3), 349–377.

Schneider, A., K. Fennel, and P. Keenan (1999, May). Medicaid eligibility for the elderly.

Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear IV regression. Technical report, NBER Working Paper.

U.S. Department of Health and Human Services (1992, April). Estimating eligibility for publicly-financed home care: Not a simple task ...

Viscusi, W. K. and W. N. Evans (1990). Utility functions that depend on health status: estimates and economic implications. *American Economic Review*, 353–374.

Whitmore, D. (2002). What are food stamps worth? *Princeton University Working Paper #468*.

Tables and Figures

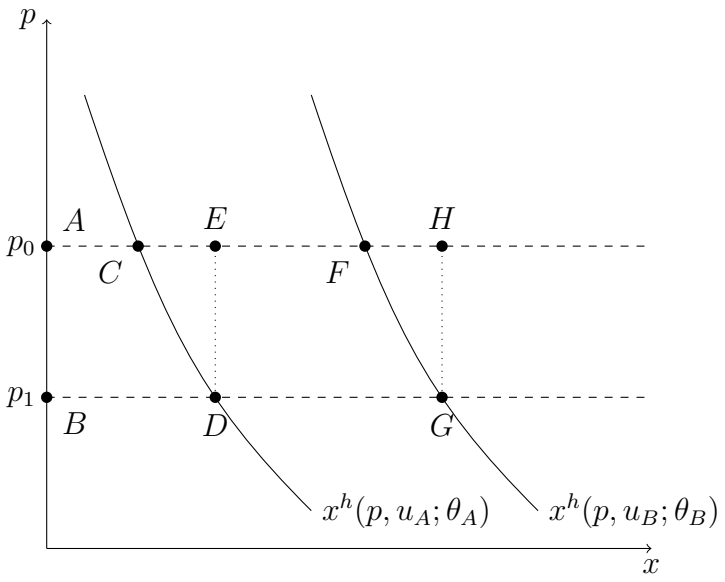


Figure 1: Equivalent variations and excess burdens of a subsidy

[Equivalent variations and excess burdens of a price subsidy that reduces the after-subsidy price from p_0 to p_1 for individuals with different levels of demand for the subsidized good. The equivalent variation of the subsidy is increasing in the level of demand for the good (individual B's equivalent variation, the area bounded by the vertices $ABGF$, exceeds individual A's equivalent variation, the area bounded by the vertices $ABDC$). The excess burdens of the subsidy are independent of the level of demand and instead depend only on the slope. The excess burden of subsidizing individual A's purchases of the good is the area bounded by the vertices CDE, and the excess burden of subsidizing individual B's purchases of the good is the area bounded by the vertices FGH.]

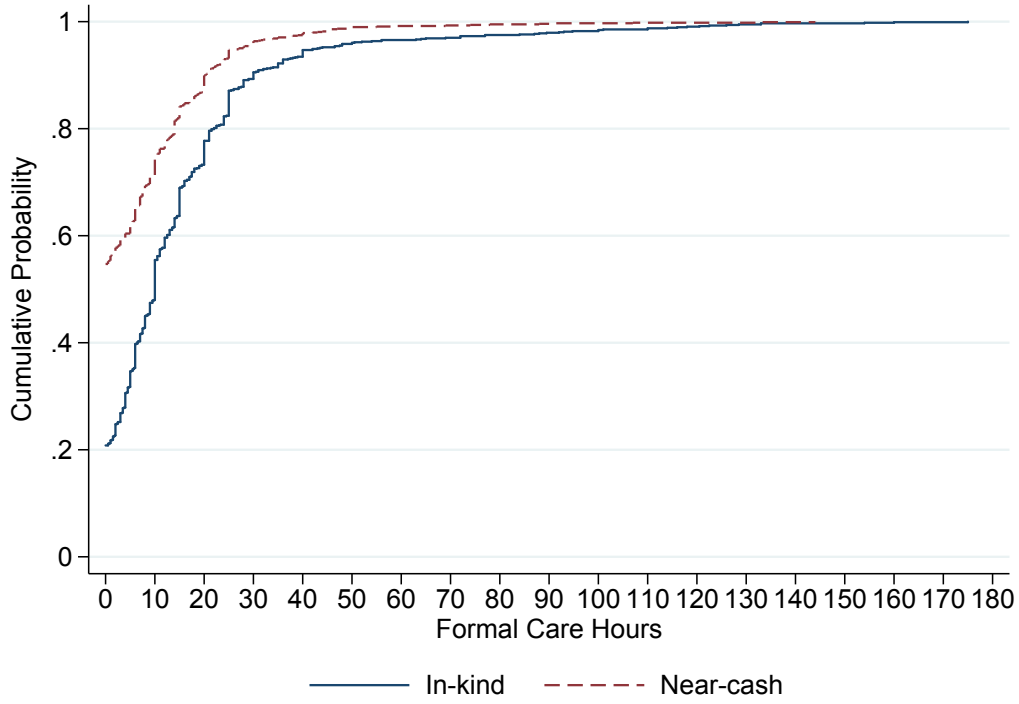
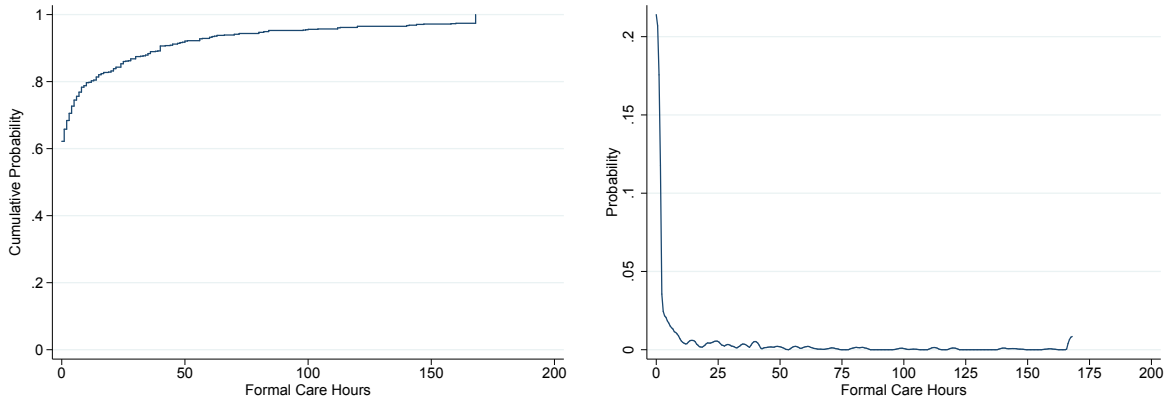


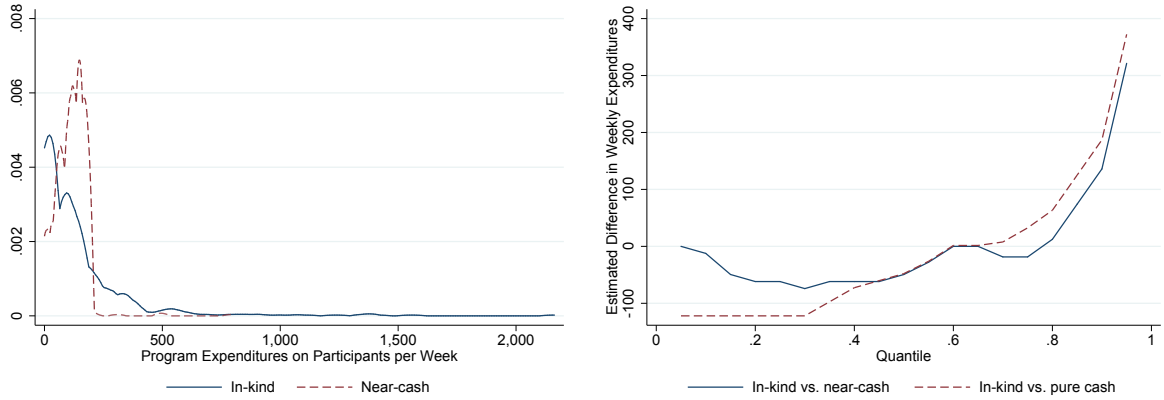
Figure 2: CDFs of Formal Care by Randomized Assignment
[Data from the Cash and Counseling follow-up survey. Hours of formal home care per week.]



(a) CDF (b) PDF

Figure 3: Distribution of Formal Care, NLTCS

[Data from the 1999 National Long-Term Care Survey. Hours of formal home care per week. One individual consumed more than 168 hours of care per week and has been omitted from the figures. Conditional on positive hours of formal care, median consumption is 14 hours per week, the 75th percentile is 40 hours per week, the 90th percentile is 120 hours per week, the 95th percentile is 168 hours per week, and the 99th percentile is 168 hours per week.]



