Online Appendix

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

Ethan M.J. Lieber and Lee M. Lockwood

A Approach Appendix

A.1 In-kind transfers with maximum benefit limits and incomplete take up

In the rest of the paper, in order to match our context of Medicaid home care, we consider an in-kind transfer without a binding maximum benefit limit. This is a good approximation to several important contexts, including many in-kind health care benefits. But some in-kind transfer programs have binding maximum benefit limits. We analyze this case here, while also allowing for the possibility of incomplete take up.

As discussed in Section 2, for any in-kind transfer there is a subsidy that has the same effect on recipients' budget constraints. The simplest type of in-kind transfer with a binding maximum benefit limit allows recipients to consume up to μ units of the good free of charge and does not subsidize consumption beyond that limit. Provided that resale is not possible, this has the same effect on recipients' choice sets as a piecewise-linear subsidy schedule with a 100 percent marginal subsidy rate on the first μ units of consumption and a 0 percent marginal subsidy rate on any additional units of consumption.[24]

Consider a benefit program that combines a cash benefit, b, with a 100 percent subsidy on the first μ units of consumption of good k and no subsidy on additional consumption beyond μ. The individual automatically receives the cash benefit, regardless of the state

[24]The nature of resale opportunities, if any, is an important determinant of the effects of in-kind benefit programs. The better are resale opportunities, the more cash-like is an in-kind benefit. In the case of home care benefits, resale is impossible. In the case of food stamps, by contrast, resale does occur, albeit at a discount from face value (Whitmore, 2002). Another important consideration is whether recipients can "top up" their consumption of the good beyond the in-kind benefit by spending their own resources. Schooling vouchers, for example, can generally be topped up, whereas public schooling cannot. Here we consider a situation in which resale is impossible and individuals can top up their consumption of the good provided in kind by purchasing it in the market.

of the world, but may or may not take up the in-kind benefit. Any in-kind benefit the individual receives cannot be resold.

Consider a cost-neutral shift toward in-kind provision. This increases the maximum benefit limit, μ, while decreasing the cash benefit to maintain the same expected spending. The marginal benefit of the increase in μ is

$$MB = \frac{-\partial E\left(v(p,m,\mu)\right)/\partial\mu}{\partial E\left(v(p,m,\mu)\right)/\partial m} = \frac{E\left(\lambda V\right)}{E\left(\lambda\right)} = E\left(V\right) + Cov\left(\widehat{\lambda}, V\right),$$

where V is the ex post marginal value of the increase in μ in a particular state:

$$V = \begin{cases} 0 & \text{if individual does not take up} \\ 0 & \text{if takes up and reaches satiation, } x_k < \mu \\ p_k^0 & \text{if takes up and is inframarginal to the in-kind transfer, } x_k > \mu \\ MRS_{k,A} \in [0, p_k^0] & \text{if takes up and is marginal to the in-kind transfer, } x_k = \mu, \end{cases}$$

where $MRS_{k,A}$ is the marginal rate of substitution between good k and "all other goods," i.e., the marginal value of good k in units of income.[25] This is a slightly adjusted version of the marginal benefit equation in the text. Whereas the ex post marginal benefit of a reduction in the price of good k in any state equals consumption of good k in that state, the ex post marginal benefit of an increase in the maximum benefit limit is $V \in [0, p_k^0]$, which is 0 in states in which the individual does not take up benefits and is increasing in the level of demand for good k in states in which the individual does take up benefits.

The marginal cost to the insurer of the increase in μ is

$$MC = \frac{d}{d\mu}\left\{E\left(TU \times \min\{\mu,\ x_k\}\right)\right\},$$

where $TU \in \{0, 1\}$ is an indicator of whether the individual takes up benefits. This marginal cost includes both the marginal increase in costs among those who took up before the change and the full costs of those induced to take up by the change. As Finkelstein and Notowidigdo (2018) emphasize, when take-up decisions are privately optimal, to first order changes in take up have only costs and no benefit.

Although this analysis considers a counterfactual different from that in the main text, the same core tradeoff of in-kind provision arises and the same considerations apply. The

[25] What matters for the ex post marginal value of the increase in μ is whether the individual was taking up benefits *before* the change, since the marginal benefit of changes in the program to those induced to take up by those changes is zero.

targeting benefit of an increase in the benefit limit is increasing in the covariance between the value of the increase and marginal utility. The value of the increase in the benefit limit is closely related, though not identical, to the level of demand for the good. The distortion cost is increasing in the extent to which in-kind provision leads people to consume more of the good than they would when facing the market price.

A.2 Applicability of the approach

As discussed in the main text, we focus on the case of in-kind benefits insuring a risk. But with minor adjustments, the framework can be used to study other settings in which there is a tradeoff between targeting and the ex post value of the transfer. These include the many policies in which the size of the transfer an individual receives depends on her consumption of a particular good or bad.

One can view an in-kind benefit program as providing a cash benefit while at the same time imposing the 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. Hence, an in-kind transfer is just one of many possible restrictions that are imposed on recipients.

The same core tradeoff applies to other restrictions as well. Tightening a restriction on recipients weakly reduces the value of the benefit to each *potential* recipient. This has two key effects. First, it reduces the value of the benefit to actual recipients, which is costly. Second, it may reduce the extent to which some potential recipients take up benefits relative to others, which could be beneficial or costly. The value of this targeting effect depends on the extent to which the cost of meeting the restriction, which depends on the demand for the underlying good or bad, covaries with marginal utility. Since no good is a perfect "indicator" of marginal utility, the covariance between marginal utility and demand reflects two types of errors: benefits are too large in some states of the world and too small in others (see, for example, Kleven and Kopczuk, 2011). Of course, a given restriction is worth imposing only if the benefits of doing so exceed the costs and cannot be achieved at a lower cost in some other way. A minimum requirement is that there is substantial, costly-to-verify heterogeneity within the eligible population.

Although the targeting versus value tradeoff seems likely to be central for many in-kind transfers, it does not appear to be so for the U.S. Supplemental Nutrition Assistance Program (SNAP) ("food stamps"). Most recipients spend more on food than their potential food stamp benefit. For them, the benefit is inframarginal and the restriction imposed by in-kind provision is not binding. This suggests that any targeting and distortion effects of making this transfer in kind are likely to be modest (though not necessarily zero; several studies find effects of food stamps on consumption even among people who spend more than the benefit amount on food, e.g., Hastings and Shapiro, 2017). But the highly incomplete take up of food stamps among the eligible population indicates that other restrictions imposed by the program likely have important targeting effects. Food stamps depart in important ways from the theoretical ideal of a pure-cash benefit, which everyone can take up at no cost. Potential recipients must actively apply for food stamps, so awareness about the program, hassle costs of taking up, and stigma costs of receiving the benefit might all have important effects on take up and targeting (see Currie (2006) for a review and Finkelstein and Notowidigdo (2018) for recent experimental evidence).

