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

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

September 28, 2018

Abstract

Making a transfer in kind reduces its value to recipients but can improve targeting. We develop an approach to quantifying this tradeoff and apply it to home care. Using randomized experiments by Medicaid, we find that in-kind provision significantly reduces the value of the transfer to recipients while targeting a small fraction of the eligible population that is sicker and has fewer informal caregivers 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 toward more flexible benefits.

*Lieber: Department of Economics, 3049 Nanovic Hall, University of Notre Dame, Notre Dame, IN 46556 (Ethan.Lieber.2@nd.edu); Lockwood: Department of Economics, 248 McCormick Road, University of Virginia, Charlottesville, VA 22903, and NBER (leelockwood@virginia.edu). We are grateful to the Editor, Liran Einav, the anonymous referees, Norma Coe, Manasi Deshpande, Amy Finkelstein, Gopi Shah Goda, Tami Gurley-Calvez, Seema Jayachandran, Brian Melzer, Matt Notowidigdo, Mike Powell, Emmanuel Saez, Diane Schanzenbach, Jesse Shapiro, Courtney Van Houtven, our department colleagues, and many seminar participants for helpful comments. We thank Jose Carreno, Vishal Kamat, Max Rong, 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. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

In-kind transfers are a ubiquitous feature of government programs, private contracts, and charitable giving. In the United States, government spending on in-kind programs exceeds 12 percent of GDP and spending on in-kind health programs alone exceeds $1 trillion per year (Currie and Gahvari, 2008; Centers for Medicare and Medicaid Services, 2017). In domestic policy, foreign aid, and charitable giving, there are active debates about the desirability of flexible benefits such as direct cash transfers and universal basic income programs versus restrictive in-kind transfers of food, housing, medical care, and other goods.

Central to these debates is a tradeoff inherent to in-kind transfers. In-kind provision has a fundamental cost: Recipients would prefer an equal-cost cash transfer. But this cost is linked to an important potential benefit: In-kind provision can better target desired recipients by leading some people to take up more benefits than others (Nichols and Zeckhauser, 1982). In the context of insurance, if someone values a particular good more in states of the world in which marginal utility is higher, an in-kind transfer of that good can help concentrate benefits in those states and thereby better insure the risk. Although these costs and benefits of in-kind provision are crucial determinants of optimal policy, little is known about their relative magnitudes across a wide range of important contexts.

In this paper, we develop an approach to quantifying this core tradeoff of in-kind provision and apply it to home care. Home care helps people who have chronic health problems with tasks such as eating, dressing, and bathing. Its value, including care from family and friends ("informal care") as well as professional caregivers ("formal care"), is thought to exceed $200 billion per year (Arno et al., 1999). Traditionally, home care benefits have been provided as in-kind formal care. But following Medicaid's large-scale Cash and Counseling experiments in the late 1990s, many states reformed their home care programs to make benefits more flexible and cash-like (National Conference of State Legislatures, 2007).

Our approach to quantifying the welfare effect of in-kind provision involves three main ingredients. The first is the moral hazard effect, the extent to which in-kind provision increases consumption of the good. The greater this increase, the lower the value to recipients of the in-kind benefit relative to its cost. Using the randomized assignment of in-kind versus near-cash benefits in the Cash and Counseling experiments, we estimate that in-kind provision increases formal care consumption among those consuming formal care by 25 hours per week—nearly twice the average consumption in the benefit-eligible population. This suggests that many recipients value the in-kind benefit far below its cost. Our estimates imply that a recipient of the average in-kind transfer in the experiment values it at 28 percent of its cost.

The second ingredient is the distribution of consumption of the good within benefit-eligible states of the world. The greater the heterogeneity in consumption of the good, the greater the extent to which in-kind provision concentrates transfers. Using nationally-representative data on the eligible population, we find considerable heterogeneity in consumption of formal care. While 63 percent of those eligible do not consume any formal care, among those who do, there is a long right tail. An individual at the 95th percentile receives around-the-clock care, which at the average hourly price of $15 amounts to about $131,000 per year (Genworth Financial, 2005).

The third ingredient is the link between consumption of the good and the marginal utility of income. The stronger this link, the more valuable it is to shift resources to the states of the world targeted by in-kind provision. In our context, this link is likely strong. Greater costs of coping with bad health leave fewer resources for non-health consumption, which tends to increase marginal utility. Empirically, we find that in-kind provision sharply concentrates transfers on a small fraction of the eligible population that has a greater demand for formal care, is sicker, and has fewer informal caregivers than the average eligible. To the extent that such recipients tend to have relatively high marginal utility, in-kind provision could significantly improve insurance.

These results suggest that designers of home care benefits face a stark tradeoff. Restrictive in-kind benefits are valued far less than their cost ex post, but they sharply concentrate transfers in what appear to be relatively high-marginal utility states. We combine our reduced-form estimates with a structural model to quantify this tradeoff in a stylized expected utility framework. Under a wide range of assumptions, the optimal contract involves a large in-kind component and delivers substantial welfare gains over a cash-benefit contract.

Our paper complements and extends the literature on barriers to private, voluntary long-term care insurance (see Brown and Finkelstein, 2011, for a review). Our findings reveal the critical importance of two factors in determining the welfare effect of any long-term care insurance, whether public or private, voluntary or mandatory: risk within unhealthy states of the world and moral hazard. Although in-kind provision has a large moral hazard cost, the gain from insuring the considerable risk within unhealthy states appears even larger. This raises concerns about recent reforms toward cash-like benefits.

Our approach helps link the theoretical and empirical literatures on in-kind transfers, which have been largely disconnected so far (Currie and Gahvari, 2008).[1] Methodologically,

[1]The theoretical branch has investigated potential advantages of in-kind transfers in terms of paternalism (Musgrave, 1959), targeting (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988), tax system efficiency (Munro, 1992), and the Samaritan's Dilemma (Bruce and Waldman, 1991). Much of the empirical branch estimates the effects of in-kind transfers on consumption (e.g., see Hoynes and Whitmore Schanzenbach, 2016, for a review of the effects of food transfers). Other work examines the effects of in-kind transfers

the equivalence of the effects of an in-kind transfer and a corresponding price subsidy on a recipient's choice set allows us to use ideas from the literatures on optimal taxation and health insurance to quantify a core tradeoff of in-kind provision. Substantively, the key feature of in-kind transfers that gives rise to their targeting and distortion effects is that they reduce the recipient's cost of consuming the good over some range of quantities, thereby "loosening" the budget constraint more for recipients who consume more of the good.[2] This feature is shared by a wide range of other policies, including vouchers, conditional cash transfers, benefit programs with ordeals, insurance policies, and commodity taxes and subsidies.

This paper also contributes to the literature on targeting in benefit programs such as housing assistance (Reeder, 1985), Medicaid (Cutler and Gruber, 1996), Supplemental Security Income (Benitez-Silva et al., 2004), disability insurance (Low and Pistaferri, 2015; Deshpande and Li, 2017), and food stamps (Finkelstein and Notowidigdo, 2018).[3] In many programs, only a small fraction of the eligible population takes up benefits. While low take up can be undesirable, our findings suggest that it can also significantly increase welfare through better targeting.

2 Approach

This section describes an approach to quantifying the targeting-distortion tradeoff of in-kind provision. This tradeoff has previously been analyzed in theory (e.g., Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988), using models that provide clear insights about the economic factors involved but that are not well-suited to empirical implementation. Our approach to empirical implementation has close parallels in the literatures on optimal taxation and health insurance (e.g., Zeckhauser, 1970; Mirrlees, 1971; Manning and Marquis, 1996; Saez, 2001). These parallels arise from an economic equivalence: For any in-kind transfer, there is a subsidy that has the same effect on the recipient's budget set.[4]

The key feature of an in-kind transfer that gives rise to the targeting-distortion tradeoff

on poverty (Smeeding, 1977), targeting (Reeder, 1985; Cutler and Gruber, 1996; Jacoby, 1997), and prices (Cunha et al., 2011). Our approach complements those of Finkelstein et al. (2015), who analyze the welfare effect of Medicaid health insurance coverage for prime-age adults.