(a) Distributions of Expenditures

(b) Differences of Expenditures

Figure 4: Distributions and Differences of Expenditures in the Arkansas Cash and Counseling Experiment

[Data from the Arkansas Cash and Counseling experiment. Expenditures are the dollar value of benefits transferred to the recipient per week. Kernel density estimates shown for each group on the left. A pure cash program is a hypothetical transfer that provides everyone in the Arkansas Cash and Counseling experiment the average expenditure, approximately $122 per week. The figure to the right shows expenditures for the near-cash or pure cash transfer subtracted from the in-kind transfer.]

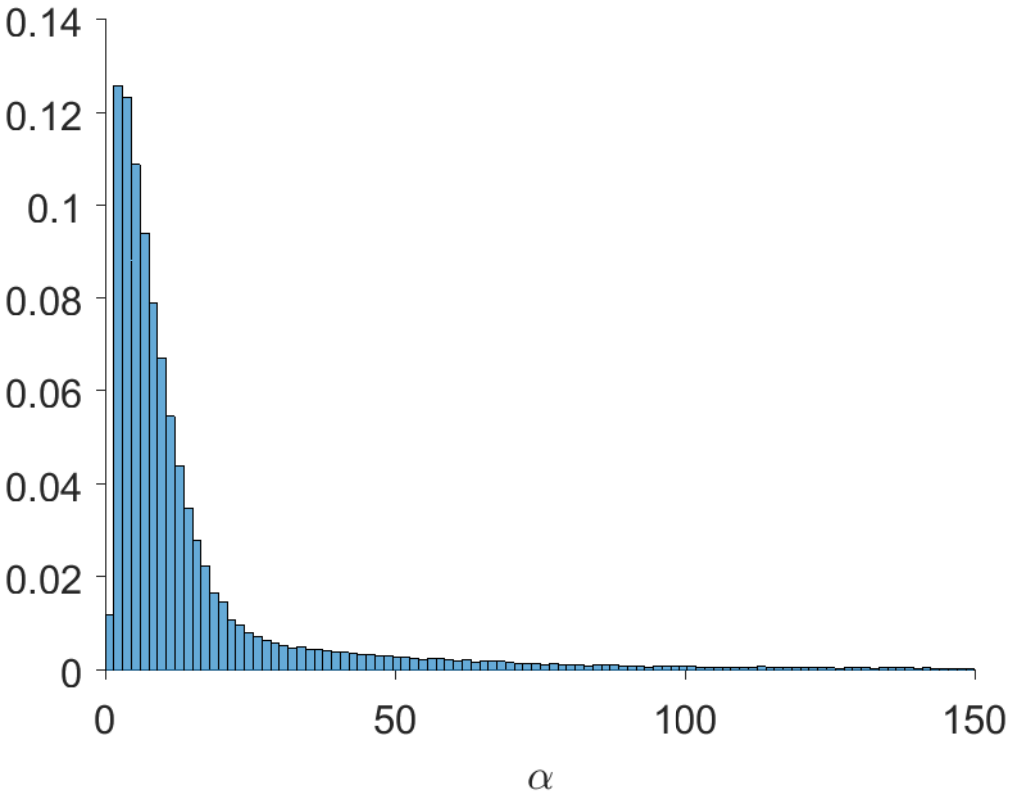


Figure 5: Distribution of the demand for formal care

[Simulated distribution of formal care satiation points, α, in hours per week. The population is people age 65 and older with at least two activities of daily living limitations. The mean is 21 hours per week.]

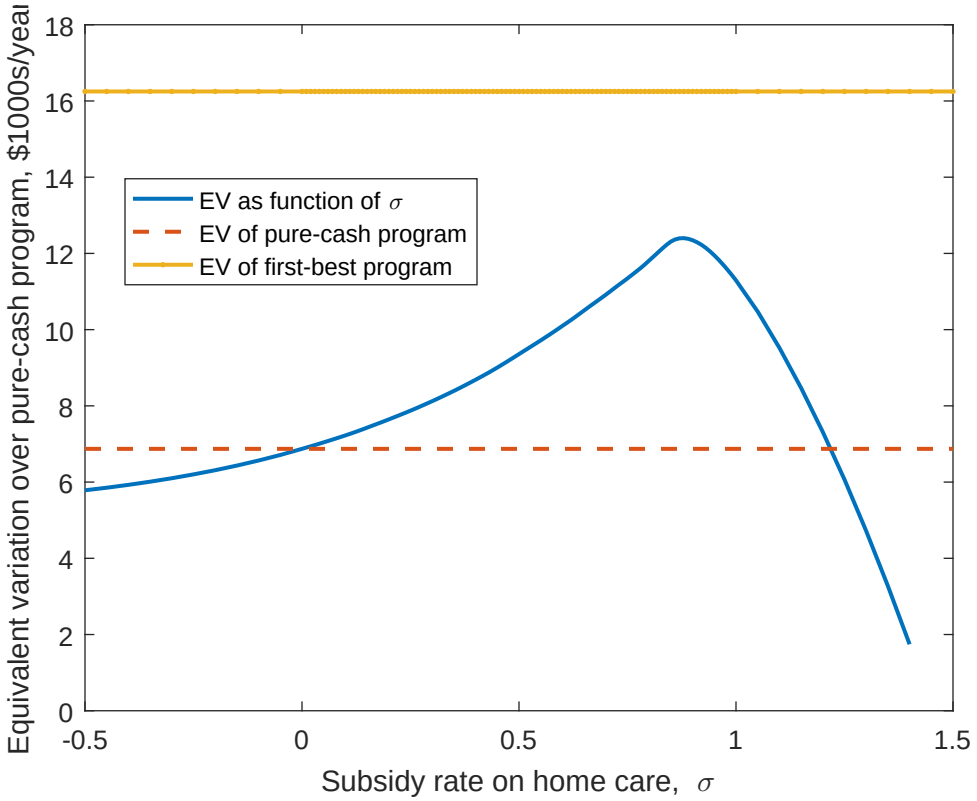


Figure 6: Equivalent variation of mixed cash/in-kind program as function of subsidy rate, σ
[Equivalent variation of mixed cash/in-kind program as function of the subsidy rate, σ. Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total program spending.]

Table 1: Summary Statistics and Balance Tests

	Arkansas			Florida			New Jersey		
	Cash	In-kind	Difference p-value	Cash	In-kind	Difference p-value	Cash	In-kind	Difference p-value
Formal care hours, baseline	8.61	7.85	0.31	13.16	13.26	0.95	16.10	15.39	0.52
Number unpaid caregivers, baseline	2.12	2.13	0.90	1.93	2.05	0.41	2.02	2.15	0.36
Age	78.73	79.33	0.25	79.00	79.87	0.22	77.63	77.81	0.77
Male	0.15	0.17	0.56	0.19	0.22	0.46	0.19	0.22	0.31
White	0.61	0.63	0.60	0.67	0.72	0.21	0.50	0.57	0.11
Less than high school degree	0.67	0.67	0.96	0.37	0.38	0.78	0.66	0.63	0.44
High school degree	0.29	0.26	0.36	0.44	0.46	0.67	0.19	0.21	0.61
College degree or more	0.03	0.05	0.13	0.17	0.14	0.33	0.11	0.13	0.36
Health	3.25	3.28	0.70	3.12	3.22	0.17	3.13	3.24	0.07
Lives alone	0.33	0.33	0.85	0.26	0.32	0.10	0.32	0.37	0.23
Observations	428	432	.	245	237	.	312	292	.

Means by state and type of transfer. P-value is for test of equality of means across the cash and in-kind groups within the state. Formal care hours and number of unpaid caregivers are for the baseline survey at the time of randomization. Remaining variables are measured at the nine-month follow-up.

Table 2: Summary Statistics for NLTCS

	Mean	Standard Deviation
Formal care hours	3.34	17.53
Number of ADLs	0.76	1.50
Health	2.37	0.94
Age	78.92	7.83
Female	0.65	0.48
Unmarried	0.57	0.50
Number in household	1.84	0.93
Any children	0.80	0.40
Income eligible for Medicaid	0.59	0.49

Means and standard deviations presented for variables from the 1999 NLTCS Community survey. Number of observations is 5,147 for most variables. Income eligible for Medicaid indicates whether the person's income would qualify her for Medicaid HCBS in her state.

Table 3: Average Hours of Formal Care by Treatment Group

	Cash	In-kind	Difference p-value
Overall	7.11	14.76	0.00
Arkansas	6.94	11.00	0.00
Florida	7.79	19.35	0.00
New Jersey	6.81	16.60	0.00

Means for formal home care hours per week. Cash indicates group received cash transfer; in-kind indicates group received traditional Medicaid home care. P-value for test of equality across groups shown in last column. Rows denote different samples.