B Medicaid Home Care and the Cash and Counseling Experiments

B.1 Medicaid home care

Medicaid plays a major role in financing home care. Medicaid home care programs have grown rapidly in recent years, from 1.9 million recipients in 1999 to nearly 3 million recipients in 2013, and from 18 percent of Medicaid's long-term care spending in 1995 to 51 percent in 2014 (Ng et al., 2016). Summaries of Medicaid-provided home care services are available in LeBlanc et al. (2001) and Ng et al. (2011).

Eligibility for Medicaid home care is determined by financial- and health-related criteria. An individual must have sufficiently low income and assets and must have at least two ADL limitations that are expected to last at least 90 days. Medicaid is financed jointly by the federal and state governments, and Medicaid policies vary somewhat across states. 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. 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 (SSI) amount (LeBlanc et al., 2001). States with more restrictive income limits use 100 percent of the SSI amount.

In principle, the amount of Medicaid home care for which an individual qualifies is determined by a medical exam with an approved medical care provider. During this exam, the individual and provider compile a list of activities with which the individual needs assistance as well as how much time that assistance will take. That information is compiled in the individual's care plan. The applicant's health care provider then submits the care plan to the relevant state agency for approval. Once approval is given, the individual and an approved formal home care agency develop a schedule for the needed care. The individual and her care plan must be evaluated at regular intervals, often every six months. In many states, the amount of care people can receive is also limited by maximum benefit rules. In practice, however, it appears that in the Cash and Counseling experiments, neither care plans nor maximum benefit rules limited the transfers of recipients of the in-kind benefit (see Appendix Section B.4).

B.2 Cash and Counseling Experiments

The Cash and Counseling experiments were large-scale experiments conducted by the Medicaid programs of Arkansas, Florida, and New Jersey in the late 1990s and early 2000s (for more details see Brown et al., 2007). Participants were enrolled beginning in 1998 in Arkansas, 1999 in New Jersey, and 2000 in Florida. In New Jersey and Florida, only individuals who were currently receiving Medicaid home care were eligible to participate in the experiments. Arkansas allowed a limited number of individuals who qualified for but were not receiving Medicaid home care to participate.[26] Both non-elderly and elderly individuals were enrolled and there was no screening on whether the individual had or would be able to find sources of care. Participants were given a baseline survey and then randomized to the traditional in-kind benefit or an experimental near-cash benefit, each with a 50 percent probability. Participants were surveyed 4–6 months after enrollment and again 9 months after enrollment. We use data from the baseline and 9-month follow-up surveys.

Each individual's near-cash transfer was slightly less than the cashed-out cost of the individual's care plan. This stemmed from a requirement that the experimental cash treatment be budget-neutral, which meant that the costs of paying the counselors who helped treatment group members manage their care came out of the cash allowances. In New Jersey, for example, 10 percent of the value of the care plan was set aside to cover program costs. Counselors were available to help participants develop plans for spending their benefit, secure caregiver services, issue checks to caregivers and other service providers, handle paperwork associated with being an employer (e.g. payroll taxes), and maintain the necessary records. Recipients had to submit receipts documenting that they spent at least 90 percent of their

[26]These individuals had to verbally commit to seeking the in-kind benefit if they were randomly assigned to it.

benefits on personal care services. The idea was that the remaining 10 percent could be spent on services that could not be readily invoiced, like payments to a neighbor for mowing the lawn.

Appendix Table E.1 provides summary statistics on the Cash and Counseling participants and balance tests of the randomization. We restrict the sample to people who are at least 65 years of age and who have non-missing data on age, sex, race, education, and self-rated health. Our final sample includes 2,470 individuals, of whom 30 are missing data on formal care consumption at follow-up. This leaves us with 2,440 individuals for analyses that require this variable. At baseline, average formal care consumption ranges from 9 (Arkansas) to 16 (New Jersey) hours per week, and the average number of informal caregivers is two. The average age is in the upper 70s, the majority of participants are female, and education levels are low. Although non-negligible fractions of the treatment and control groups attrited from the experiment before the nine-month follow-up survey (20 and 35 percent, respectively), of the 30 balance tests, none of the differences between treatment and control groups are statistically significant at the 5 percent level and only one is significant at the 10 percent level.

Not surprisingly, participants in the experiments are somewhat different from the broader population of Medicaid home care users in the US. Appendix Table E.2 shows that compared to Medicaid home care users in the US, participants in the experiments are similar in terms of age (around 79 on average) and health status (about three-quarters self report fair or poor health), but they have lower formal care consumption (12 vs. 36 hours per week) and are less likely to be living alone (32 vs. 39 percent). The differences could arise from selection into the experiment, differences in the generosity of states' Medicaid home care programs, or from differences in the composition of Medicaid home care users across states. Unfortunately, the NLTCS has too few Medicaid home care recipients in Arkansas, Florida, and New Jersey to address this directly. We discuss issues related to the internal and external validity of our analysis in more detail in Appendix Section C.

B.3 Estimating Take-up of Medicaid Home Care

Take-up rates are notoriously difficult to estimate both for means-tested programs in general and for Medicaid in particular (U.S. Department of Health and Human Services, 1992; Currie, 2006; Sommers et al., 2012). Eligibility rules often are complex, vary from state-to-state, and depend on household characteristics that are unobservable to the researcher. We estimate take-up rates of Medicaid home care by combining data from the NLTCS, the size of the 65-and-older population, and administrative estimates of the number of Medicaid home care users from LeBlanc et al. (2001). 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). To be eligible, someone must have at least two ADL limitations and meet income and asset requirements. The main source of uncertainty in our estimated take-up rate is the incompleteness of the information on household assets in the NLTCS. Given this data limitation, we aim to bound the true eligibility rate. Our less restrictive eligibility threshold uses the income limits from Schneider et al. (1999) and limits eligibility to households with fewer than two cars. Our more restrictive eligibility threshold uses (much) more restrictive income and asset requirements than the actual limits in the vast majority of states: Household income must be no more than 100 percent of the SSI benefit and the household must have no cars (car value is one of the primary inputs to the asset tests). The more restrictive the eligibility definition, the greater the implied take-up rate among eligibles. Given that our most restrictive eligibility estimate likely understates eligibility substantially, the implied take-up rate of 19 percent likely exceeds the true take-up rate.

B.4 Benefit Limits in the Cash and Counseling Experiments

In this subsection, we provide suggestive evidence that benefit limits do not appear to have been binding for recipients of the in-kind benefit in the Cash and Counseling experiments.

In the Arkansas Cash and Counseling data, approximately 30 percent of Medicaid home care recipients consume more formal care than the number of hours in their care plans. This is true both for the care plan created at the baseline as well as the care plan in effect 12 months after baseline. Because we measure consumption nine months into the experiment, it is possible that some of these individuals had a different care plan in operation when their consumption was measured. However, the strong correlation between care plan hours at baseline and 12 months later, 0.86, makes it unlikely that this can explain much of the excess of consumption over care plan hours. And if care plans were binding, it is not clear what incentive physicians might have to restrict care plan hours below what the recipient, their patient, would like. Physicians' professional norms and ethos emphasize acting as an agent of the patient, not Medicaid or other parties.