[2]In-kind transfers that are inframarginal for most potential recipients are unlikely to have large targeting and distortion effects. This seems likely to be the case for the food transfer program in the U.S., though there is an ongoing debate about the effect of the transfer on patterns of spending (Hoynes and Whitmore Schanzenbach, 2016).

[3]See Currie (2006) for a review. A related literature in the developing world investigates the targeting effects of ordeals (Alatas et al., 2016), subsidized prices (Cohen and Dupas, 2010), and delegating authority over the distribution of benefits to local leaders (Alatas et al., 2012; Basurto et al., 2017). Kleven and Kopczuk (2011) analyze the role of program complexity in determining take up.

[4]For example, an in-kind transfer that offers recipients up to a fixed amount of the good free of charge has the same effect as a non-linear subsidy of 100 percent up to that fixed amount and 0 percent thereafter.

is that it reduces the recipient's cost of consuming the transferred good. One consequence is that recipients over-consume the good and value the transfer less than its cost to the provider. This is the moral hazard cost of in-kind provision. Another consequence is that, because the cost reduction is more valuable to someone consuming more of the good, it targets states of the world or types of people with relatively high consumption of the good. If these states or types have relatively high marginal utility, that is a targeting benefit of in-kind provision.

2.1 Theory

An individual faces a risk that potentially affects prices, income, and preferences. The eventual state of the world is uncertain ex ante and non-contractible ex post. As a result, an insurance contract cannot target high-marginal utility states directly by offering larger benefits in those states. Instead, any targeting must be indirect, relying on differential take up of a single benefit. A natural candidate is an in-kind transfer of an "indicator good," a good consumed in greater quantities in higher-marginal utility states of the world (Nichols and Zeckhauser, 1982).

Consider an in-kind transfer with no quantity limit, a linear subsidy. An increase in the subsidy rate reduces the after-subsidy price to consumers. We focus on this case for simplicity and because it best matches our empirical application. Appendix Section A.1 analyzes the case with a binding benefit limit; the core tradeoff is the same.

Ex post indirect utility in the realized state of the world is

$$v(p,m) \equiv \max_{x_k, x_{-k}} u(x_k, x_{-k}) \text{ subject to } p_k x_k + \sum_{i \neq k} p_i x_i \leq m,$$

where p is the vector of prices, m is income, x_k is the good being transferred in kind, and x_{-k} is the vector of all other goods. By the envelope theorem (Roy's Identity), the ex post marginal value of a reduction in the price of good k is the individual's consumption of good k in that state: $\frac{-\partial v(p,m)}{\partial p_k} \Big/ \frac{\partial v(p,m)}{\partial m} = \frac{\lambda x_k}{\lambda} = x_k$, where λ is the marginal utility of income. The individual's consumption of x_k is the amount by which the price reduction "loosens" the individual's budget constraint in that ex post state.

The ex ante expected marginal benefit of a reduction in the price of good k is

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

where $E\left(v(p,m)\right)$ is expected (indirect) utility and $\widehat{\lambda}$ is the marginal utility of income normalized so that its mean is one.[5] The ex ante value of the price reduction is its mean ex post value, $E(x_k)$, plus a "correction term" for the relationship, if any, between marginal utility and consumption of good k, $Cov\left(\widehat{\lambda}, x_k\right)$. This term arises because the benefits of the subsidy are greater in states with greater consumption of good k. If $Cov\left(\widehat{\lambda}, x_k\right) > 0$, the targeting of benefits to states with greater consumption of good k also tends to target benefits to high-marginal utility states, providing insurance that a cash transfer does not. In this case, the ex ante expected marginal benefit of the subsidy exceeds its mean ex post value. This covariance term—the insurance value of the subsidy's differential targeting of states with greater consumption of good k—is the targeting benefit of in-kind provision.

The expected cost to the insurer of the in-kind benefit is $(p_k^0 - p_k)E(x_k)$, where p_k^0 is the un-subsidized price and p_k is the consumer's net-of-subsidy price.[6] The marginal cost to the insurer of a reduction in the price of good k is

$$MC = E\left(x_k\right) + (p_k^0 - p_k)E\left(-\frac{dx_k}{dp_k}\right).$$

The first term is the insurer's additional spending due to the increase in the subsidy rate, holding fixed consumption ("mechanical effect"). The second term is the insurer's additional spending on the subsidy due to the induced change in consumption ("moral hazard effect").[7]

Figure 1 plots the marginal cost, marginal benefit, and mechanical effect of reductions in the price of good k as functions of the subsidy rate, s, where $p_k = (1 - s)p_k^0$. The targeting benefit is the vertical distance from the mechanical effect to the marginal benefit. The distortion cost is the vertical distance between the mechanical effect and the marginal cost. The marginal cost of the subsidy exceeds its mean ex post value (the mechanical effect) due to moral hazard; in each state, the subsidy is less valuable than an equal-cost

[5]The second equality comes from the envelope theorem. The final equality comes from noting that $E(\lambda x_k) = E(\lambda)E(x_k) + Cov(\lambda, x_k)$ and that $Cov(\lambda, x_k)/E(\lambda) = Cov(\lambda/E(\lambda), x_k) \equiv Cov\left(\widehat{\lambda}, x_k\right)$.

[6]This assumes that the supply of every good is perfectly elastic. In this case, an increase in the subsidy reduces the individual's after-subsidy price of good k one-for-one and has no effect on the prices of other goods. This marginal cost does not include second-best considerations from other distortions in the economy, such as substitution from subsidized nursing home care. The problem can therefore be viewed as that of a private insurer offering a stand-alone home care benefit, which would not account for such effects. We discuss the likely impact of some of the main second-best considerations in the conclusion.

[7]The policy counterfactual, in particular who pays for the subsidy, affects the size of the moral hazard effect. The moral hazard effect is the *total* derivative, i.e., the combined effect of the price reduction and any accompanying change in nominal income, $\frac{dx_k}{dp_k} = \frac{\partial x_k(p,m)}{\partial p_k} + \frac{\partial x_k(p,m)}{\partial m} m'(p_k)$, where $x_k(p, m)$ is Marshallian demand for good k and $m'(p_k)$ is the accompanying change in nominal income. We focus on cost-neutral shifts in a mixed in-kind/cash benefit, which pair an increase in the subsidy with a reduction in the uniform cash benefit that holds fixed total spending on recipients. Such cost-neutral shifts isolate the welfare effect of in-kind provision from that of redistribution between recipients and other parties. This means that the marginal cost and marginal benefit are in the same units: income in the hands of recipients.

cash benefit due to the change in consumption it induces. At the optimum, the marginal targeting benefit equals the marginal distortion cost and both exceed the mechanical effect. The optimal contract leaves some risk uninsured since the benefit of insuring it is smaller than the cost.

2.2 Empirical Implementation

Our approach to quantifying the targeting-distortion tradeoff is based on three ingredients: the price-sensitivity of demand for the good, dx_k/dp_k; the distribution of consumption of the good within benefit-eligible states of the world, $F(x_k)$; and the link between consumption of the good and marginal utility.

The price-sensitivity of demand determines the moral hazard cost of in-kind provision, the excess of the cost of the benefit over its value to recipients ex post.