Table 4: Price Sensitivity of Formal Care, First Stage Estimates

	(1)	(2)
Assigned to near-cash	8.14***	8.07***
	(0.25)	(0.25)
Controls	No	Yes
F-Statistic	1,066	1,046
Mean market price	13.73	13.73
Adjusted R-squared	0.35	0.37
Observations	1,946	1,946

Dependent variable is the marginal price of formal care. Data are from the Cash and Counseling experiments. Controls described in text are included in column (2). Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table 5: The Price Sensitivity of Formal Care

	(1)	(2)
Price	-1.85***	-1.82***
	(0.17)	(0.17)
Controls	No	Yes
Mean hours, in-kind	14.76	14.76
Observations	1,946	1,946

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Columns (1) and (2) are IV Tobits where formal care hours are censored at zero. Controls described in text are included in column (2). Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table 6: Summary Statistics by Takeup Decision, NLTCS

	(1) Take-up = 0	(2) Take-up = 1	(3) Difference p-value
Age	81.29	82.34	0.27
Four or more ADLs	0.48	0.61	0.01
Health fair or poor	0.64	0.76	0.02
Female	0.68	0.78	0.04
Unmarried	0.57	0.76	0.00
Has children	0.80	0.75	0.40
Lives alone	0.25	0.45	0.00

Means presented separately for those who had not taken up Medicaid home care (column (1)) and those who had (column (2)) as well as p-value from test of equality of means across groups. Data from the 1999 NLTCS. Only those who had at least two ADLs are included. This leads to a sample of 849 individuals, 104 of whom took up. Alternative to health fair or poor is health good or excellent.

Table 7: Predicting Formal Care Use

	(1) Health	(2) Informal Care	(3) Income	(4) Care Plan
OLS in-sample	0.078	0.103	0.131	0.063
OLS out-of-sample	0.032	0.052	0.059	0.024
Machine learning out-of-sample	0.038	0.063	0.063	0.012

Rows present the R^2 from separate models. Columns (1)-(3) indicate which sets of variables are used to predict outcomes and are cumulative. Thus column (2) uses both health and informal care variables to predict price-adjusted formal care use. Columns (1)-(3) use data from the NLTCS. Column (4) uses data from the Arkansas Cash and Counseling experiment. Only care plan hours are used to predict formal care in column (4).

	(1)	(2)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
			β		Unid'd α's		State-dependent utility						α distribution	
	Baseline	0	25	50	0	Max	1	2	3	4	$\gamma = 1$	$\bar{c}$ =\$2.5k	Drop α >50	$\alpha/2$
Optimal policy														
Subsidy rate, σ^*	0.88	1.00	0.77	-0.50	0.94	0.86	1.30	-0.50	0.69	1.10	-0.50	0.89	0.59	0.75
Equivalent variation over pure-cash policy, \$1,000s														
Optimal subsidy policy	5.53	7.60	0.71	0.00	6.09	4.28	>38.13	0.67	1.86	35.74	0.23	21.07	1.74	2.31
First-best policy	9.38	7.60	1.96	0.22	8.35	9.52	-	-	-	-	-0.08	25.39	2.63	3.34
Non-care consumption, \$1,000s														
Mean, optimal subsidy	15.83	15.00	20.75	21.87	15.14	15.98	12.89	20.38	16.97	14.30	20.38	15.77	17.30	16.34
Mean, pure-cash policy	19.56	16.16	21.87	21.87	17.99	21.14	19.56	19.56	19.56	19.56	19.56	19.56	18.61	17.56
Std. dev., optimal subsidy	1.24	0.00	1.17	0.00	0.65	1.34	3.20	6.03	3.00	1.05	6.03	1.13	1.56	1.22
Std. dev., pure-cash policy	5.61	5.98	0.13	0.00	5.34	5.84	5.61	5.61	5.61	5.61	5.61	5.61	2.56	3.09
Consumption distortion														
Total CV over total cost	0.48	0.65	0.88	-	0.44	0.51	0.02	-	0.63	0.26	-	0.52	0.85	0.73
$E(q_{FC}$|optimal subsidy)	14.01	15.80	2.84	0.00	11.67	17.73	20.31	3.78	11.55	17.30	3.78	14.15	5.37	4.97
$E(q_{FC}$|pure-cash policy)	5.72	13.38	0.02	0.00	5.51	6.26	5.72	5.72	5.72	5.72	5.72	5.72	2.33	2.18
Targeting benefit														
Corr(marg. utility, CV)	0.84	0.90	0.24	-0.68	0.90	0.81	0.11	0.19	0.64	0.85	-0.94	0.81	0.74	0.81
E(1(subsidy $\succ$ cash pol.))	0.16	0.24	0.04	0.00	0.18	0.14	0.14	0.83	0.15	0.15	0.83	0.16	0.18	0.13

Table 8: Policy analysis and robustness. Column 1 presents results based on the baseline assumptions. Columns 2–5 vary the value of β away from the baseline value of 1.8. Columns 6 and 7 vary the values of the α's corresponding to people who consume no formal care when facing a positive price (which are only partially-identified). Column 6 sets these α's to zero. Column 7 sets these α's to the maximum value consistent with these individuals' choices to consume no formal care when facing a positive price. Columns 8–11 use different models of state-dependent utility in which $\mu(\alpha)$ is linear in α and in which the multiplier factors $\mu(\alpha)$ vary by a factor of 100, $\max_\alpha\{\mu(\alpha)\}/\min_\alpha\{\mu(\alpha)\} = 100$. In columns 8 and 9, the utility function is "inner state-dependent." In columns 10 and 11, the utility function is "outer state-dependent." In columns 8 and 10, $\mu(\alpha)$ is decreasing, and in columns 9 and 11, $\mu(\alpha)$ is increasing. See Appendix D for more details about state-dependent utility. Column 12 sets the coefficient of relative risk aversion to one (log utility), whereas the baseline coefficient of relative risk aversion is three. Column 13 sets the consumption floor to \$2,500, whereas the baseline value is \$5,000. Column 14 drops values of α (formal care satiation levels) that exceed 50 hours per week. Column 15 cuts every α value in half. Subsidy rates are constrained to be no smaller than -0.5 (a 50 percent tax) and no greater than 1.5 (a 150 percent subsidy, under which individuals are paid 50 percent of the market price to consume units of formal care). "Total CV over cost" is the total ex-post compensating variation of benefits under the optimal program as a fraction of the total cost of these benefits. Mean values of formal care consumption, $E(q_{FC})$, are in hours per week. "Corr(marg. utility, CV)" is the correlation between marginal utility in the absence of any policy and the ex-post compensating variation of benefits under the optimal subsidy. "E(1(subsidy $\succ$ cash pol.))" is the fraction of people who prefer the optimal subsidy to the pure-cash policy benefit ex post.

	Tag: Lives alone		Tag: Number of ADL limitations		
	No	Yes	2–4	5	6
Average formal care consumption, h/w	9.3	25.1	8.8	19.3	27.3
Optimal policy, $s in $1,000s					
Tagged pure-cash benefits, $(B = b)$	5.67	10.46	6.08	6.4	7.92
Tagged mixed benefits, (B, σ, b)	$(5.67, 0.87, 1.4)$	$(10.46, 0.9, 1.77)$	$(6.02, 0.87, 1.52)$	$(8.53, 0.9, 1.11)$	$(7.92, 0.92, 0.68)$
Equivalent variation over untagged policy, $1s					
Tagged pure-cash benefits, $(B = b)$	227		4		
Tagged mixed benefits, (B, σ, b)	44		8		
Targeting benefit					
Corr(marg. utility, tagged pure-cash benefit)	0.20		0.05		

Table 9: Tags analysis. Average formal care consumption, in hours per week, is estimated in the NLTCS. The sample consists of people age 65 and older with at least two activities of daily living limitations. Subsidy rates are constrained to be no smaller than -0.5 (a 50 percent tax) and no greater than 1.5 (a 150 percent subsidy, under which individuals are paid 50 percent of the market price to consume units of formal care). "Corr(marg. utility, tagged pure-cash benefit)" is the correlation between marginal utility in the absence of any policy and the optimal tagged pure-cash benefits.

Appendices

A Theory Appendix

A.1 The optimal mix of in-kind and cash benefits

Consider a planner choosing how to allocate a given budget, B, between cash and in-kind benefits. The planner's goal is to choose the benefits package that maximizes expected utility:

$$\max_{\sigma} EU(\sigma) = \int_{\Theta} v(p(\sigma), m(\sigma, B); \theta) f(\theta) d\theta.$$

The first-order condition, which holds with equality at an interior optimum, σ^*,[37] is

$$\begin{aligned} \frac{dEU(\sigma^*)}{d\sigma} &= \int_{\Theta} \frac{dv(p(\sigma^*), m(\sigma^*, B); \theta)}{d\sigma} f(\theta) d\theta = E_{\Theta}\left(\lambda(\sigma^*; \theta) \frac{dV(\sigma^*; \theta)}{d\sigma}\right) = 0 \\ &\iff Cov_{\Theta}\left[\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0\right] = (\sigma^* p_X^0) E_{\Theta}\left(\lambda(\sigma^*; \theta)\right) E_{\Theta}\left(\frac{dx_X(\sigma^*; \theta)}{d\sigma}\right). \quad (3) \end{aligned}$$

The second version of Equation 3 shows that, at the margin at an optimum, the covariance between marginal utility and the level of demand for X must be the same sign as the mean marginal change in X due to the shift in benefit composition, i.e., $\text{sign}\left(Cov_{\Theta}\left[\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0\right]\right) = \text{sign}\left(E_{\Theta}\left(\frac{dx_X(\sigma^*; \theta)}{d\sigma}\right)\right)$. This is the classic equity-efficiency tradeoff. Absent distortions, $E_{\Theta}\left(\frac{dx_X(\sigma^*; \theta)}{d\sigma}\right) = 0$, the optimal benefit composition fully eliminates the covariance between marginal utility and the demand for X, $Cov_{\Theta}\left[\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0\right] = 0$. More generally, the greater is the marginal distortion cost of shifting toward in-kind provision, the greater must be the marginal targeting benefit.