State Medicaid programs' maximum benefit limits do not appear to have been binding either. LeBlanc et al. (2001) survey Medicaid home care programs and discuss several explicit mechanisms for granting exceptions to the limits. For example, recipients in New Jersey, where the statutory limit was 25 hours per week, could with prior authorization receive up to 40 hours of care per week and with central office approval could receive as much care as "needed."

Appendix Figures E.1–E.3 present the distribution of formal care consumption among people randomized to the in-kind benefit in each of the three Cash and Counseling states. The distribution of formal care consumption in Arkansas shows no apparent signs of having been influenced by the statutory limit of 16 hours per week. Nearly one-fifth of the sample consumed more than the limit and there is no apparent bunching at that quantity: Only 1 percent of recipients consume 16 hours per week, whereas 10 percent consume 10 hours per week and 4 percent consume 15 hours per week. The distribution of formal care consumption in New Jersey exhibits some bunching at the statutory limit: Ten percent of people consume the statutory limit of 25 hours per week. But this bunching is only slightly greater than that at other round-number amounts. For example, 7 percent of people consume 15 hours per week and 9 percent consume 20 hours per week. In addition, about one-sixth of people consume more than the statutory limit. 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, increasing the price from zero to the market price. If benefit limits were binding, one would expect them to be highly salient, which might reduce attenuation from reporting error.

C Moral Hazard Effects of In-Kind Provision: Robustness and Generalizability

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. This price sensitivity is important for other questions as well, including the extent to which insurance contracts that subsidize formal care suffer from a "moral hazard tax." In this section, we address issues related to the interpretation of our estimates and their internal and external validity.

C.1 Interpretation

Throughout the paper, our analysis is based on the standard revealed-preference assumption that observed choices correspond to people's preferred allocations in their opportunity sets. The usual concern that actual choices might diverge from utility-maximizing ones may be especially relevant in this context. People with chronic health problems may have more trouble than most in consuming their most-preferred bundles. This seems likely to reduce formal care consumption since they may have difficulty finding and coordinating care, particularly for those who consume large amounts of care. This would affect both of the key empirical ingredients in the analysis.

First, to the extent that optimization frictions reduce formal care consumption, the observed distribution of formal care consumption understates the extent of the risk. This would tend to reduce the targeting benefit of in-kind provision and so work against our conclusion that the targeting benefit is large.

Second, it could affect our estimate of the price sensitivity of demand. The key issue is whether such frictions are larger for people receiving the traditional in-kind benefit or the experimental near-cash benefit. Under the traditional in-kind benefit, Medicaid bears many of the costs of finding and coordinating with formal home care providers. This presumably helps recipients get the care they want. Under the near-cash benefit, recipients have more control of and so responsibility for coordinating their own care. They receive help from their counselors, but they may still face higher costs of contracting with formal care providers than recipients of the traditional in-kind benefit. If so, our analysis would understate the true difference in costs of consuming formal care between the two groups and our estimates would overstate the price sensitivity of demand. This particular bias is limited to some extent by the fact that members of the near-cash group could revert to the traditional in-kind benefit at any time. If the costs of securing formal care became high enough, the participant could simply switch to the in-kind benefit and let Medicaid bear those costs for them.

Previous evaluations of the experiments have found that those randomized to the near-cash benefit had health outcomes no worse than those of participants randomized to the in-kind benefit (Lepidus Carlson et al., 2007). There were 11 measures of health examined: whether the individual fell; saw a doctor for a fall; saw a doctor for a cut, burn, or scald; was injured while receiving paid help; had contractures develop or worsen; had bedsores develop or worsen; had shortness of breath develop or worsen; had a urinary tract infection; had a respiratory infection; was in poor health; and was hospitalized or in a nursing home in the previous two months. In each case, either there were no statistical differences in outcomes or those randomized to the near-cash benefit did better. Had those in the near-cash group experienced significantly worse outcomes, it would have been consistent with other costs limiting their ability to secure care. That they experienced somewhat better outcomes suggests that they did not face significantly greater costs of getting care.

C.2 Internal validity

There are two main threats to the internal validity of our estimate of the price sensitivity of demand for formal care: quantity constraints in the in-kind benefit and the distributional assumptions we make in the estimation.

Quantity constraints could have limited the formal care consumption of those receiving the traditional in-kind benefit. If recipients of the in-kind benefit faced binding quantity constraints, 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 Section 3 and Appendix Section B, the amount of Medicaid home care that someone can consume is determined by a care plan written by their 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 patient's satiation point, then Medicaid home care recipients may not be able to reach satiation. Although maximum benefit limits and care plans do not appear to have constrained consumption in our context (see Appendix Section B.4), we assess the robustness of the estimated price sensitivity to different assumptions about how binding these might have been.

Appendix Table E.4 shows estimates of the price sensitivity of demand separately for each state. The first row shows that the IV Tobit estimates range from -1.04 (Arkansas) to -2.74 (Florida). In the second row, we impose the upper bounds on care hours implied by the Arkansas and New Jersey benefit 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. The differences across states are similar to those found with the standard IV Tobit. Because average care consumption varies somewhat 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 10 percent in Arkansas, 10 percent in New Jersey, and 15 percent in Florida.

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.

The other 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 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 several different assumptions about the distribution of the error term, ε_i. In each case, we continue to instrument for price as in the main analysis. The results are reported in Appendix Table E.5. They show that the estimated price sensitivity changes somewhat from one specification to the next but not dramatically. The first three columns show results that vary the distribution of the error term while maintaining the assumption, as in the baseline specification, that observed consumption reflects a latent demand that is censored to be non-negative. The next three columns assume instead that everyone with $q_i = 0$ has a marginal value of care of exactly p_i, the maximum consistent with their behavior. Because the fraction of people with $q_i = 0$ is much greater in the cash group than in the in-kind group, this assumption increases (latent) consumption more for the cash group. This reduces the consumption difference between the cash and in-kind groups and so the implied price sensitivity. Under these distributional assumptions, we tend to find a price sensitivity around -1, though under the negative binomial assumption the price sensitivity is only -0.35. While there is some variation in the estimates, only price sensitivities far greater than any of the estimates can overturn the result that the optimal subsidy on formal care in the model in Section 6 is significantly greater than zero.

C.3 External validity

The generalizability of the results from the Cash and Counseling experiments to other contexts depends on the similarity of the policies and populations, especially in terms of characteristics that affect the price sensitivity of demand for formal care. This section discusses these issues. But as emphasized in Section 6, our main conclusions are robust to even large changes in the price sensitivity, so any issues of generalizability are less central to the key conclusions of our paper.