The remaining two ingredients, the distribution of consumption of the good and the link between consumption of the good and marginal utility, determine the targeting benefit of in-kind provision. Letting σ_X be the standard deviation of variable X across states of the world, we can decompose the marginal targeting benefit as $Cov\left(\widehat{\lambda}, x_k\right) = \sigma_{x_k}\, \sigma_{\widehat{\lambda}}\, Corr\left(\widehat{\lambda}, x_k\right)$. The distribution of consumption of the good determines σ_{x_k}. This in turn determines the extent to which in-kind provision concentrates benefits in some states and not others since the ex post marginal benefit of a shift toward in-kind provision is proportional to x_k. Given the distribution of consumption of the good, the link between consumption of the good and marginal utility determines both the extent of risk ($\sigma_{\widehat{\lambda}}$) and the extent to which the states targeted by in-kind provision have relatively high-marginal utility $\left(Corr\left(\widehat{\lambda}, x_k\right)\right)$. This decomposition splits the targeting benefit into two parts: (i) the *targeting effect*, the estimable effect on the distribution of transfers, and (ii) the value of this targeting, which depends on the unobservable link between consumption of the good and marginal utility.

This decomposition isolates assumptions about marginal utility from the rest of the analysis. It facilitates analyses that incrementally build from reduced-form estimations that shed light on the key magnitudes to sufficient-statistics and structural approaches that quantify the net welfare effect. Without any assumptions about marginal utility, straightforward estimations of the price sensitivity of demand and the distribution of consumption reveal the moral hazard cost and targeting effect of in-kind provision. With a qualitative sense of the link between consumption of the good and marginal utility—as presumably exists for a good being transferred in kind—the distribution of consumption is also informative about the extent of risk and the potential targeting benefit of in-kind provision. With a model of marginal utility, the net welfare effect can be quantified as well. The theoretical considerations and empirical evidence that can help inform this important modeling choice will vary by context. Appendix Section A.2 discusses the applicability of the approach.

3 Home Care, Medicaid, and the Cash and Counseling Experiments

Chronic health problems are one of the most important risks people face over the life cycle. Roughly 15 percent of Americans over age 50 have at least one person helping them perform activities of daily living (ADL) such as bathing, eating, and dressing (Barczyk and Kredler, 2017). Eighty-seven percent of those receiving help live in the community and 74 percent of all care hours occur in private homes (Barczyk and Kredler, 2017). Spending on formal home care was $88 billion in 2015, and the total cost of home-based care, including (hard-to-measure) informal care from family and friends, is thought to exceed the total cost of formal long-term care (Arno et al., 1999; Centers for Medicare and Medicaid Services, 2017). Despite the magnitude of this risk, just 10 percent of people 65 and older own private long-term care insurance. As a result, a large share of the costs of long-term care in general and home care in particular are paid by the means-tested Medicaid program.

Medicaid home care programs are an important source of care for many people. In 2013, Medicaid spent $57 billion on the home-based care of more than 3 million recipients. This is about half of Medicaid's total spending on long-term care and about two-thirds of all spending on formal home care. 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. The traditional Medicaid home care benefit is an in-kind benefit of formal home care from a Medicaid-approved agency. The amount of care an individual can receive free of charge is determined by a "care plan" created by her physician or nurse following a medical examination, though in the specific cases we analyze there does not appear to be a binding upper limit. Appendix Section B discusses evidence on this and provides additional information about Medicaid home care.

In recognition of the importance of informal care and other ways of dealing with chronic health problems, many state Medicaid programs have implemented reforms toward more flexible, cash-like benefits (Doty et al., 2010).[8] These programs typically allow recipients to spend their benefits on a wide range of personal care goods and services including assistive

[8] Early versions of the Affordable Care Act included a long-term care insurance program that would have paid cash benefits. This program, the CLASS (Community Living Assistance Services and Supports) Act, was eventually dropped due to concerns about its cost.

devices, home modifications, and, most important, informal care from family or friends. More flexible, cash-like benefits are increasingly common in other countries as well. 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).

An important milestone in the debate about more- versus less- flexible benefits, and an important source of evidence in our paper, is the Cash and Counseling experiments. These were large-scale experiments run by Medicaid programs in Arkansas, Florida, and New Jersey that began in 1998. Participants were drawn primarily from the population of Medicaid home care recipients and were randomized to either the traditional in-kind home care benefit or a near-cash benefit, each with 50 percent probability. Participants randomized to the near-cash benefit could revert to the standard in-kind benefit at any time; those randomized to the in-kind benefit could not switch to the near-cash benefit. Each recipient of the near-cash benefit received a budget for spending on care-related goods and services roughly equal to the cost of the care in her care plan. She also received "counseling" services to help manage her benefit. These services included help with planning how to spend the benefit, hiring and paying caregivers (and paying payroll taxes), and maintaining records. The aim was to make it as easy to receive care for the near-cash group as for the in-kind group. The restriction that the near-cash benefit had to be spent on care-related goods and services was unlikely to be binding for most recipients because of the broad definition of care-related goods and services, especially the inclusion of informal care.[9]

The main goal of the experiments was to test whether recipients could effectively manage their near-cash benefits and receive "enough" care. The results were almost uniformly positive. Members of the near-cash treatment group reported greater satisfaction with their care and with their lives as a whole (Foster et al., 2003; Brown et al., 2007). They also had similar or better health outcomes across a wide range of measures such as mortality, nursing home entry, falls, urinary tract infections, and respiratory infections (Lepidus Carlson et al., 2007). In the official final report on the experiments, Brown et al. (2007) conclude that the near-cash benefit had overwhelmingly positive effects on recipients.

[9]The vast majority of participants had been receiving enough informal care at baseline to more than exhaust their benefit. At follow up, 86 percent of recipients of the near-cash benefit used it to pay for informal care (Brown et al., 2007). Appendix Section B.2 contains more information about the Cash and Counseling experiments. Appendix Table E.1 reports summary statistics of Cash and Counseling participants and balance tests; these provide evidence of a valid randomization. Appendix Table E.2 compares Cash and Counseling participants both to the broader population of people eligible for home care benefits and to those who take up Medicaid home care. Our analysis uses data on the 2,470 participants age 65 or older.

4 Moral Hazard Cost of In-Kind Provision

In this section, we estimate the price sensitivity of demand for formal care. This, the first ingredient in our approach, is the key parameter for quantifying the moral hazard cost of in-kind provision. We use the Cash and Counseling experiments, which have two major advantages for this purpose.[10] First, the randomization solves an especially difficult simultaneity problem: Many factors that shift the supply of formal care are also likely to shift the demand for formal care by changing the opportunity cost of informal care.[11] Second, the variation in the price of formal care spans the full range most relevant for policy, from zero to the market price.

The experimental results suggest that in-kind provision of home care has a large moral hazard cost. Table 1 shows that being randomized to in-kind benefits doubles average formal care consumption from 7 to 14 hours per week. Figure 2 shows the distributions of formal care consumption for those randomized to the in-kind versus near-cash benefits. In-kind provision increases formal care consumption throughout the distribution, more than doubling both the fraction of people who consume formal care (from 24 to 55 percent) and the fraction who consume more than 20 hours per week (from 9 to 22 percent).

We estimate the price sensitivity of demand for formal care taking into account censoring at zero and imperfect compliance. We account for censoring by treating an individual's observed hours of care, q_i, as the outcome of a censored, latent demand for care, $q_i = \max\{0,\ q_i^*\}$. We account for imperfect compliance—some people assigned to the near-cash benefit reverted to the traditional in-kind benefit and some left Medicaid home care altogether—by using the randomized assignment as an instrument for the price each participant faced. Participants who receive the near-cash benefit or who leave Medicaid home care face the market price in their state. Participants who receive the in-kind benefit face a price

[10]Previous research on the Cash and Counseling experiments has focused on the distinction between *paid* and *unpaid* home care, where paid home care includes care from family and friends as well as from professionals, so long as the recipient pays for it (e.g., Brown et al., 2007). We focus on the distinction between formal care, provided by professionals, and informal care, provided by family and friends, regardless of whether the recipient pays the caregiver. This is the relevant distinction for comparing in-kind formal care benefits to more flexible benefits that can be spent on informal care.