The first version of Equation 3 implies that, at the margin at an interior optimum, the benefit to some types from shifting toward greater in-kind provision must be exactly offset by the cost to other types of this shift. Suppose there are just two types, L and H. Then at an interior optimum, the end of the first row of Equation 3 implies that at the margin the planner optimally imposes

$$\frac{\left|\frac{dV(\sigma^*; \theta_L)}{d\sigma}\right|}{\frac{dV(\sigma^*; \theta_H)}{d\sigma}} = \frac{p_H \lambda_H}{(1 - p_H)\lambda_L}$$

dollars' worth of costs on L types in exchange for \$1 worth of benefits to H types. The marginal willingness to pay in terms of costs imposed on Ls in order to help Hs by \$1 is increasing in the ratio of the expected marginal utility of Hs to the expected marginal utility

[37] In certain contexts, including possibly home care, it might be feasible to subsidize formal care at more than a 100 percent rate, so that consumers face a negative net-of-subsidy price of formal care. In this case, the subsidy rate σ can take any real value and the first-order condition holds with equality. A necessary condition for a greater-than-100-percent subsidy to be feasible is that recipients are not able to freely dispose of the good.

of Ls.

A.2 First best

In the first-best case, an individual's type, θ, is verifiable. In this case the planner can choose different (b, σ) benefit bundles for different types. The total derivative of type θ's indirect utility with respect to the in-kind component of its benefit, σ, is

$$\frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma} = \lambda(\sigma; \theta) \left[x_X(\sigma; \theta) p_X^0 - x_X(\sigma; \theta) p_X^0 - (\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right]$$
$$= -\lambda(\sigma; \theta)(\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma},$$

which is negative for all positive subsidy rates. When type is verifiable, a pure cash contract is optimal, and the cash benefits for each type are chosen to equalize each type's marginal utility. Verifiable types means that the planner can redistribute across types without resorting to distortions, so there is no motive for introducing a distortion in this case.

B Robustness and Generalizability of the Estimate of the Demand for Formal Home Care

As we discuss in Section 6, the key conclusion about the desirability of subsidizing formal care is robust to a wide range of values of the price sensitivity of demand for formal care. But the magnitudes of the optimal subsidy and the welfare gains from in-kind provision depend on the particular value of the price sensitivity of demand. The price sensitivity of demand for care is important for other questions as well, including the extent to which private long-term care insurance contracts that subsidize formal care suffer from a "moral hazard tax." In this section, we address issues related to both the internal and external validity of our estimates of the price sensitivity of demand for formal care.

B.1 Internal validity

There are two main threats to the internal validity of our estimate of the price sensitivity of demand for formal care. The first is quantity constraints that might limit consumption of traditional Medicaid home care. If quantity constraints bind, the first stage of our IV overstates the change in prices (marginal values) associated with being randomized to the cash group and thereby leads us to underestimate the price sensitivity of demand. Quantity constraints may have taken two main forms in this context: supply constraints and statutory or de facto limits on Medicaid home care benefits.

Supply constraints are thought to have faced Medicaid home care recipients in Arkansas during the period of the Cash and Counseling experiment (Brown et al., 2007). These

constraints apparently arose from some combination of Medicaid paying below-market prices and the local home care market being in disequilibrium around the time of the experiment. To the extent that such issues were important, ignoring them would tend to lead us to underestimate the true price sensitivity of demand. The simplest way to avoid this issue is to drop Arkansas from the analysis and instead focus on Florida and New Jersey.

Quantity constraints may also have arisen from statutory or de facto limits on how much Medicaid home care people can use. Both Arkansas and New Jersey had statutory limits on Medicaid home care—16 hours per week in Arkansas and 25 hours per week in New Jersey. (Florida had no statutory limit.) Moreover, as discussed in the text, the amount of Medicaid home care that someone can consume is determined by a care plan written by the individual's physician. If physicians, whether in an effort to be "good agents" of Medicaid or for other reasons, prescribe care plans whose hours fall short of their patients' satiation points, then Medicaid home care recipients may not be able to reach their satiation points.

Although in principle the combination of maximum benefit limits and care plan limits could limit the quantity of Medicaid home care available to recipients, in practice it does not appear that either one of these constraints significantly constrained consumption. On care plans, many recipients consume strictly less than their care plan hours, and it is not clear what incentive physicians may have to restrict hours. If anything, physicians' professional norms and ethos might lead them to act as an agent of the patient rather than Medicaid. Maximum benefit limits also appear to be less binding than might have been expected. LeBlanc et al. (2001) survey Medicaid home care programs and discuss several explicit mechanisms for granting exceptions to the limits. For example, in New Jersey, where the statutory limit was 25 hours per week, with prior authorization a recipient could receive between 26 and 40 hours of care per week and with central office approval a recipient could receive as much care as "needed." Consistent with these or other mechanisms relaxing quantity limits, the distributions of formal care hours among Cash and Counseling participants receiving traditional Medicaid home care do not exhibit much bunching around these limits. If the limits were binding, one would expect significant bunching because a binding limit causes a convex kink in the budget constraint between formal care and all other goods.[38] Figures E.1–E.3 present the CDFs of formal care hours for people randomized to the in-kind group in each of the three Cash and Counseling states. In Arkansas (Figure E.1), there is no apparent bunching that would suggest that consumption was constrained by the state's limit. In addition to there not being a large mass point at 16 hours, nearly one-fifth of the sample consumed more care than the state's limit. In New Jersey (Figure E.3), there is bunching at certain points in the CDF of care hours, but this appears to be more of a function of rounding than any limits being imposed. The mass points at 15 and 20 hours (8 and 9 percent of the distribution, respectively) are similarly sized to the mass point at the statutory limit of 25 hours (11 percent).

In Table E.1, we present estimates of the price sensitivity of formal care for each state. The first row shows that the IV Tobit estimates range from -0.96 (Arkansas) to -2.79 (Florida). In the second row, we impose the upper bounds on care hours implied by the Arkansas

[38] Of course, any test of bunching faces the limitation that measurement error lessens observed bunching. A useful feature of our context in this regard is that the tested-for kink in the budget constraint is quite sharp, from zero up to the market price. To the extent that care limits were truly binding, one might expect the limits to be highly salient to recipients and as a result perhaps less attenuation from reporting error.

and New Jersey limits. We censor observations above those cutoffs and use the IV Tobit to re-estimate the price sensitivity. The additional censoring reduces our estimated price sensitivity in Arkansas but increases it in New Jersey. (We exclude Florida since care hours are not limited there.) The differences across states are similar to those found with the standard IV Tobit.

Generally, the results are consistent with the concern that quantity constraints—whether from supply constraints in Arkansas or statutory limits in Arkansas and New Jersey—might be biasing our price sensitivity estimates towards zero. The state without limits (Florida) consistently displays greater price sensitivity than the other states. Because average care consumption is so different across states, it is also useful to consider the percentage changes implied by the coefficients. A one-dollar increase in the price of formal care is estimated to increase formal care consumption by 9 percent in Arkansas, 14 percent in Florida, and 10 percent in New Jersey. The results also reveal important heterogeneity in price sensitivity across states above and beyond that which appears to be due to quantity constraints. We return to this issue in our discussion of external validity below.

The second main threat to the internal validity of our estimate of the price sensitivity of demand for formal care is the distributional assumptions we make in the estimation. The key assumption we make is that the unobservables are jointly normally distributed (particularly that ε_i, the residual in the latent demand function, is normal). This assumption is important because the majority of the cash group and a large minority of the in-kind group do not consume any formal care. People who do not consume any formal care are at a corner, so revealed preference analysis only bounds their level of demand. The Tobit normality assumption is one way among many to deal with this missing data problem.

We test the sensitivity of our results to a number of different distributional assumptions on ε_i. In each case, we continue to instrument for price as we did in the main analysis. These results can be found in Table E.2. As seen in columns (2) through (4), the estimated price sensitivity changes somewhat from one specification to the next but not dramatically so.