Appendix Table E.2 compares Cash and Counseling participants to various representative samples of Americans from the NLTCS. As discussed in Appendix Section B.2, Cash and Counseling participants are similar to the broader population of Medicaid home care recipients in terms of age and health status, but they have lower formal care consumption and are less likely to be living alone. These differences are consistent with negative selection on demand for formal care into the Cash and Counseling experiments. This is unsurprising given that the gain from a more flexible benefit is decreasing in the demand for care. Compared to the broader population of people eligible for home care benefits (column marked "2+ ADLs"), Cash and Counseling participants are in worse health, are more likely to be female, and are more likely to be unmarried. These differences are consistent with the strong selection into Medicaid home care among the eligible population of those who are sicker and who have worse informal care options, as shown in Table 3, overcoming any selection into the Cash and Counseling experiments among Medicaid home care recipients of those who are healthier and who have better informal care options.

It is unsurprising, given the incentives involved, that Cash and Counseling participants differ from the broader populations of people eligible for home care benefits and from people who take up Medicaid home care. Fortunately, what matters for the generalizability of our estimate of price sensitivity is not the level of demand for formal care, but its slope. Since little is known about this slope in different populations, in the remainder of the section we discuss what seem likely to be the most important issues.

There are two key issues that tend to offset each other. First, people whose demand was more sensitive to the composition of benefits had a greater incentive to participate in the experiment. It is therefore natural to expect that participants 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 population of recipients of Medicaid home care.

Second, the nature of the experiment—especially its unexpected occurrence and uncertain duration—likely reduced the sensitivity of demand to the composition of benefits relative to its likely value under an anticipated, permanent change in policies. Care-giving arrangements, for which people often make important investments like 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 opportunities available to them, and their decisions about where to live and work and how much formal and informal home care to consume likely depend on which home care benefits they might be eligible for. 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 it soon revert 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 welfare analysis to even large changes in the price sensitivity of demand for formal care greatly limit this concern in our context.

D Targeting Effects of In-Kind Provision: Additional Evidence from the Cash and Counseling Experiments

Those who take up Medicaid home care benefits are a highly selected subset of the population eligible for benefits, in terms of both their observable and unobservable determinants of demand for formal care (see Table 3 and Appendix Table E.6). Among those who take up Medicaid home care, recipients whose observable characteristics would normally suggest a low demand for formal care are likely to have unobservable characteristics that are strongly associated with having high demand for formal care; otherwise they would have been unlikely to take up benefits.[27] Such selection complicates comparisons of benefits received by different groups of recipients based on their observable characteristics.

For example, although being married is associated with having below-average demand for formal care in the population as a whole, among Medicaid home care recipients, being married could be associated with having above-average demand for formal care since the married people who actually take up benefits presumably have other characteristics that lead them to have a high demand for formal care. By the same logic, although in-kind provision will tend to target unmarried people relative to married people in the population as a whole, among Medicaid home care recipients in-kind provision could target married people relative to unmarried people. Whether such "reversals" arise depends on features of the joint distribution of observable and unobservable characteristics and the nature of selection into Medicaid home care and the Cash and Counseling experiments.

[27] Of course, there may be important heterogeneity in participation costs and awareness of the program as well.

Since selection could significantly bias such levels comparisons, we pursue a differences-in-differences approach that likely mitigates, though does not eliminate, this issue. We also separately analyze the subset of participants of the Cash and Counseling experiments who had not been receiving Medicaid home care before the experiments, who are likely to be more representative of the eligible population as a whole. Even so, selection issues are a major caveat to the results that follow, which provide only suggestive evidence of the effects of in-kind provision on targeting on the intensive margin. This is one reason why our preferred evidence is on targeting by formal care demand (see discussion in Section 5.3).

Using data from the Arkansas Cash and Counseling experiment, we run regressions of the form

$$benefits_i = \beta_0 + \beta_1 inkind_i + \beta_2 X_i + \beta_3(inkind_i * X_i) + \varepsilon_i \tag{1}$$

where $benefits_i$ 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 people with greater values of X_i receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do people with lower values of X_i. For example, if X_i is the number of ADL limitations, $\beta_3 > 0$ would imply that those with more ADL limitations receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do those with fewer ADL limitations. This compares the in-kind benefit to the Cash and Counseling tagged near-cash benefit. Because of the tagging, the near-cash benefit likely targets resources more than a hypothetical pure (untagged) cash transfer would. As a result, this analysis likely understates the degree to which in-kind provision targets particular groups relative to a pure cash transfer.

Appendix Table E.7 reports the effects of in-kind provision on average benefits, estimated with OLS regressions, and on the right tail of the benefit distribution, estimated with quantile regressions. 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 concentrates transfers, the OLS estimates will reflect an average of negative effects at the bottom of the benefits distribution and positive effects at the top. The quantile regressions, by contrast, estimate the effects at the top of the distribution, where targeting is likely to have the greatest impact on utility.

Column 1 shows that in-kind provision differentially targets people who are older and who have more ADL limitations. There are no significant differential targeting effects by self-rated health, sex, and marital status. In-kind provision differentially targets people who lived with others at baseline. This may be because living with others signals worse health, which may more than offset the likely effect of living with others on having better informal care options. This interpretation is consistent with the fact that those who lived with others had a greater average cost (\$129 vs. \$107 per week). Columns 2 through 4 show effects on the 90th, 95th, and 99th quantiles. In-kind provision differentially targets people with more ADL limitations, women, and the unmarried, all to a greater extent higher up in the benefits distribution.

Columns 5 through 8 repeat the analysis for the subset of participants who had not been in the Medicaid home care program at baseline. This group is likely more representative of the roughly 90 percent of eligibles who do not take up Medicaid home care. The patterns are qualitatively similar, though with larger standard errors. This is suggestive that, on the intensive margin among recipients, in-kind provision targets recipients in worse health and with worse informal care options.

E Welfare Analysis: Further Details and Robustness

E.1 The utility function, marginal utility, and optimal first-best insurance

As discussed in Section 6, the utility function nests as a special case the widely-used model in which health spending is equivalent to a wealth shock. As β approaches 0, formal care consumption approaches θ ($F(p, m; \theta) \to \theta$, ignoring corner solutions), and the indirect utility function approaches $v(p, m; \theta) = u(m - p\theta)$. For $\beta > 0$, the demand for formal care is sensitive to its price and the indirect utility function is

$$v(p, m; \theta) = \begin{cases} u\left(m - \frac{\theta^2}{2\beta}\right), & \text{if } \theta < \beta p; \\ u\left(m - p(\theta - \beta p) - \frac{\beta p^2}{2}\right), & \text{if } \theta \geq \beta p. \end{cases}$$

This differs from the benchmark case in which health spending is a wealth shock by just a slight adjustment, which is necessary to accommodate a non-zero price sensitivity of demand for formal care.