[11]Consider using changes in minimum wage laws as instruments for the price of formal care. Many formal home care workers earn roughly the minimum wage, so changes in the minimum wage likely shift the supply of formal care. But at the same time, changes in the minimum wage also likely change the opportunity cost of informal care-giving by changing the wage or employment prospects of some potential informal care-givers. This likely shifts the demand for formal care since formal and informal care are closely-related goods.

of zero.[12] 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 near-cash treatment, and X_i includes indicators for sex, education level, race, self-rated health at baseline, living alone at baseline, five-year age bins, and state. The key parameter of interest is β, the effect on formal care consumption of an increase in its net-of-subsidy price. As a starting point, we assume that (ε_i, ν_i) are jointly normal and estimate this system using an instrumental variables Tobit specification.

The first stage relationship is economically and statistically large. Being randomized to the in-kind benefit decreases the average price of formal care by approximately \$7.70, with a first-stage F-statistic of over 1,100 (see Appendix Table E.3). The instrumental variables estimate of β is presented in Table 2. It implies that a one-dollar increase in the hourly price of formal care reduces consumption by 1.8 hours per week. This corresponds to an elasticity of -1.7 at the sample means. The conclusion that the demand for formal care is highly sensitive to its price holds in each of the three states and is robust to a wide range of alternative assumptions about the distribution of the error terms and benefit limits (see Appendix Tables E.4 and E.5). See Appendix Section C for details and a discussion of the generalizability of the results to other populations and policies of interest.

The estimates imply that in-kind provision has a large moral hazard cost. An individual consuming the average amount of formal care in the in-kind group would consume no formal care without the subsidy and values the care she does receive at just 28 percent of its cost.[13]

5 Targeting Benefit of In-Kind Provision

In this section, we provide evidence on the two ingredients that determine the targeting benefit of in-kind provision: the distribution of formal care consumption and the link between

[12] In principle, care plans or maximum benefit rules could limit the amount of formal care that those receiving the in-kind benefit could consume free of charge and thereby raise the shadow price of formal care above zero. In practice, a variety of evidence suggests that recipients of the traditional in-kind benefit were able to consume as much care as they wished free of charge. See Appendix Section B.4 for additional details and evidence.

[13] With $\beta = -1.8$ and no income effects, someone consuming 14 hours of care per week has an equivalent variation of formal care benefits of \$54 per week. Medicaid's cost of that care is \$192 per week.

formal care consumption and the marginal utility of income. We discuss the implications for the targeting benefit of in-kind home care and conclude with evidence on the targeting benefit of Medicaid home care.

5.1 Distribution of formal care consumption

We use data from the National Long Term Care Survey (NLTCS) to estimate the distribution of formal care consumption among the home care-eligible population. The NLTCS is a nationally representative survey of Americans 65 and older who are eligible for Medicare (see Appendix Table E.2 for summary statistics). We use the standard eligibility criterion for home care benefits: having at least two ADL limitations. A subset of this population with low enough income and assets is also eligible for Medicaid home care.

Figure 3 shows the distribution of formal care consumption in the home care-eligible population. Even within this group of people with severe chronic health problems, there is significant heterogeneity in formal care consumption.[14] Sixty-three percent do not consume any formal care. Among those who do there is a long right tail. For that group, the 95th percentile is around-the-clock care, almost 17 times the median among those consuming care. At the average hourly price, that volume of care would cost $131,000 per year.

The significant heterogeneity in formal care consumption implies that in-kind provision has a large targeting effect, sharply concentrating transfers on the small subset of the eligible population with high formal care consumption. The standard deviation of formal care consumption, σ_{x_k}, is 35 hours per week. At the average market price, that implies a standard deviation of annual spending—and so of the ex post marginal benefit of increasing the subsidy rate on formal care—of more than $27,000.

5.2 Link between formal care consumption and marginal utility

Both theoretical considerations and empirical evidence suggest a strong link between formal care consumption and marginal utility. In theory, formal care consumption will tend to be positively linked to marginal utility through the budget constraint: Greater spending on formal care leaves fewer resources available for non-care consumption.[15] Empirically, private long-term care insurance contracts typically subsidize formal care consumption, and people

[14]The cross-sectional distribution is not a pure measure of risk; it reflects predictable heterogeneity as well as heterogeneity in ex post realizations of risk. In the welfare analysis (Section 6), we test robustness to large changes in risk.

[15]The idea is that formal care consumption is a poor substitute for "regular," non-care consumption. This is the idea underlying the link between health spending and marginal utility in standard models of health spending risk. See Cutler and Zeckhauser (2000) for a review.

provide significant informal care and financial support to family members with high formal care consumption.

A strong link between formal care consumption and marginal utility implies that in-kind provision of home care would target relatively high-marginal utility states (high $Corr(\widehat{\lambda}, x_k)$). Together with the considerable heterogeneity in formal care consumption, a strong link also implies substantial risk within the benefit-eligible population (high $\sigma_{\widehat{\lambda}}$). Altogether, this suggests that the marginal targeting benefit of in-kind provision of home care would be large (large $\sigma_{x_k}\ \sigma_{\widehat{\lambda}}\ Corr(\widehat{\lambda}, x_k)$).

5.3 Targeting of Medicaid home care

Although the combination of highly-concentrated formal care consumption and a strong link between formal care consumption and marginal utility would imply a large targeting benefit of in-kind provision, the targeting of Medicaid home care depends not only on in-kind provision but also on factors such as awareness of the program and hassles.

Table 3 investigates the targeting of Medicaid home care within the eligible population using nationally representative data from the 1999 NLTCS. The first three rows of the table present estimates of the take-up rate among those eligible for benefits. Differences in the estimates are due to differences in the estimated size of the eligible population. The estimates range from 5 to 19 percent, with 19 percent likely overstating the true rate (see Appendix Section B.3 for details). Compared to an equal-cost program with complete take up, low take up of Medicaid home care increases benefits per recipient by a factor of 5–20. Combining the concentration of benefits from incomplete take up with that from differences in formal care consumption among those who take up implies a large targeting effect of Medicaid home care. The standard deviation of Medicaid-financed formal care is 27 hours per week.[16]

The next several rows of Table 3 compare the characteristics of those who do versus do not take up benefits among the eligible population, using the "Income eligible, $<$ 2 cars" eligibility criteria. People who take up have much greater demand for formal care: If everyone faced a common price, those who take up would be predicted to consume 12 hours per week more formal care on average.[17] Consistent with this, those who take up are sicker (66 vs. 46 percent have four or more ADL limitations) and have fewer "prime" potential informal caregivers (67 vs. 59 percent are unmarried and 39 vs. 29 percent live alone). The correlation

[16]This assumes that all of the formal care consumed by those who take up Medicaid home care is paid for entirely by Medicaid and that all Medicaid home care benefits are in-kind formal care, not cash (cash benefits were rare at the time). That the data lack information on the transfer from Medicaid increases the uncertainty in the calculation but does not obviously bias it toward greater or lesser concentration.

[17]We use our estimated price sensitivity from Section 4 to predict what each individual's consumption would have been had she faced a price of \$18.50 per hour, the maximum price in the data.

of benefits with formal care consumption is 0.62, with number of ADL limitations is 0.13, and with living alone is 0.19.

We turn to investigating targeting in the Cash and Counseling experiments. Unlike take up of Medicaid home care, the experimental design isolates the effect of in-kind provision. We focus on participants in Arkansas, the only state in which we can calculate each individual's near-cash benefit.

Figure 4 shows the distributions of transfers separately for those randomized to the in-kind and near-cash benefits.[18] The in-kind benefit concentrates transfers significantly relative to the near-cash benefit. Transfers to those assigned to the near-cash benefit cluster tightly around the median of \$147 per week. Transfers to those assigned to the in-kind benefit are much more dispersed, with a standard deviation more than twice as large and a much greater likelihood of being very large or very small.