In the next four columns of Table E.2, we assume that everyone who is potentially at a corner solution has a marginal value of care of exactly p, the maximum consistent with their behavior. As seen in Figure 2, those in the cash group were more likely to consume zero hours of care than those in the in-kind group. In a Tobit model, this greater mass at the censoring point tends to reduce the (latent) mean of the care hours distribution for the cash group relative to the in-kind group. The 2SLS model does not have this feature and, as a result, tends to produce smaller mean differences between the cash and in-kind groups. In our setting, this translates into a smaller price sensitivity. Again, we instrument for the price of care with each participant's randomly assigned transfer type. Under these assumptions, we tend to find a price sensitivity around -1.

As we show in Section 6, only values of the price sensitivity far greater than any we find in this appendix section can overturn the result that the optimal subsidy on formal care is significantly greater than zero.

B.2 External validity

The generalizability of the results from the Cash and Counseling experiments to other contexts depends on the similarity of the experiments' participants to various populations of interest (in terms of price sensitivity of demand for formal care) and how well the experiments match various policies of interest.

Cash and Counseling participants are unlikely to be representative of Americans 65 and older in bad health. Most participants selected into Medicaid home care, and Medicaid home care recipients have a greater demand for formal care than the population as a whole. The participants are also unlikely to be representative of the broader population of Medicaid home care recipients. Participation in the Cash and Counseling demonstrations is voluntary and the benefits are increasing in the price sensitivity of demand for formal care. By participating, an individual gains the possibility of receiving in cash roughly the cost to Medicaid of providing their formal care benefit. The extent to which an individual values the cash benefit more than the in-kind benefit is increasing in the sensitivity of the individual's demand for formal care to its price. It is natural to expect that participants in the experiments were more sensitive to the price of formal care than the broader population of Medicaid home care recipients in the Cash and Counseling states. This tends to increase our estimate of the price sensitivity of demand for formal care relative to what we would expect to find among the broader population of recipients of Medicaid home care.

Another reason the results of the Cash and Counseling experiments might not generalize well to other contexts is the nature of the experiment itself. Care-giving arrangements, for which people often make important investments such as moving or adjusting their labor supply, likely depend on both the past history of policies and expectations about future policies. People arrange their lives in order to make the best of the choices available to them, and their decisions about where to live and work and whether to use formal or informal home care likely depend on the nature of any home care benefits for which they might be eligible. The Cash and Counseling experiments likely came as a surprise to many participants, and it is unclear what participants might have expected about the persistence of this policy—would it continue indefinitely or would they soon be reverted back to traditional Medicaid home care? Both the surprise aspect and the uncertainty about how long cash benefits might last likely dampened responses relative to what they would have been under an anticipated, permanent policy.

These considerations suggest caution in applying the results of the Cash and Counseling experiments to other contexts. But the robustness of our main conclusions to even large changes in the price sensitivity of demand for formal care greatly limit this concern in our context. And the strengths of the Cash and Counseling experiments—the large, exogenous price variation—make it a valuable piece of evidence about the demand for formal care and the effects of alternative home care-related policies.

C Differential Targeting of In-kind and Near-cash Transfers in the Arkansas Cash and Counseling Experiment

We provide additional evidence on who is targeted by in-kind provision with the Arkansas Cash and Counseling experiment. Note that the experiment's in-kind provision is being compared to the near-cash transfer, not a pure cash transfer. Because the near-cash transfer is tagged, it likely targets resources to the same set of eligibles targeted by the in-kind transfer. As a result, this analysis likely understate the degree to which in-kind provision targets particular demographic groups relative to a pure cash transfer. We run regressions of the form

$$expenditures_i = \beta_0 + \beta_1 inkind_i + \beta_2 X_i + \beta_3 (inkind_i * X_i) + \varepsilon_i \tag{4}$$

where $expenditures_i$ is the dollar cost of benefits received by participant i, $inkind_i$ is an indicator for whether i was randomized to the in-kind group, and X_i is a particular demographic characteristic. The coefficient of interest, β_3, tells us whether those with more of the characteristic X_i receive differentially greater transfers in the in-kind group than do those with lower values of X_i. For example, if X_i indicates having more disabilities, $\beta_3 > 0$ would indicate that those who are more disabled gain more from the in-kind program than do individuals who are less disabled, i.e. that the in-kind program targets those who are more disabled to a greater extent than the near-cash benefit. Table E.3 reports the average effects estimated via OLS as well as impacts in the right tail of the distribution estimated via quantile regression. The right tail of the distribution is of particular importance because that is where there is the greatest scope for targeting to provide insurance value. If in-kind provision targets transfers, then the OLS estimates will reflect an average of the negative effects in one tail with the positive effects in the other tail. The quantile regression, however, will only reflect what is happening in the far right tail and could be more informative about which types of individuals are being targeted. Robust standard errors are reported for the OLS regressions while bootstrapped standard errors are presented for the quantile regressions.

Column (1) of Table E.3 reports that older and sicker individuals received differentially more transfers in the in-kind program. Measures of self-rated health, gender, and a proxy for availability of informal care do not appear to associated with the type of transfer provision. Those who lived alone at the baseline receive differentially fewer resources in the in-kind group than those who did not live alone at the baseline. While living alone could signal having fewer informal care options, it could also signal being in better health since the individual is able to live alone. This latter interpretation appears to be more apt in our context because those who lived alone at baseline had lower costs than those who did not live alone ($107 per week for those who lived alone vs. $129 for those who did not live alone).

Columns (2) through (4) show results at the 90th, 95th, and 99th quantiles respectively. The point estimates in the second row suggest that more disabled participants received differentially more benefits from the in-kind transfer than under the near-cash transfer. This difference appears to be growing as we move further out into the tail of the distribution

where the targeting benefits from in-kind provision could have their largest impacts. We find similar patterns for women and those with less access to informal care, the unmarried.

Columns (5) through (8) present the same analyses for the subset of participants in Arkansas who had not been in the Medicaid home care program at baseline. This group is more representative of the roughly 90 percent of eligibles who do not take up Medicaid home care. Again, we find that the Cash and Counseling in-kind program appears to target more resources towards those with worse health and fewer informal care options. Although we can not observe the marginal utility of income directly, it seems likely that the sicker and those with fewer informal care options are likely to be those with greater marginal utility.[39]

D Welfare Analysis

D.1 Optimal first-best insurance

To better understand the nature of the risk that people face and the desired insurance transfers, consider the benchmark of a first-best insurance program. The first-best transfer schedule satisfies:

$$b(\theta; B) = \begin{cases} b(B) + \frac{\max\{\alpha,0\}^2}{2\beta}, & \text{if } \alpha < \beta p; \\ b(B) + p(\alpha - \beta p) + \frac{\beta p^2}{2}, & \text{if } \alpha \geq \beta p, \end{cases}$$

where B is average per-person spending on people eligible for home care benefits and $b(B)$ is the cash transfer that makes total program spending equal B. The first-best transfer is increasing in α, first quadratically then linearly. With these transfers, indirect utility is

$$v_{FB}(p, m, B; \theta) = u\left(m + b(B)\right),$$

which is independent of θ. The first-best contract does not distort consumption, and it fully insures all risk. By making larger transfers to people with larger demands for formal care, it fully compensates people for their expenditures on formal care and any residual utility costs they face from coping with their health problems.

D.2 Estimating the distribution of demand for formal care

As discussed in the text, we use the observed distribution of formal care consumption together with our estimate of the price sensitivity of demand for formal care to infer the latent distribution of the level of demand for formal care. We express the level of demand for formal care in terms of satiation points, α. The only tricky part of this calculation is that observed formal care consumption does not point-identify α for people consuming zero formal

[39]There is some evidence of health-dependent utility (e.g. Viscusi and Evans, 1990; Finkelstein et al., 2013). If being sicker reduces an individual's marginal utility of income, then it is less clear that a social planner would want to target resources to sick states. We explore the sensitivity of welfare analysis to health-dependent utility in the next section.

care, it only bounds it: $\alpha_i \leq \beta p_i$. We estimate the full α distribution, including the α's of people who consume zero formal care, in three steps.

The first step involves using the observed distribution of formal care consumption, q, to infer the partially-unobserved distribution of latent demand, q^*, where $q_i = \max\{0, q_i^*\}$. In the baseline specification, we fill in the censored values of q_i^* corresponding to the $q_i = 0$ cases by linearly extrapolating the observed q density among people with small positive quantities. In particular, we calculate the number of people in each of two groups: those who consume more than zero and less than five hours of care per week and those who consume more than five and less than ten hours of care per week. Based on the shares of people in each group, we estimate the implied (constant) slope of the probability density function over this range and its level at $q^* = 0$. We assume that this slope remains constant at lower values of q^*, which amounts to assuming that the left part of the underlying latent quantity distribution has a triangular distribution. For each censored q^* (corresponding to an individual who consumed no formal care at market prices), we draw the underlying latent q^* from the truncated triangle distribution based on the estimated slope. Figure E.4 shows the underlying distribution of formal care consumption on which this calculation is based.