"Net consumption," non-care consumption net of any residual care costs, is

$$NC(p, m; \theta) = \begin{cases} m - \frac{\theta^2}{2\beta}, & \text{if } \theta < \beta p; \\ m - p\theta + \frac{\beta p^2}{2}, & \text{if } \theta \geq \beta p. \end{cases}$$

The targeting benefit of in-kind provision is increasing in the ratio of marginal utility in high-demand states of the world to marginal utility in low-demand states of the world. When $u(\cdot)$ is constant relative risk aversion, as in the text, the ratio of marginal utility in one state of the world relative to another is a power function of the ratio of net consumption in those states:

$$\frac{MU(\theta_H)}{MU(\theta_L)} = \left(\frac{NC(p,m;\theta_L)}{NC(p,m;\theta_H)}\right)^{\gamma}.$$

Here we show that this ratio of marginal utility in high- relative to low-demand states is decreasing in β, other things equal, and so is maximized in the limiting case in which $\beta = 0$—the standard case in the literature in which health spending is equivalent to a wealth shock. There are three cases to consider.

(i) $\theta_H \geq \theta_L \geq \beta p$: In this case,

$$\frac{NC(p,m;\theta_L)}{NC(p,m;\theta_H)} = \frac{m - p\theta_L + \beta p^2/2}{m - p\theta_H + \beta p^2/2},$$

and a marginal increase in β has the following effect on this ratio:

$$\frac{NC(p,m;\theta_H)p^2/2 - NC(p,m;\theta_L)p^2/2}{NC(p,m;\theta_H)^2} = \frac{p^2[NC(p,m;\theta_H) - NC(p,m;\theta_L)]}{2NC(p,m;\theta_H)^2} \leq 0.$$

(ii) $\theta_H \geq \beta p \geq \theta_L$: In this case,

$$\frac{NC(p,m;\theta_L)}{NC(p,m;\theta_H)} = \frac{m - \theta_L^2/(2\beta)}{m - p\theta_H + \beta p^2/2},$$

and a marginal increase in β has the following effect on this ratio:

$$\frac{NC(p,m;\theta_H)\theta_L^2/(2\beta^2) - NC(p,m;\theta_L)p^2/2}{NC(p,m;\theta_H)^2} \leq \frac{p^2[NC(p,m;\theta_H) - NC(p,m;\theta_L)]}{2NC(p,m;\theta_H)^2} \leq 0.$$

(iii) $\beta p \geq \theta_H \geq \theta_L$: In this case,

$$\frac{NC(p,m;\theta_L)}{NC(p,m;\theta_H)} = \frac{m - \theta_L^2/(2\beta)}{m - \theta_H^2/(2\beta)},$$

and a marginal increase in β has the following effect on this ratio:

$$\frac{NC(p,m;\theta_H)\theta_L^2/(2\beta^2) - NC(p,m;\theta_L)\theta_H^2/(2\beta^2)}{NC(p,m;\theta_H)^2} = \frac{m(\theta_L^2 - \theta_H^2)}{2NC(p,m;\theta_H)^2\beta^2} \leq 0.$$

Increasing β reduces the ratio of net consumption in low- relative to high-demand states, which reduces the ratio of marginal utility in high- relative to low-demand states, which reduces the targeting benefit of in-kind provision. As a result, the baseline case with $\beta > 0$ contains a weaker link between demand for formal care and marginal utility—and so a smaller targeting benefit from in-kind provision—than the standard model in which health spending is equivalent to a wealth shock.

To better understand the utility function, the nature of the risk the individual faces, and desired insurance transfers, consider the benchmark of a first-best insurance program. The first-best transfer schedule satisfies:

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

where B is expected spending on someone eligible for home care benefits and $b(B)$ is the cash transfer that makes expected spending equal B. The first-best transfer is increasing in θ, 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 the risk. By making greater transfers in states of the world with greater demand for formal care, it fully compensates the individual both for her expenditures on formal care and for any residual utility costs she faces from coping with her health problems.

E.2 Estimating the distribution of demand for formal care

As discussed in the text, we use our estimate of the price sensitivity of demand for formal care, β, to convert the observed joint distribution of formal care consumption and formal care prices in the NLTCS into a distribution of the level of demand for formal care in the benefit-eligible population, $G(\theta)$. We express the level of demand for formal care in terms of satiation points, θ. The main challenge is that observed formal care consumption does not point-identify θ for people consuming zero formal care, it only bounds it: $\theta_i \leq \beta p_i$. We estimate the full θ distribution, including the θ'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 as well as 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), we draw the underlying latent q^* from the truncated triangle distribution based on the estimated slope. Appendix 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 θ using the estimated price sensitivity of demand for formal care, $\theta_i = q_i^*(p) + \hat{\beta}p$. 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 θ distribution. Figure 6 shows the resulting θ 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.

In the quantitative analysis, we further constrain θ to be non-negative and, as a baseline, no larger than 150 hours per week. A negative satiation point is not implausible in theory; someone might wish to consume no formal care even if they were paid to consume it. But a negative satiation point is awkward in practice with the baseline utility function, since someone with $\theta < 0$ would be worse off than someone with $\theta = 0$. Moreover, behavior when $\theta < 0$ is identical to behavior when $\theta = 0$ as long as the net-of-subsidy price of formal care is non-negative. We truncate the baseline θ distribution at 150 hours per week in order to reduce the influence of outliers. To the extent that such large values are valid, excluding them tends to reduce the targeting benefit relative to the distortion cost and so leads us to understate the optimal subsidy. Given the importance of right-tail risks for insurance, we also report results under different assumptions about the right tail of the θ distribution.

We test the robustness of our results to making a worst-case assumption about the unidentified θ values. We set all of the partially-identified θ's to their (point-identified) upper bound, $\theta_i = \hat{\beta}p_i$.

E.3 State-dependent utility

Any state-dependence in utility that is correlated with the demand for formal care affects the value of in-kind provision by affecting the value of targeting states of the world with greater demand for formal care. State dependence that increases marginal utility in states with greater demand for formal care relative to states with lower demand for formal care increases the attractiveness of in-kind formal care transfers, whereas state dependence that decreases marginal utility in states with greater demand for formal care relative to states with lower demand for formal care decreases the attractiveness of in-kind formal care transfers

People in worse health likely have different utility functions from people in better health; they likely have a lower level of utility, for example. But what matters for insurance is marginal utility, and a priori it is not clear in which direction a reduction in health might shift marginal utility. On one hand, activities like eating out and traveling likely become less attractive, which tends to reduce marginal utility. On the other hand, home upgrades and equipment likely become more attractive, which tends to increase marginal utility.

The importance of state-dependent utility for our analysis is lessened by the nature of our counterfactuals of interest, which vary the type of benefit available to people in bad health (those with two or more ADL limitations) while holding fixed spending on these bad health states as a whole. Since home care benefits are limited to states of the world with fairly severe chronic health problems, the relative marginal utility of healthy versus sick people is irrelevant; only relative marginal utility within bad-health states matters. Although this lessens the likely importance of state-dependent utility in our context, we test the robustness of our results to different possibilities about state-dependent utility within bad-health states.