Figure 5 shows the extent to which each benefit type concentrates transfers on people with the greatest demand for formal care. For each benefit type, we rank those randomly assigned to that benefit by their formal care consumption. Then we calculate the average transfer, in dollars, received by people at different ranks of the distribution. The in-kind transfers are highly concentrated on those with the greatest demand for formal care. Whereas the average in-kind transfer is \$133 per week, individuals between the 91st and 95th percentiles of the formal care distribution receive an average of \$350 per week and individuals above the 95th percentile receive an average of \$843 per week—almost 7 times the average benefit. The near-cash transfers, by contrast, are roughly constant throughout the formal care distribution, despite being based on individual medical exams. Appendix Section D provides suggestive evidence that in-kind provision concentrates benefits on recipients who are sicker and have fewer informal caregivers than the average recipient as well.

Taken as a whole, these results indicate that in-kind provision sharply concentrates transfers on a small fraction of the eligible population who are sicker, have fewer informal caregivers, and have a greater demand for formal care than the average eligible. To the extent that such recipients tend to have relatively high marginal utility, in-kind provision could have a large targeting benefit.

[18] The near-cash transfers are calculated as the product of care plan hours and the hourly price of care. The in-kind transfers are calculated as the product of hours of care used and the hourly price. For both groups, if the individual leaves Medicaid home care we set their transfer to zero. We censor transfers at \$600 for the figure but not elsewhere.

6 Welfare Effect of In-Kind Provision: Targeting Benefit Versus Moral Hazard Cost

This section uses a stylized expected utility model to quantify the net welfare effect of the targeting benefit and moral hazard cost of in-kind home care benefits. As discussed in Section 2, the key ingredients for the analysis are the price sensitivity of demand for formal care, the distribution of formal care consumption, and the link between formal care consumption and marginal utility. The first two are readily estimable; the third is provided by the model.

6.1 Model, policy counterfactual, and welfare measure

An individual faces risk about her health and her costs of coping with bad health. Together, these determine the level of her demand for formal care. The amount of formal care at which she reaches satiation (i.e., how much she would consume if facing a price of zero) is $\theta \in \mathbb{R}_+$. θ is known to be drawn from the distribution $G(\theta)$, but the particular realization of θ is not contractible ex post. Once θ is realized, the individual chooses formal care consumption, F, and non-care consumption, A ("all other goods," the numeraire), to maximize utility subject to a budget constraint that depends on the policy in operation. Indirect utility is

$$v(p, m; \theta) = \max_{A \geq 0, F \geq 0} u\left(A - \frac{(\theta - F)^2}{2\beta}\right) \text{ subject to } A + pF = m,$$

where p is the net-of-subsidy price of formal care and m is total after-transfer income, including any cash benefit from the home care program and any transfer from a means-tested program that provides a consumption floor. The corresponding Marshallian demand for formal care is

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

$\beta \geq 0$ determines the utility cost of consuming levels of care other than the satiation level θ and thereby determines the sensitivity of the demand for formal care to its price.

This utility function is motivated by key evidence from our setting. It produces a simple demand function for formal care that is consistent with some people in bad health not consuming any formal care, with formal care consumption being sensitive to its price, and with people becoming satiated at finite levels of formal care consumption.[19] It has an intuitive interpretation: Utility is decreasing in any unmet, residual care demand, $(\theta - F)$, the size

[19]The most direct evidence of satiation is that among the Cash and Counseling participants for whom we observe care plans, 43 percent consume less care than their care plans entitle them to. Intuitively, satiation might arise from a demand for privacy or space, since home care involves close contact with caregivers in one's home.

of which is increasing in the level of demand for formal care and decreasing in formal care consumption. This captures the idea that certain health problems are costly for people to cope with on their own. It nests as a special case the widely-used model in which health spending is equivalent to a wealth shock and shares with that model the implication that formal care consumption is linked to marginal utility mainly through the budget constraint: Greater spending on formal care means lower non-care consumption and so greater marginal utility.[20]

We analyze cost-neutral shifts in a mixed in-kind/cash-benefit policy that combines a linear subsidy rate s and a cash benefit b. For any policy, indexed by s, the cash benefit $b(s)$ adjusts to hold fixed total spending on recipients, which is the sum of spending on the subsidy, the cash benefit, and the consumption floor program. Take up of all benefits is automatic and there are no participation costs. This policy counterfactual isolates the effect of in-kind provision from other sources of incomplete take up, and it isolates the insurance-moral hazard tradeoff of in-kind provision from redistribution between recipients and other parties.

We measure the welfare effect of policy s as its ex ante equivalent variation gain over an equal-cost pure-cash policy, $EV(s)$. Expected (indirect) utility is $EU(s,b) = E\left(\max\{u(\bar{c}), v(p(s), m+b; \theta)\}\right)$, where $u(\bar{c})$ is utility when relying on the consumption floor. The equivalent variation gain of policy s is the extra income the individual would need over an equal-cost pure-cash policy to be as well off in expected utility as she is under s,

$$EU(0, b(0) + EV(s)) = EU(s, b(s)).$$

6.2 Empirical inputs and other parameter values

The key empirical inputs are the first two ingredients described in Section 2: the price sensitivity of demand for formal care and the distribution of formal care consumption. Our baseline value of the price sensitivity of demand is our main estimate from the Cash and Counseling experiment, $\beta = \left|\hat{\beta}_{C\&C}\right| = 1.8$.[21] This estimate implies that each \$1 increase in the hourly price of formal care reduces formal care consumption by 1.8 hours per week.

[20] As β decreases to zero, demand for formal care becomes less elastic, indirect utility approaches $u(m - p\theta)$, and spending on formal care becomes equivalent to a negative wealth shock—the standard case in the literatures on long-term care and health spending risks more generally. Compared to this standard case, our baseline model with $\beta > 0$ implies a weaker link between formal care consumption and marginal utility, which, other things equal, reduces the targeting benefit of in-kind provision. See Appendix Section E.1 for details.

[21] As described in Section 2, the moral hazard cost of in-kind provision depends on the total response of demand to the policy change. Absent income effects on demand for formal care, $\beta_{C\&C}$ is the correct parameter for evaluating any policy that affects the relative price of formal care. With non-zero income effects, $\beta_{C\&C}$ is the right parameter for analyzing policies like those in the Cash and Counseling experiments, which roughly

The high sensitivity of formal care demand to its price means that the moral hazard cost of in-kind provision will be large, especially at high subsidy rates.

Our baseline value of the distribution of formal care consumption is the observed distribution among non-institutionalized individuals age 65 and older who have two or more ADL limitations in the NLTCS. Restricting to people with two or more ADL limitations follows standard practice for Medicaid home care and private long-term care insurance contracts. We use β to convert the observed joint distribution of formal care consumption and formal care prices into the distribution of the level of demand for formal care, $G(\theta)$. Appendix Section E.2 contains details of this procedure.

Figure 6 presents our main estimate of the density of the level of demand for formal care, $g(\theta)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits substantial dispersion and has a long right tail. Most of the mass reflects low demand for care; about 56 percent of the θ satiation values are less than 10 hours per week. For those θs, an individual facing the average market price would consume no formal care. But some states have high demand. The 90th percentile, for example, is about 37 hours per week. The substantial heterogeneity in demand implies that in-kind provision will concentrate transfers significantly. Together with the model, it also implies substantial heterogeneity in non-care consumption and so in marginal utility. This suggests that the targeting benefit from in-kind provision could be large.