Second, we convert each q^* to its corresponding α using the estimated price sensitivity of demand for formal care, $\alpha_i = q_i^*(p) + \beta p$. This adjusts (potentially latent) formal care consumption by our estimate of the impact of the price on consumption. Finally, we estimate the kernel density of the implied α distribution. Figure 5 shows the resulting α distribution. It is mostly just a rightward-shifted version of the observed distribution of formal care consumption, with adjustments for the censoring of people who consume no formal care.

For the tags analysis, we repeat the same procedure for estimating the α distribution separately for different groups of people, as defined by their tagged characteristics. Figures E.5 and E.6 show the α distributions of people who do vs. do not live alone and for people with different numbers of activities of daily living limitations. All of the distributions are similarly-shaped, and they exhibit the expected differences in levels. The demand for formal care is greater among people who live alone than among people who live with others, and it is greater among people with more activities of daily living limitations.

We test the robustness of our results to making different extreme assumptions about how to fill in the unidentified α values. In one case, we set every unidentified α value to zero, which is equivalent to assuming that anyone who consumed no care when facing market prices would also consume no care when facing a price of zero. In the other extreme, we set all of the partially-identified α's equal to their (point-identified) upper bound, $\alpha_i = \hat{\beta} p_i$.

D.3 State-dependent utility

As discussed in the text, any state-dependence in utility that is correlated with formal care consumption is centrally important for the value of in-kind provision, since it affects the value of redistribution across people with different levels of demand for formal care. State-dependence that increases the marginal utility of people with greater demand for formal care relative to people with lower demand for formal care increases the attractiveness of in-kind formal care transfers, whereas state-dependence that decreases the marginal utility of people

with greater demand for formal care relative to people with lower demand for formal care decreases the attractiveness of in-kind formal care transfers. Given the possibility that people with different demands for formal care might have systematically different utility functions, it is therefore important to test the robustness of the results to different possibilities about state-dependent utility.[40]

Two natural ways in which to model state-dependent utility are to introduce a scaling factor on the outside or inside of the utility function:

$$U(c;\theta) = \begin{cases} \mu(\theta)u(c), & \text{"outer state-dependence"}; \\ u(\mu(\theta)c), & \text{"inner state-dependence"}. \end{cases}$$

"Outer state-dependence" multiplies the standard, type-independent component of the utility function by a factor $\mu(\theta) \geq 0$, which is potentially correlated with demand for formal care. This type of state dependence has a straightforward effect on the value of redistribution across types. Types with greater scaling factors have greater marginal utility for any given level of net consumption. "Inner state-dependence" multiplies net consumption (non-care consumption net of any utility costs of residual health problems) inside the standard, type-independent utility function. Unlike "outer state-dependence," "inner state-dependence" can have a subtle effect on the marginal utility of a given level of net consumption. On the one hand, types with greater scaling factors are more effective at converting income into net consumption ("effective consumption" is $\mu(\theta)c$, which is increasing in $\mu(\theta)$ for any c), which tends to increase the marginal utility of income. On the other hand, types with greater scaling factors have greater effective consumption for any given level of net consumption, which tends to reduce the marginal utility of income due to marginal utility diminishing in the level of effective net consumption. With log utility, these two effects cancel out, and "inner state-dependence" has no effect on the marginal utility of income. With preferences in which marginal utility diminishes more rapidly in effective consumption, such as constant relative risk aversion preferences with a coefficient of risk aversion greater than one, the latter effect dominates and types with greater scaling factors have lower marginal utility for any given level of net consumption.

D.4 Robustness

This section provides additional intuition for and discussion of the robustness tests reported in Table 8 and discussed in the main text.

The reason that the results are robust to large changes in the distribution of demand for formal care among people with low demand is that the key driver of the targeting benefit from in-kind provision is the shape of the other tail of the formal care distribution: people

[40] Although health-dependent utility is a natural concern, in the context of home care benefits its importance is somewhat diminished by the fact that most home care benefit programs limit eligibility to people with at least two activities of daily living limitations. This ensures that home care benefits go only to people who have fairly severe chronic health problems. As a result, the type of state-dependence of utility that is relevant for the design of home care benefits (taking as given the eligibility criteria for home care benefits) is state-dependence within the set of (sick) people eligible for benefits, not between people in good vs. bad health.

with high demand for care. The distribution of demand among people with a low demand for care matters mainly for determining the distortion cost of in-kind provision.

The robustness of the results to changes in the right tail of the distribution of demand for formal care partially addresses possible biases from modeling a dynamic situation in a static model. The static nature of the model means that formal care costs must be financed by reducing non-care consumption in that period; they cannot be smoothed over time by saving and borrowing. To the extent that shocks are not entirely persistent, this tends to leads us to overstate the welfare cost of uninsured risk and so the value of insurance against it. This issue is less relevant for Medicaid home care—with its strict asset tests—than for private long-term care insurance. It also addresses possible biases from ignoring other risk-sharing arrangements, e.g., informal family insurance.

That a combination of relatively low risk aversion together with a relatively generous consumption floor can overturn the optimality of a large subsidy on formal care reflects the undesirability of *any* insurance—including a first-best contract—in situations in which means-tested programs are sufficiently attractive. The final column of the table shows that if risk aversion is relatively low ($\gamma = 1$) and the consumption floor is relatively generous ($\bar{c} = \$5,000$), the first-best insurance policy that provides complete insurance without distorting consumption is dominated by an alternative uniform pure-cash benefit that provides no insurance at all. The reason that even a first-best, actuarially-fair insurance contract is dominated by the no-insurance alternative in this case is the high rates of implicit taxation from the consumption floor. Without insurance, the consumption floor pays for much of the care of people with the greatest demand for care. As a result, insurance reduces average consumption among the insured by reducing the transfers they receive from consumption-floor programs. This is similar to Brown and Finkelstein's (2008) findings about how Medicaid can crowd out purchases of even actuarially fair long-term care insurance by a large part of the wealth distribution. It should be noted that while the first-best contract is dominated by no insurance from the perspective of people eligible (or potentially eligible) for home care, the first-best contract is better from the perspective of society as a whole. From the perspective of society as a whole, the home care benefit should internalize any effects alternative home care benefits might have on the rest of society, including government or private consumption-floor programs.

E Appendix Figures and Tables

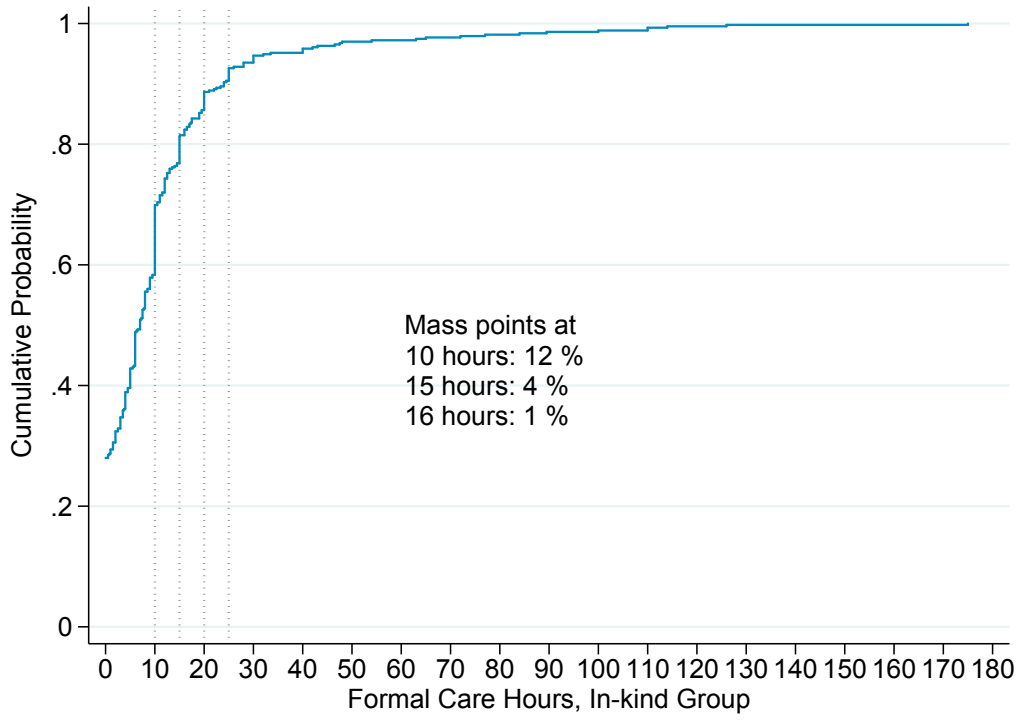


Figure E.1: CDF of Formal Care in Cash and Counseling States, Arkansas

[Data from the Cash and Counseling follow-up survey of the in-kind group in Arkansas. Formal care is measured in hours per week. Arkansas had a regulation that limited care to 16 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