We analyze the effects of state-dependent utility based, as closely as possible, on the estimates of Finkelstein et al. (2013). Finkelstein et al. (2013) estimate the state-dependence of utility in the number of chronic health problems someone has.[28] It is important to emphasize that their estimates do not map perfectly to our context, whether to the level of demand for formal care, θ, or to the number of ADL limitations someone has. But it is the best evidence on the likely extent of state-dependent utility in a related context.

Finkelstein et al. (2013) estimate that a one-standard deviation increase in the number of chronic health problems is associated with a 10–25 percent decline in marginal utility. We adapt this evidence to our setting by assuming that a one-standard deviation increase in the number of chronic health problems corresponds to a one-standard deviation increase in the level of demand for formal care, θ. We assume that $U(c;\theta) = \mu(\theta)u(c)$, with $\mu(\theta)$ linearly decreasing in θ at a rate that corresponds to the upper endpoint of their preferred range of estimates. So the marginal utility multiplier, $\mu(\theta)$, decreases by 25 percent for every one-standard deviation increase in the demand for formal care, θ.

E.4 Additional robustness tests

Appendix Table E.8 shows the results of the welfare analysis under the baseline specification and five other specifications not shown in the main text. Consistent with the results in Table 4, the results in Appendix Table E.8 show that the welfare gain from in-kind provision is

[28] Viscusi and Evans (1990) and Evans and Viscusi (1991) also estimate the state-dependence of utility in health, but they do so for a younger, less disabled population.

highly robust to plausible changes in the model. Column 2 shows that even a price sensitivity of demand for formal care over five times larger than that consistent with the Cash and Counseling experiments ($\beta = 10$) does not overturn the conclusion that the optimal subsidy is large. Columns 3 and 4 show the importance of the right tail of the distribution of demand for formal care in determining the targeting benefit and so the optimal subsidy. The longer the right tail, the more valuable is the insurance benefit of in-kind provision. But the optimal subsidy remains large even when the right tail of the distribution is chopped off or when all of the θ values are scaled down (as shown in Table 4). Column 5 shows that even if the distribution of partially-identified θ values is in the "worst-case" configuration (i.e., each θ_i equals the maximum value consistent with i's behavior), the optimal subsidy rate is still 86 percent. Column 6 shows the importance of the consumption floor. Cutting the level of the floor in half, from $5,000 to $2,500, increases the equivalent variation gain from the optimal policy substantially, from $6,416 to $21,854.

References

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.

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

Evans, W. N. and W. Viscusi (1991). Estimation of state-dependent utility functions using survey data. *Review of Economics and Statistics 73*(1), 94–104.

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

Finkelstein, A. and M. J. Notowidigdo (2018, May). Take-up and targeting: Experimental evidence from SNAP. NBER Working Papers 24652, National Bureau of Economic Research, Inc.

Hastings, J. S. and J. M. Shapiro (2017, January). How are SNAP benefits spent? Evidence from a retail panel. NBER working paper 23112, National Bureau of Economic Research, Inc.

Kleven, H. J. and W. Kopczuk (2011, February). Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy 3*(1), 54–90.

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.

Ng, T., C. Harrington, and J. Howard (2011). *Medicaid home and community-based service programs: Data update*. Henry J. Kaiser Family Foundation.

Ng, T., C. Harrington, M. Musumeci, et al. (2016). Medicaid home and community-based service programs: 2013 data update. *Issue Paper. Washington, DC: Kaiser Commission on Medicaid and the Uninsured*.

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

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

Sommers, B., M. Roberts Tomasi, K. Swartz, and A. M. Epstein (2012). Reasons for the wide variation in Medicaid participation rates among states hold lessons for coverage expansion in 2014. *Health Affairs 31*(5), 909–919.

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 80*(3), 353–374.

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

Appendix Figures and Tables

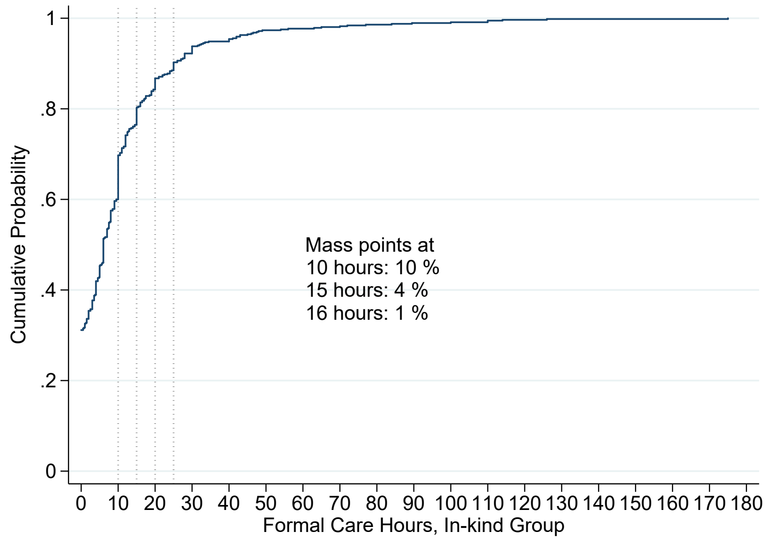

Figure E.1: CDF of Formal Care Consumption in Cash and Counseling, Arkansas

[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 in principle limited formal care benefits to 16 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

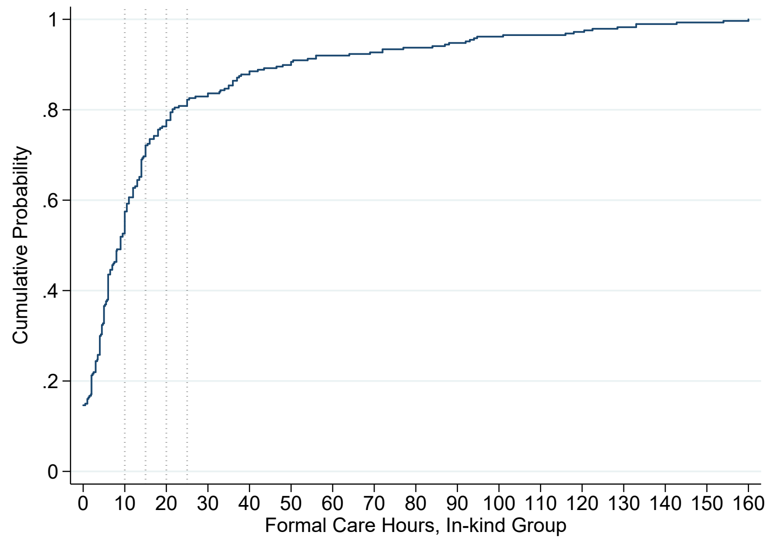

Figure E.2: CDF of Formal Care Consumption in Cash and Counseling, Florida

[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 formal care benefits (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