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; Ameriks et al., 2011). In our model, the argument c is "net consumption," non-care consumption net of any residual coping costs, $c = A - \frac{(\theta-F)^2}{2\beta}$. We follow Brown and Finkelstein (2008) and others in taking as a baseline value a coefficient of relative risk aversion, γ, of three. Income before transfers is \$15,000 per year. The distribution of before-subsidy prices of formal care is the empirical distribution observed in the NLTCS. If the individual cannot achieve net consumption of at least $\bar{c}$ =\$5,000 per year, she receives transfers that enable her to reach exactly that living standard. This consumption floor is meant to approximate the combined effects of means-tested government programs like Medicaid and Supplemental Security Income as well as any non-governmental charity care. The higher the consumption floor, the smaller the gains from insurance.[22] The policy

held fixed Medicaid's spending on each participant of the experiments, but not policies with different cash benefits. Cash and Counseling's near-cash benefits were on average greater than those under the policy counterfactual we consider here, which holds fixed total spending on the entire eligible population. With positive income effects on demand for formal care, this estimate will tend to understate slightly the true moral hazard effect of in-kind provision in these policy counterfactuals.

[22]In many contexts, a sizable fraction of insurance transfers displace means-tested transfers rather than increasing consumption. As a result, greater insurance (a higher subsidy rate in our context) is implicitly

counterfactuals hold fixed total spending on recipients at its expected level under a pure in-kind benefit ($s = 1, b = 0$), which is \$7,150.

6.3 Welfare effects of in-kind provision

As a benchmark and to get a sense of the extent of the risk within benefit-eligible states, we first calculate the welfare gain from a hypothetical (infeasible) first-best contract that provides state-dependent cash transfers. The equivalent variation gain from this contract is \$10,687 (see Appendix Section E.1 for details).

Figure 7 shows the marginal benefit and marginal cost of a cost-neutral reduction in the price of formal care as a function of the subsidy rate, the quantitative analogue of Figure 1. The marginal benefit far exceeds the marginal cost for subsidy rates up to about 75 percent before falling somewhat below the marginal cost for subsidy rates above 90 percent.[23]

The first column of Table 4 reports statistics about the optimal policy. The optimal subsidy rate is 87 percent, close to that of a pure in-kind program. The optimal subsidy increases welfare substantially. Its equivalent variation gain over the equal-cost cash contract is \$6,416. For a cash-benefit contract to achieve the same expected utility, it would have to cost 90 percent more than the optimal contract. Though not optimal, a pure in-kind benefit program with a 100 percent subsidy and no cash benefit also improves substantially on the pure-cash program, with an equivalent variation gain of \$5,265.

6.4 Robustness and intuition

To assess the robustness of the main conclusions and the relative importance of different factors in driving them, we summarize the effects of changes in each of the three key ingredients of the analysis.

Price sensitivity of demand for formal care. Our baseline estimate of the price sensitivity of demand implies a large moral hazard cost of in-kind provision. As a result, the optimal contract achieves only 60 percent of the gain from the first-best policy. This shortfall comes from costs along two dimensions. First, recipients over-consume formal care and as a result value the benefit less than its cost in each state ex post. The optimal benefit nearly triples

taxed by the means-tested program. In the context of long-term care, this implicit taxation of private insurance by means-tested programs is quite large (Brown and Finkelstein, 2008). In our baseline specification, the individual relies on the floor in states of the world in which her spending on formal care would otherwise be at least \$10,000, the amount by which income exceeds the consumption floor.

[23]The marginal benefit is steeper than the mechanical effect at small subsidy rates because increasing the subsidy increases the set of states in which the individual consumes any formal care. This tends to increase the covariance between marginal utility and formal care consumption. At higher subsidy rates, this effect is dominated by the effect of increasing the subsidy on reducing the variance in marginal utility.

formal care consumption and its mean ex post value is just 49 percent of its cost (column 1 of Table 4). Second, the optimal contract leaves some risk uninsured, since the benefit of insuring it would be more than offset by the moral hazard cost.

Although it is not possible to sign the difference between our estimate of the price sensitivity and the long-run price sensitivity to permanent policies (see Appendix Section C.3), our finding of a large gain from in-kind provision in the baseline specification makes tests of the robustness of the results to even greater price sensitivity of the most interest. Column 2 of Table 4 shows results based on $\beta = 5$, about three times greater than our main estimate. Although the optimal subsidy and the gain from in-kind provision are smaller than in the baseline specification, they remain large. The optimal subsidy remains large even when $\beta = 10$, over five times greater than our main estimate and greater than seems plausible even in the long run (see Appendix Table E.8).

Distribution of formal care consumption. The observed cross-sectional distribution of formal care consumption exhibits substantial heterogeneity with a long right tail. This is the key empirical fact driving the large targeting benefit of in-kind provision. Using the observed distribution to proxy for the (unobservable) counterfactual distribution facing an individual follows much of the literatures on optimal taxation and health care but is limited by the fact that it treats all of the observed heterogeneity as reflecting the result of an exogenous, uncertain process. In reality, some of the observed heterogeneity reflects measurement error, some is predictable, and some is endogenous—people make many choices that affect their future demand for formal care, including investing in their health and relationships with potential caregivers. For these reasons and others, the observed distribution of formal care consumption is an imperfect measure of the risk facing any particular individual.

Columns 3 and 4 of Table 4 show results based on specifications with less variation in the demand for formal care than is implied by the observed distribution of formal care consumption. Column 3 cuts off the right tail of the θ distribution, dropping states with $\theta > 50$. Column 4 scales down the θ distribution, replacing each θ with $\theta/2$ (and thereby reducing the variance of the distribution to one-fourth of its baseline value). Both changes reduce the optimal subsidy and the gain from in-kind provision, but in both cases the optimal subsidy and the gain from in-kind provision remain large.

Link between formal care consumption and marginal utility. Although our baseline model involves a weaker link between formal care consumption and marginal utility than the benchmark model in which health spending is treated as a wealth shock, in-kind provision is quite effective at targeting states of the world with relatively high marginal utility. Absent insurance (i.e., under a pure-cash policy), the correlation between an individual's marginal utility and formal care consumption is 0.89 (column 1 of Table 4). The strength of this link depends on two sets of factors. One is factors that affect the extent to which greater spending on formal care reduces non-care consumption, including saving and dissaving and informal insurance arrangements. Such factors are most naturally modeled as reducing the dispersion in the distribution of demand for formal care, the effects of which we just discussed.

The other set of factors is features of the utility function, in particular its curvature and any state dependence. Column 5 reduces the coefficient of relative risk aversion to one (log utility). This significantly reduces the extent to which heterogeneity in net consumption translates into heterogeneity in marginal utility and so significantly reduces the value of insurance. Both the optimal subsidy and the gain from in-kind provision are much smaller: 35 percent and \$133, respectively. This largely reflects the low cost of home care risk in this specification. The gain from the first-best contract is just \$1,215 in this case, 11 percent of its value in the baseline specification. Column 6 considers state-dependence in utility. It reduces relative marginal utility in higher-demand states to an extent designed to match the upper end of the most relevant estimates of state dependence, those of Finkelstein et al. (2013) (see Appendix Section E.3). This reduces both the extent of risk and how well in-kind provision targets transfers. The worse targeting is revealed by the reduction in the correlation between marginal utility and formal care consumption from 0.89 to 0.54 (columns 1 and 6 of Table 4). As a result, both the optimal subsidy and the welfare gain from in-kind provision are smaller than in the baseline specification. But they remain large in absolute terms: 60 percent and \$1,505, respectively.

With one exception, the gain from in-kind home care is large and robust to changes in the key inputs that appear to span the range of plausible values. Appendix Section E.4 shows that the results are robust to several other changes as well. The one exception is risk aversion. If risk aversion is low enough, the cost of uninsured risk becomes small enough that even first-best insurance is not that valuable. This points to the key role of the curvature of the utility function in determining the cost of the risk. Provided utility is such that home care risk is important, in-kind provision appears to produce a large welfare gain.