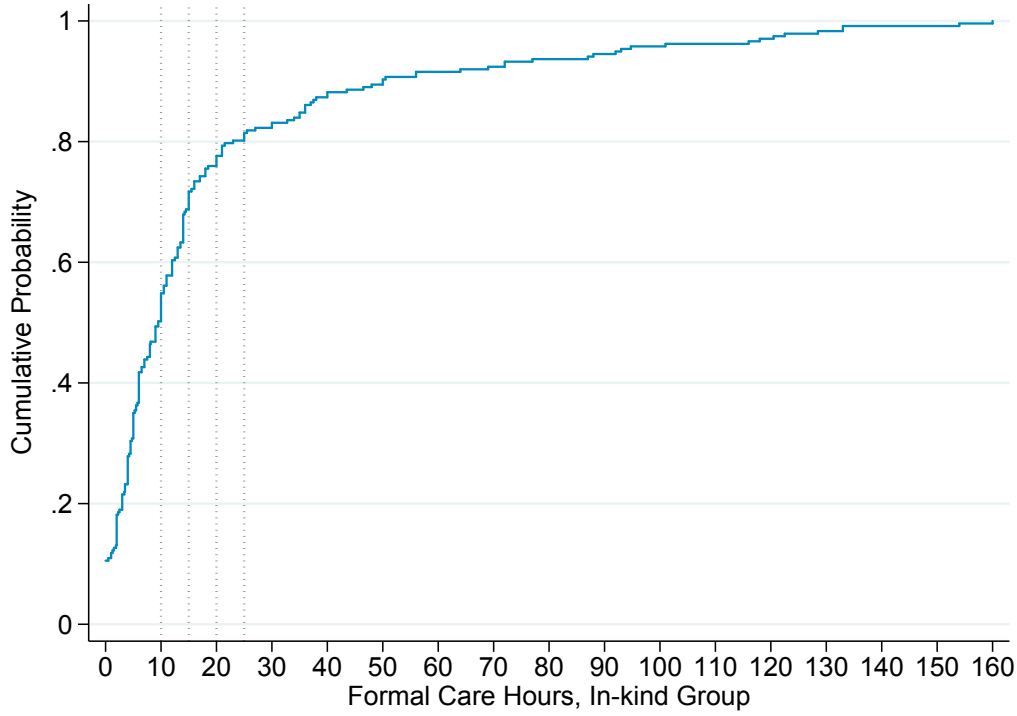


Figure E.2: CDF of Formal Care in Cash and Counseling States, Florida

[Data from the Cash and Counseling follow-up survey of the in-kind group in Florida. Formal care is measured in hours per week. Florida had no regulation limiting care hours (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

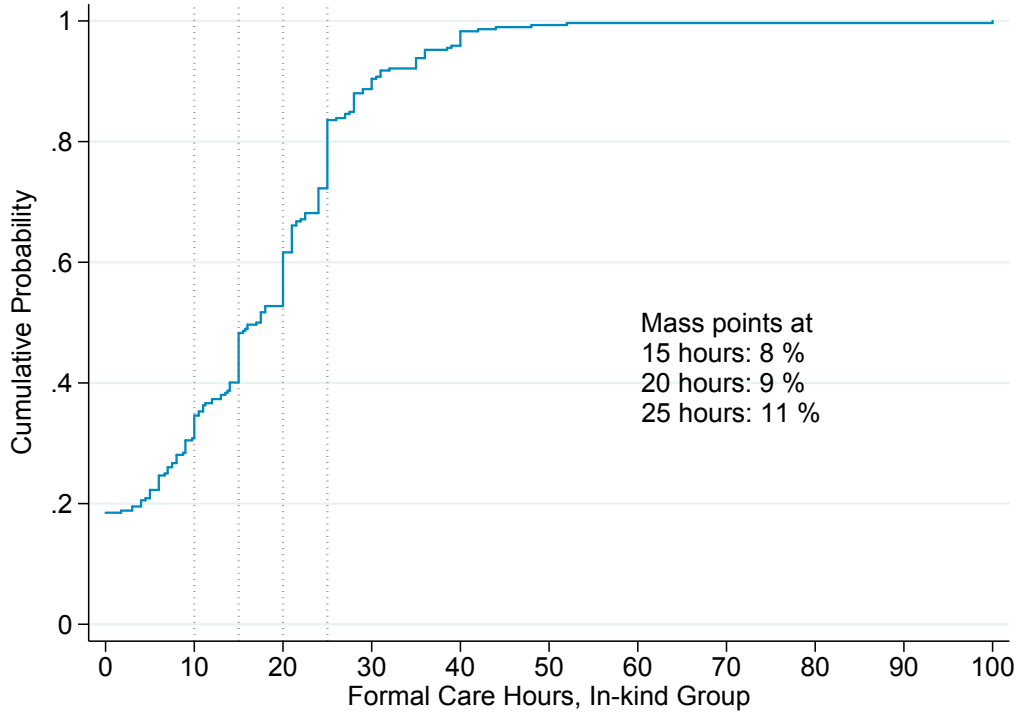


Figure E.3: CDF of Formal Care in Cash and Counseling States, New Jersey

[Data from the Cash and Counseling follow-up survey of the in-kind group in New Jersey. Formal care is measured in hours per week. New Jersey had a regulation that limited care to 25 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

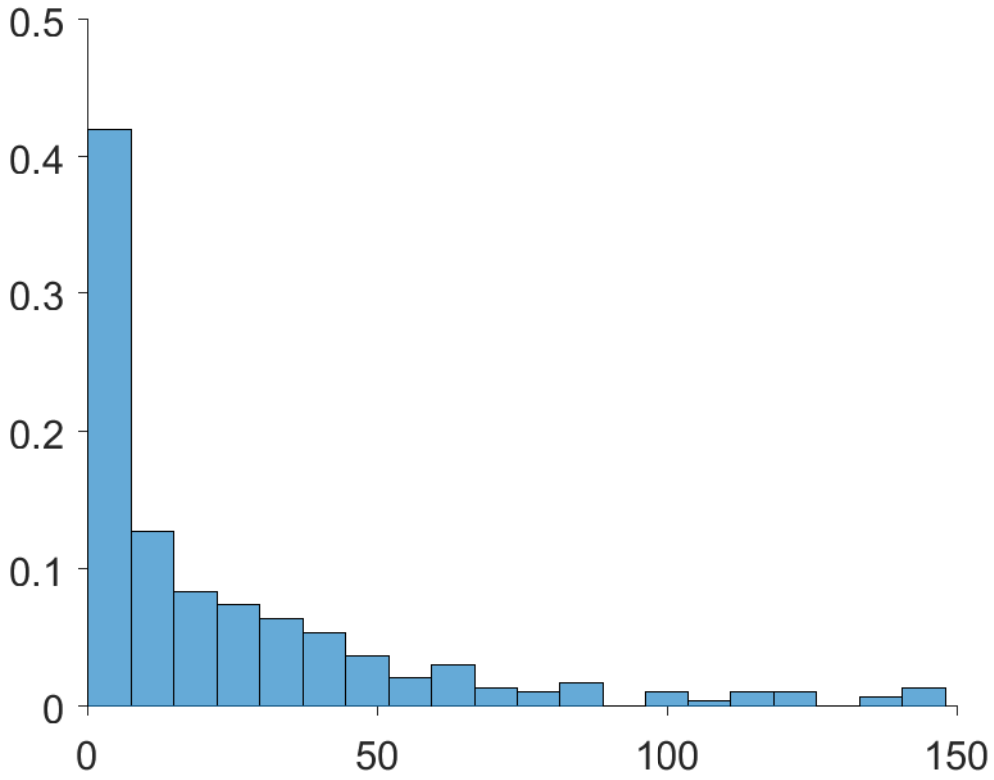


Figure E.4: Distribution of formal care consumption among people with two or more ADL limitations

[Distribution of formal care consumption among people with two or more activity of daily living limitations in the NLTCS. The figure omits the 65 percent of people who report consuming no formal care and the 3 percent of people who report consuming more than 150 hours per week of formal care for readability. The mean of the full distribution is 12 hours per week.]