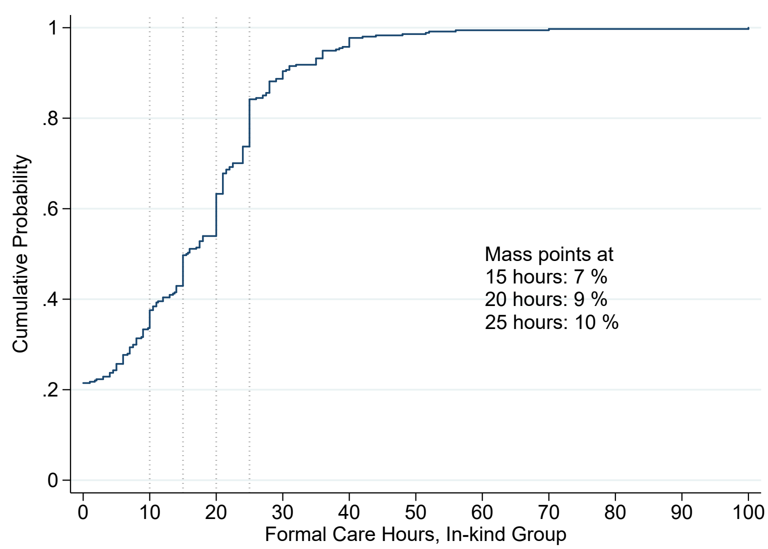

Figure E.3: CDF of Formal Care Consumption in Cash and Counseling, New Jersey

[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 in principle limited formal care benefits to 25 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]

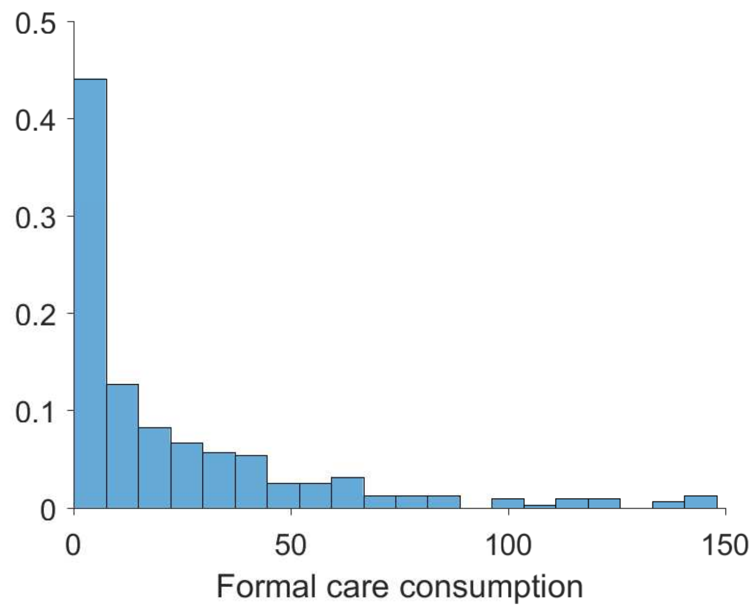

Figure E.4: Distribution of Formal Care Consumption in Benefit-Eligible Population

[Empirical density of formal care consumption among the non-institutionalized population aged 65 and older with two or more ADL limitations. Data from the NLTCS. For readability the figure omits the 63 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. The mean of the full distribution is 12.5 hours per week.]

Table E.1: Summary Statistics and Balance Tests for the Cash and Counseling Experiments

	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	9.05	9.02	0.96	12.99	13.01	0.99	16.22	15.56	0.52
Number unpaid caregivers, baseline	2.19	2.12	0.42	1.95	2.05	0.42	2.04	2.10	0.64
Age	78.94	79.02	0.87	79.02	79.89	0.18	77.54	77.78	0.67
Male	0.17	0.17	0.93	0.18	0.21	0.47	0.18	0.22	0.17
White	0.62	0.64	0.38	0.67	0.71	0.27	0.50	0.56	0.11
Less than high school degree	0.66	0.67	0.98	0.35	0.37	0.53	0.66	0.65	0.77
High school degree	0.28	0.26	0.45	0.47	0.47	1.00	0.18	0.20	0.53
College degree or more	0.03	0.05	0.10	0.16	0.14	0.51	0.10	0.11	0.67
Health, baseline	3.19	3.22	0.54	3.13	3.06	0.28	3.19	3.16	0.69
Lives alone, baseline	0.32	0.31	0.64	0.25	0.31	0.16	0.33	0.38	0.19
Unmarried	0.85	0.85	0.94	0.77	0.81	0.27	0.79	0.76	0.28
Observations	564	565	.	302	287	.	368	354	.

Means by state and treatment assignment. P-values test the equality of means of the cash and in-kind groups within each state. Formal care hours (per week), Number unpaid caregivers, Health, and Lives alone are all measured in the baseline survey, at the time of randomization. Remaining variables are measured at the nine-month followup.

Table E.2: Summary Statistics for NLTCS and Cash and Counseling Samples

	NLTCS					Cash and Counseling
	(1) All 65+	(2) 2+ ADLs	(3) Eligible for Medicaid home care	(4) 2+ ADLs, not on Medicaid home care	(5) 2+ ADLs, on Medicaid home care	(6) All 65+
Formal care hours/week	2.87	12.52	16.27	9.28	35.91	12.04
Age	77.43	79.14	80.14	79.14	79.89	78.76
Number of ADLs*	0.70	3.66	3.67	3.55	4.16	2.11
Health fair or poor	0.41	0.70	0.70	0.68	0.77	0.79
Female	0.63	0.66	0.70	0.65	0.74	0.82
Lives alone	0.36	0.25	0.32	0.22	0.39	0.32
Unmarried	0.53	0.55	0.61	0.52	0.70	0.81
Has children	0.81	0.78	0.74	0.80	0.80	.
Income ($/month)	1242.71	1149.27	813.87	1197.79	835.37	.
Observations	5,147	887	460	736	113	2,470

Summary statistics for NLTCS and Cash and Counseling samples. Columns denote sample restrictions. Column 1 includes all non-institutionalized individuals aged 65 and older in the NLTCS data. Column 2 additionally restricts the sample to those with at least two ADLs. Column 3 restricts to those who are eligible for Medicaid home care based on our "Income eligible, < 2 cars" criteria. Column 4 further restricts to those with 2+ ADLs who are not on Medicaid home care. Column 5 restricts to those with 2+ ADLs who are on Medicaid home care. The sum of the observations in columns 4 and 5 is lower than those in column 2 because of missing values for Medicaid home care participation. Column 6 restricts the sample to those in the Cash and Counseling data who are at least 65 years of age. Cash and Counseling data are from the baseline survey. Formal care is in hours per week. The Cash and Counseling surveys did not collect information on the number of children or income.

*The "Number of ADLs" measures, unlike the other variables, are not directly comparable across datasets. In the NLTCS, this is the total number of affirmative answers to *six* questions about *needing* help: eating, getting out of bed, getting around, dressing, bathing, and using the toilet. In Cash and Counseling, this is the total number of affirmative answers to *three* questions about *receiving* help: getting out of bed, bathing, and using the toilet.