7 Conclusion

We develop an approach to quantifying a central tradeoff of in-kind provision—it can improve targeting at the cost of being less valuable to recipients ex post—and apply it to home care. Despite the ubiquity of in-kind transfers and the centrality of this tradeoff to their welfare effects, little is known about the magnitude of these costs and benefits in many important contexts. We find that in the context of home care, the targeting benefit of in-kind provision appears to exceed its large moral hazard cost. This conclusion is fundamentally driven by the substantial risk within benefit-eligible states of the world.

In focusing on the targeting-distortion tradeoff in the stand-alone home care context, we have not explicitly modeled substitution with informal care or nursing homes and we have omitted administrative and takeup costs. Available evidence is suggestive that these considerations would tend to reinforce the net advantage of in-kind provision. In-kind provision likely reduces informal care, which likely increases the labor supply and net tax payments of would-be informal caregivers (Ettner, 1995; Van Houtven et al., 2013; Skira, 2015)—a positive fiscal externality. In-kind provision might also benefit would-be informal caregivers by reducing and insuring their share of the costs of the recipient's bad health. In-kind provision is unlikely to have much effect on usage of nursing homes, given the limited substitution between home care and nursing homes (see Grabowski, 2006, for a review). Although cash benefits typically involve lower administrative and take-up costs than in-kind benefits, in home care many of the cash-like benefits that have been implemented in practice involve medical exams, counseling, monitoring, and similar features that likely make them exceptions to this general rule. Whether in-kind provision's targeting benefit could be achieved in less-costly ways is an important question for future research.

Several recent policy reforms and proposals make restrictive in-kind benefits more flexible and cash-like. A major impetus for these reforms is the view that recipients would much prefer equal-cost cash transfers, a view that is consistent with our findings about Medicaid home care. But such reforms also change the distribution of benefits within the eligible population. If a more flexible benefit worsens targeting, the targeting loss should be weighed against the gain from making the benefit more valuable to recipients ex post.

Optimal benefit design is a central policy issue, as many major programs involve in-kind transfers of schooling, housing, food, health care, and other goods. Although home care shares much in common with other important contexts, especially other types of health care, the desirability of in-kind provision is necessarily context-specific. Evaluating the costs and benefits of alternative benefit designs is critically important, and our approach to quantifying this tradeoff could prove fruitful in other contexts as well.

References

Alatas, 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 Economy 124*(2), 371–427.

Alatas, 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.

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.

Barczyk, D. and M. Kredler (2017). Evaluating long-term care policy options, taking the family seriously. *Review of Economic Studies 85*(2), 766–809.

Basurto, P. M., P. Dupas, and J. Robinson (2017). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi. *NBER Working Paper 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 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 (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 (2017, June). National health expenditure accounts. Technical report.

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.

Cunha, J. M., G. De Giorgi, and S. Jayachandran (2011). The price effects of cash versus in-kind transfers. *NBER Working Paper 17456*.

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.

Cutler, D. M. and R. J. Zeckhauser (2000). The anatomy of health insurance. In *Handbook of Health Economics*, Volume 1, pp. 563–643. Elsevier.

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.

Deshpande, M. and Y. Li (2017). Who is screened out? Application costs and the targeting of disability programs. *NBER Working Paper 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.

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., N. Hendren, and E. F. Luttmer (2015). The value of Medicaid: Interpreting results from the Oregon health insurance experiment. Technical report, National Bureau of Economic Research, Inc.

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.

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*.

Genworth Financial (2005). Costs of care survey. Technical report.

Grabowski, D. C. (2006). The cost-effectiveness of noninstitutional long-term care services: Review and synthesis of the most recent evidence. *Medical Care Research and Review 63*(1), 3–28.

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.

Hoynes, H. and D. Whitmore Schanzenbach (2016). U.S. food and nutrition programs. In *Economics of Means-Tested Transfer Programs in the United States*, Volume I, pp. 219–302. University of Chicago Press.

Jacoby, H. G. (1997). Self-selection and the redistributive impact of in-kind transfers: An econometric analysis. *Journal of Human Resources 32*(2), 233–249.

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.

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.

Manning, W. G. and M. S. Marquis (1996). Health insurance: the tradeoff between risk pooling and moral hazard. *Journal of health economics 15*(5), 609–639.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies 38*(2), 175–208.

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

Musgrave, R. A. (1959). *The Theory of Public Finance*. New York and London: McGraw-Hill.

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

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.

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

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies 68*(1), 205–229.

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

Skira, M. M. (2015). Dynamic wage and employment effects of elder parent care. *International Economic Review 56*(1), 63–93.

Smeeding, T. M. (1977). The antipoverty effectiveness of in-kind transfers. *The Journal of Human Resources 12*(3), 360–378.

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 ...

Van Houtven, C. H., N. B. Coe, and M. M. Skira (2013). The effect of informal care on work and wages. *Journal of Health Economics 32*, 240–252.

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*.

Zeckhauser, R. J. (1970). Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives. *Journal of Economic Theory 2*(1), 10–26.

Tables and Figures

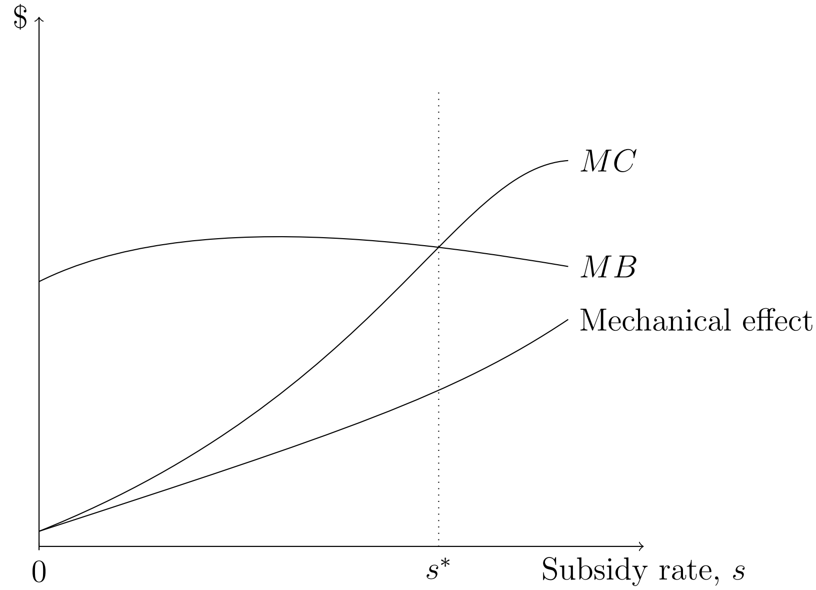

Figure 1: Costs and Benefits of a Subsidy

[MB is the marginal benefit of an increase in the subsidy on good k. MC is the marginal cost. The "mechanical effect" ($E(x_k)$) is both (i) what the marginal benefit would be if the targeting benefit were 0 (i.e., if $Cov\left(\widehat{\lambda}, x_k\right) = 0$) and (ii) what the marginal cost would be if the moral hazard cost were 0 (i.e., if $E\left(-dx_k/dp_k\right) = 0$).]