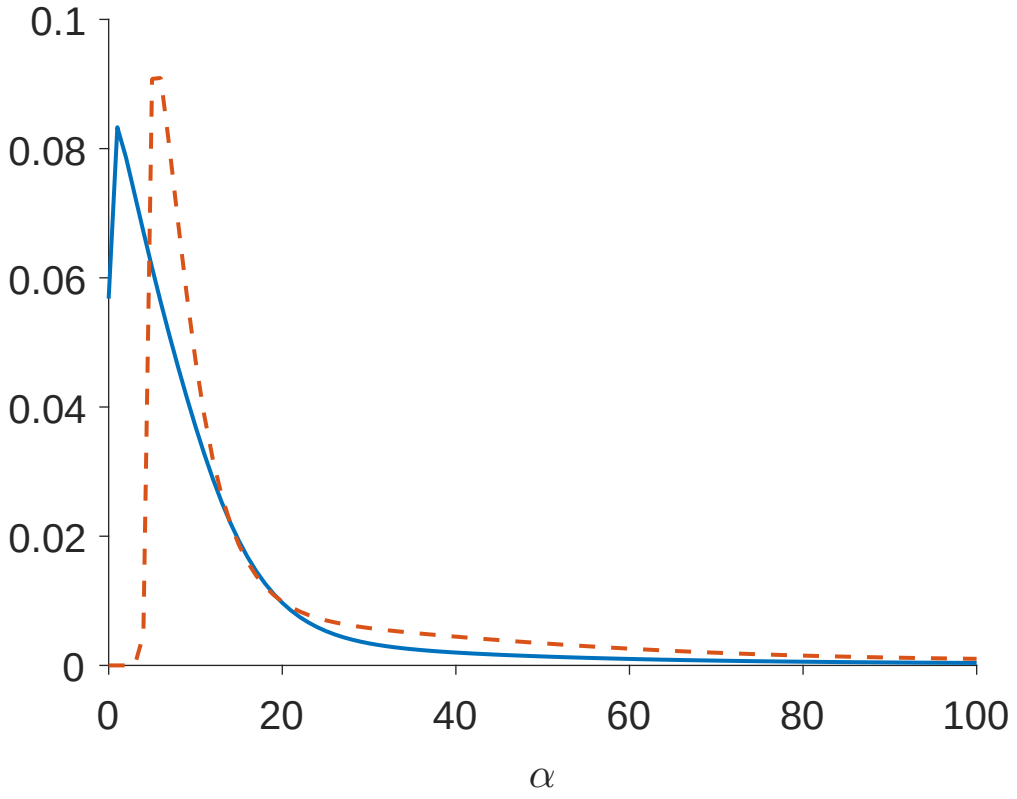


Figure E.5: Distribution of demand for formal care by whether someone lives alone
[Estimated probability density functions of formal care satiation points, α, for each of two groups: people who do not live alone (left-most pdf) and people who do live alone (right-most pdf). The mean of the distribution is 16 hours per week among people who do not live alone and 37 hours per week among people who do live alone.]

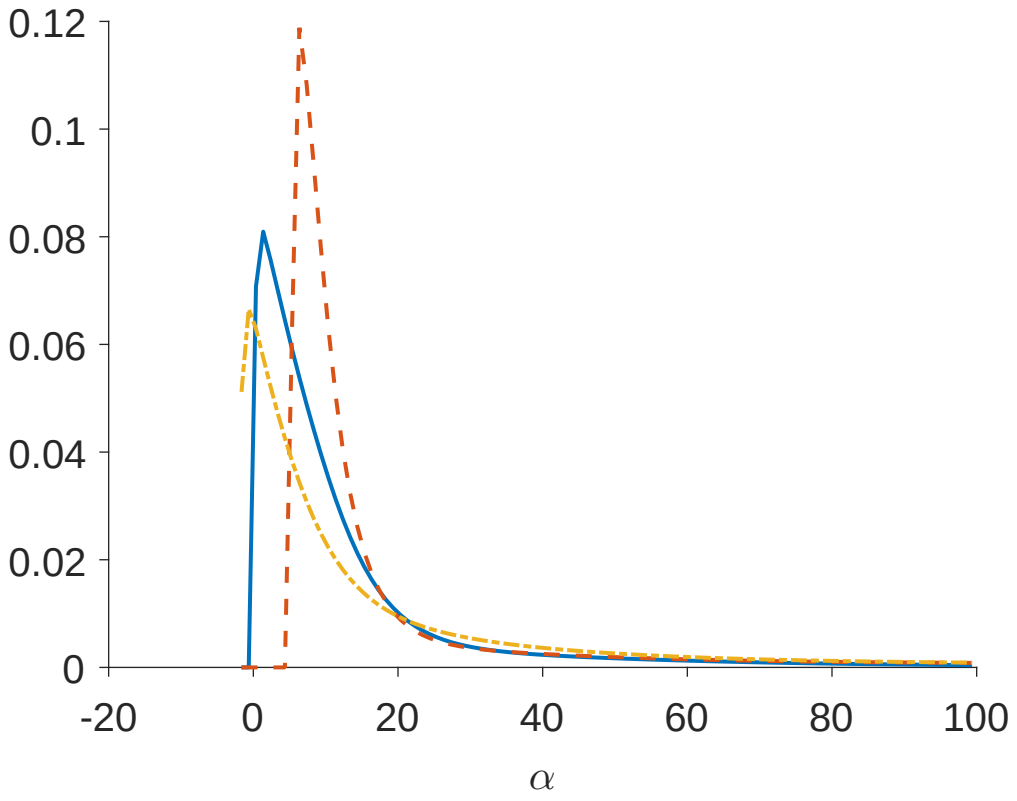


Figure E.6: Distribution of demand for formal care by number of ADL limitations
[Estimated probability density functions of formal care satiation points, α, for each of three groups: people with 2–4 ADL limitations (left-most pdf), people with five ADL limitations (middle pdf), and people with six ADL limitations (right-most pdf). The mean of the distribution is 16 hours per week among people with 2–4 ADL limitations, 31 hours per week among people with 5 ADL limitations, and 34 hours per week among people with six ADL limitations.]

Table E.1: Price Sensitivity and Statutory Limits

	(1) Arkansas	(2) Florida	(3) New Jersey
Price, IV Tobit	-0.96***	-2.79***	-1.71***
	(0.25)	(0.46)	(0.15)
Price, IV Tobit Limits	-0.45***		-1.93***
	(0.12)		(0.16)
Controls	Yes	Yes	Yes
Market price, formal care	12.36	15.09	14.59
Mean hours, in-kind group	11.00	19.35	16.60
Observations	860	482	604

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Seperate regressions run for each state with IV Tobit (first row). Second row uses IV Tobit and imposes statutory limits as upper bounds on care hours. Controls described in text are included in all regressions. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table E.2: The Sensitivity of the Demand for Formal Care to the Composition of Benefits

	Censored errors				Uncensored errors		
	Normal	Extreme Value	Logistic	T-location scale	Normal	Negative binomial	Poisson
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Price	-1.85***	-2.47***	-1.40***	-1.21***	-0.94***	-0.72***	-1.07***
	(0.14)	(0.24)	(0.10)	(0.08)	(0.09)	(0.11)	(0.21)
Mean hours	10.89	10.89	10.89	10.89	10.89	10.89	10.89
Observations	1,946	1,946	1,946	1,946	1,946	1,946	1,946

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Columns (1) - (4) are IV specifications where the error term is treated as censored on the left. Each column presents the estimated sensitivity of demand under a different distributional assumption on the underlying error term. Columns (5)-(7) use distributions that implicitly assume there is no censoring on the left. All models instrument for price with the participant's randomized treatment status and are estimated via two-stage residual inclusion. for columns (6) and (7), average marginal effects are reported. * p<0.10, ** p<0.05, *** p<0.01

Table E.3: Targeting in the Cash and Counseling Experiments, Arkansas

	Entire Sample				Not Enrolled at Baseline			
	(1) OLS	(2) 90th Quantile	(3) 95th Quantile	(4) 99th Quantile	(5) OLS	(6) 90th Quantile	(7) 95th Quantile	(8) 99th Quantile
Age $\geq$ 80	36.0*	30.9	111.2	584.0	56.6	241.0**	305.9	1,579.0***
	(18.9)	(39.8)	(121.3)	(382.2)	(38.6)	(117.3)	(324.9)	(485.9)
ADLs	20.8***	70.0***	105.1***	168.9	6.4	24.7	86.5	559.3**
	(8.0)	(13.8)	(31.1)	(152.4)	(13.2)	(51.3)	(104.5)	(277.3)
Health fair or poor	-6.6	-24.7	-123.6	435.7	-31.1	-98.9	-123.6	312.1
	(27.6)	(66.5)	(433.5)	(496.9)	(44.6)	(106.1)	(419.3)	(789.4)
Female	12.9	37.1	185.4**	543.8*	36.5	105.1	281.2**	1,053.7**
	(18.3)	(41.6)	(72.6)	(321.9)	(27.3)	(86.5)	(125.8)	(488.5)
Unmarried	-6.7	-0.0	123.6	716.9***	21.1	49.4	309.0**	1,016.6*
	(18.7)	(60.1)	(102.2)	(203.8)	(27.9)	(76.3)	(129.4)	(522.4)
Lived alone at baseline	-37.6**	-98.9*	-185.4**	-704.5*	-46.5	37.1	-222.5	-855.9
	(17.4)	(57.4)	(90.2)	(382.1)	(28.9)	(101.6)	(207.9)	(571.2)

Dependent variable is dollar costs of benefits for participants in the Arkansas Cash and Counseling experiment. Each row presents results from a separate regression. The omitted health category is health good or excellent. Columns (1) and (5) present results for OLS regressions; remaining columns for the specified quantiles. Robust standard errors presented for OLS results; bootstrapped standard errors presented for quantile regressions. Columns (1) through (4) include all Arkansas participants. Columns (5) through (8) only include Arkansas participants who had not been enrolled in Medicaid home care before the baseline. * p<0.10, ** p<0.05, *** p<0.01