Table E.3: Price Sensitivity of Demand for Formal Care, First Stage Estimates

	(1)	(2)
Assigned to near-cash	7.68	7.65
	(0.23)	(0.23)
Controls	No	Yes
F-Statistic	1,139	1,144
Mean market price	13.68	13.68
Adjusted R-squared	0.32	0.34
Observations	2,440	2,440

Dependent variable is the marginal price of formal care. Data are from the Cash and Counseling experiments. Controls included in column (2) are indicators for sex, education level, race, self-rated health, five-year age bins, and state. Robust standard errors reported.

Table E.4: Price Sensitivity of Demand for Formal Care and Statutory Limits

	(1) Arkansas	(2) Florida	(3) New Jersey
Price, IV Tobit	-1.04	-2.74	-1.61
	(0.22)	(0.42)	(0.15)
Price, IV Tobit Limits	-0.53		-1.78
	(0.12)		(0.16)
Controls	Yes	Yes	Yes
Market price, formal care	12.36	15.09	14.59
Mean hours, in-kind group	10.76	18.60	16.10
Observations	1,129	589	722

Dependent variable is formal care consumption in hours per week. Data are from the Cash and Counseling experiments. Seperate regressions are run for each state. First row is IV Tobit. Second row is IV Tobit with statutory limit as upper bound. There is no statutory limit in Florida. All regressions control for sex, education level, race, self-rated health, five-year age bins, and state. Robust standard errors reported.

Table E.5: Robustness of Price Sensitivity to the Distribution of the Error Term, ε_i

	Censored models (observe $q_i = \max\{0, q_i^*\}$)			Uncensored models (assume $q_i^* = q_i \geq 0$)		
	(1) Normal	(2) Extreme Value	(3) T-location scale	(4) Normal	(5) Negative binomial	(6) Poisson
Price	-1.78	-2.15	-1.23	-0.96	-0.35	-1.09
	(0.15)	(0.23)	(0.07)	(0.08)	(0.12)	(0.02)
Mean hours	10.48	10.48	10.48	10.48	10.48	10.48
Observations	2,440	2,440	2,440	2,440	2,440	2,440

Dependent variable is formal care consumption in hours per week. Data are from the Cash and Counseling experiments. Each column presents the estimated price sensitivity of demand under a different distributional assumption on the underlying error term. Column 1 is the baseline specification. Columns 1-3 assume, as in the baseline, that observed consumption is censored to be non-negative. Columns 4-6 assume that everyone with $q_i = 0$ has a marginal value of care of exactly p_i, the maximum consistent with their behavior (i.e., no censoring). All models instrument for price with the participant's randomized treatment status and are estimated via two-stage residual inclusion. Columns 5 and 6 report average marginal effects.

Table E.6: Level of Demand for Formal Care Among Those Who Do Vs. Do Not Take Up Medicaid Home Care

	(1) OLS	(2) 90th	(3) 95th	(4) 99th
Medicaid home care	12.57	10.85	3.39	83.42
	(7.11)	(21.70)	(25.51)	(45.61)
Age	0.60	0.02	-0.12	2.92
	(0.20)	(0.50)	(0.75)	(1.61)
Four or more ADLs	12.82	44.03	77.03	22.65
	(4.16)	(22.28)	(27.53)	(35.94)
If health fair or poor	-3.55	-8.61	-12.09	19.26
	(4.42)	(17.25)	(17.03)	(21.48)
Female	4.46	0.64	3.58	-32.40
	(4.31)	(6.06)	(9.21)	(30.64)
Lives alone	9.94	46.69	25.32	-15.21
	(6.11)	(22.06)	(25.20)	(32.10)
Unmarried	8.26	19.25	42.88	75.62
	(4.76)	(13.96)	(25.41)	(36.52)
Has children	5.31	8.11	5.65	29.18
	(5.71)	(13.47)	(15.02)	(27.13)
Income	-0.00	-0.00	-0.01	0.00
	(0.00)	(0.01)	(0.01)	(0.02)

Dependent variable is price-adjusted formal care consumption, in hours per week. Price-adjusted formal care consumption uses our estimate of the price sensitivity of demand to simulate each individual's consumption if she were to face a price of $18.50, the maximum in the data. The sample is those eligible for Medicaid home care, based on the "Income eligible, < 2 cars" measure. The sample has 448 observations. Column 1 reports results from an OLS regression with robust standard errors. Columns 2-4 present results from quantile regressions, with the quantile specified in the column heading, with bootstrapped standard errors. * p<0.10, ** p<0.05, *** p<0.01

Table E.7: 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)

Each entry is the estimate of β_3, the coefficient on $inkind_i * X_i$, in equation 1 (described in Appendix Section D) of separate regressions. The dependent variable is the weekly cost of benefits received in dollars. The estimates reveal whether those with more of the characteristic receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do those with less of the characteristic. Data from the Arkansas Cash and Counseling experiment. Columns (1) through (4) include all participants. Columns (5) through (8) only include the subset who had not been enrolled in Medicaid home care before the experiment. Columns (1) and (5) are OLS regressions with robust standard errors. Remaining columns are quantile regressions with bootstrapped standard errors. The omitted health category is health good or excellent.

Table E.8: Welfare Analysis Robustness

	(1) Baseline	(2) $\beta = 10$	(3) Drop $\theta > 100$	(4) Keep $\theta > 150$	(5) Worst-case low θs	(6) Halve consumption floor
Optimal subsidy, s^*	0.87	0.82	0.79	0.95	0.86	0.89
EV gain over pure-cash policy	\$6,416	\$3,210	\$4,723	\$9,535	\$5,276	\$21,854
E(ex post value)/E(cost)	0.49	0.71	0.62	0.34	0.51	0.51
Corr(marg. util., formal care)	0.89	0.74	0.87	0.90	0.88	0.84

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 formal care). "EV gain over pure-cash policy" is the ex ante equivalent variation gain of the optimal policy over an equal-cost pure-cash policy. "E(ex post value)/E(cost)" is the ratio of the mean ex post value of the optimal benefit to its mean cost. This is an inverse measure of the distortion cost of the optimal policy. "Corr(marg. util., formal care)" is the correlation between marginal utility and formal care consumption in the absence of insurance (under a pure-cash policy). This is a measure of how well in-kind provision targets relatively high-marginal utility states. Column 1 corresponds to the baseline assumptions. Column 2 increases the price sensitivity of demand from the baseline estimate of 1.8 to 10. Column 3 truncates the right tail of the risk by dropping all values of θ greater than 100. Column 4 leaves the right tail of the θ distribution as is, whereas the baseline specification drops all values of θ greater than 150. Column 5 sets the partially-identified θ values at the bottom of the distribution (corresponding to individuals who consumed no formal care) to the maximum value consistent with the data, which reduces the gain from in-kind provision. Column 6 cuts the consumption floor in half from \$5,000 to \$2,500.