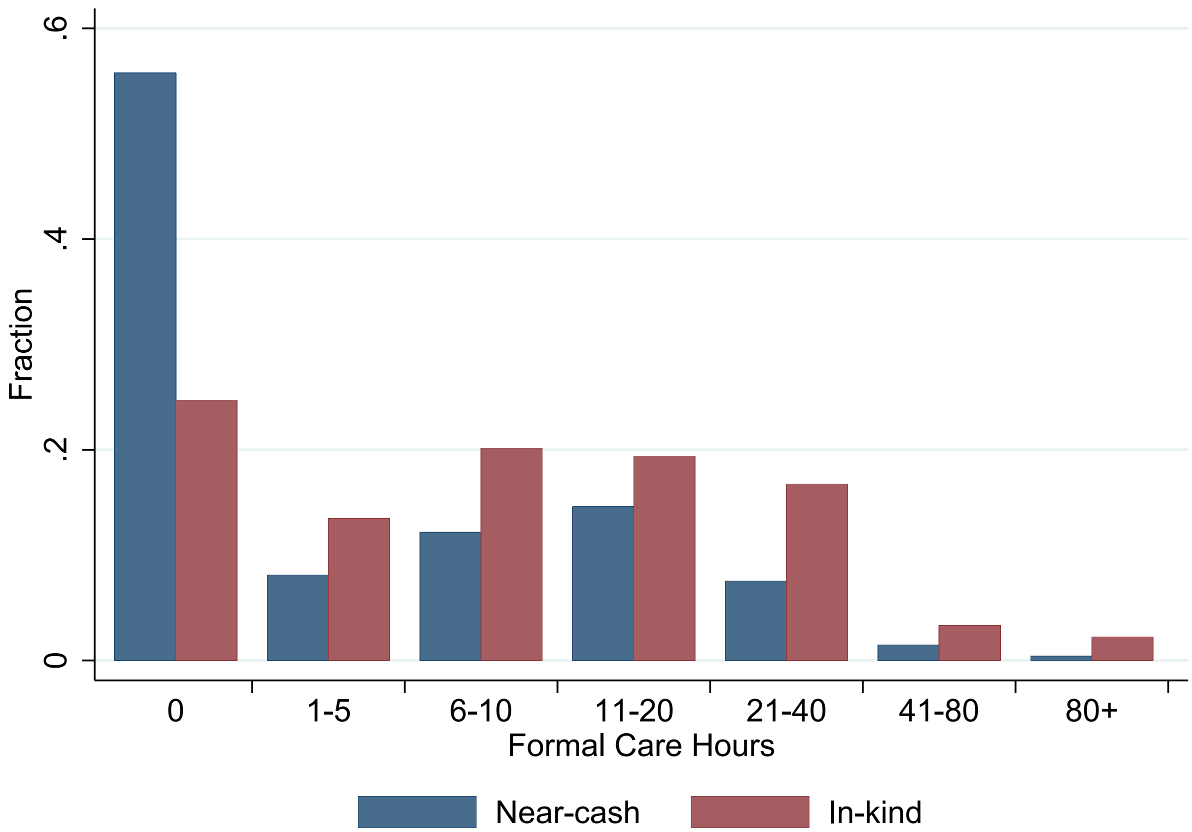

Figure 2: PDFs of Formal Care Consumption by Randomized Benefit Assignment

[Formal care consumption, in hours per week, among participants randomly assigned to the in-kind vs. near-cash benefit. Data from Cash and Counseling follow-up survey.]

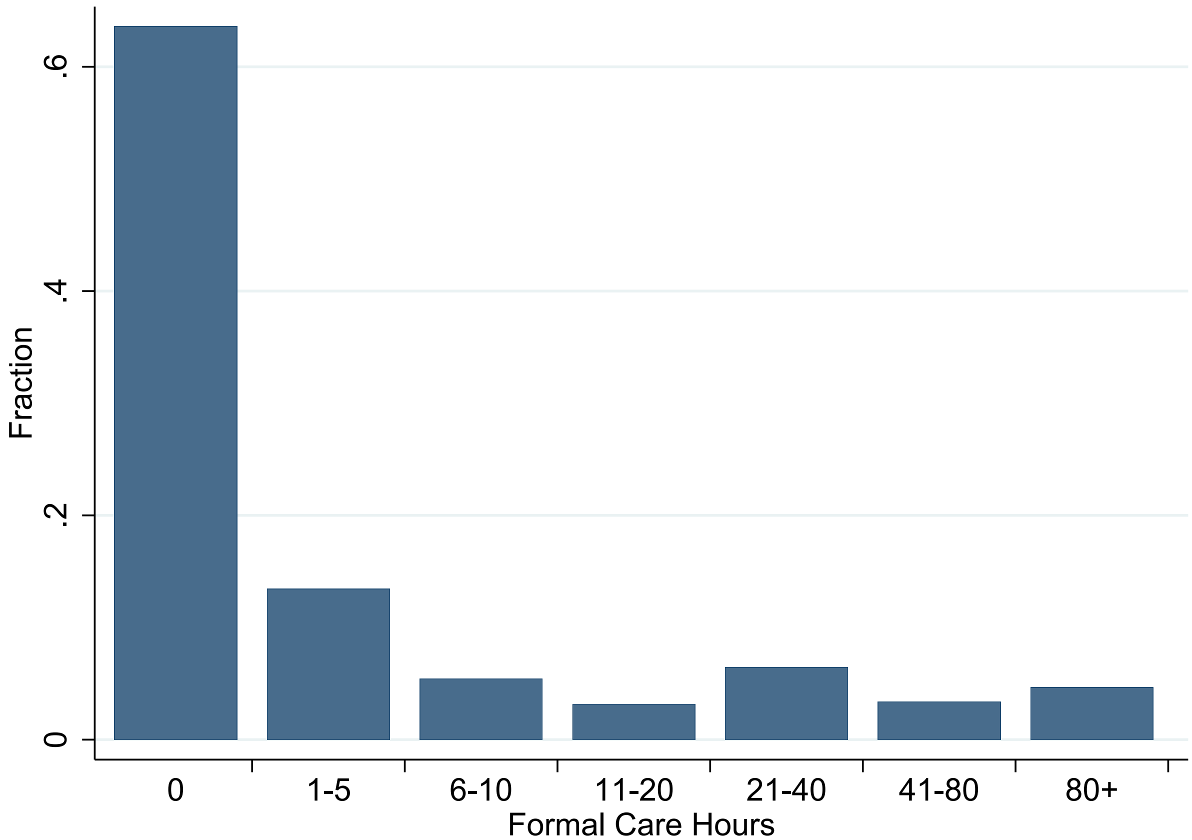

Figure 3: Distribution of Formal Care Consumption in the Benefit-Eligibile Population

[Hours per week of formal care consumption among the non-institutionalized population aged 65 and older with two or more ADL limitations. Data from the 1999 National Long-Term Care Survey. Sixty-three percent do not consume any formal care. Conditional on consuming formal care, median consumption is 10 hours per week, the 75th percentile is 40 hours per week, the 90th percentile is 120 hours per week, and the 95th and 99th percentiles are 168 hours per week (around-the-clock care). The standard deviation, σ_{x_k}, is 35 hours per week.]

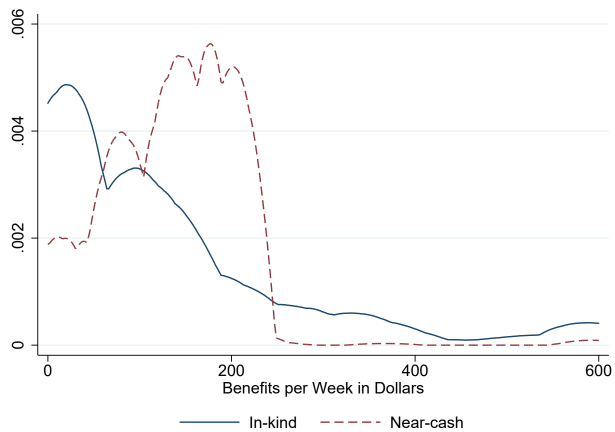

Figure 4: Targeting Effects of In-Kind Provision on the Intensive Margin

[Distributions of transfers in the Arkansas Cash and Counseling experiment. Arkansas is the only state for which we observe care plan hours, which we need in order to estimate the near-cash transfer. Transfers are measured in dollar-costs per week at market prices. We scale up the near-cash group's transfers to have the same mean as the in-kind group's in order to isolate differences in the concentration of transfers, not their average size. The average transfer is $133. Groups are based on each individual's randomized assignment. Transfers have been censored at $600 for the figure.]

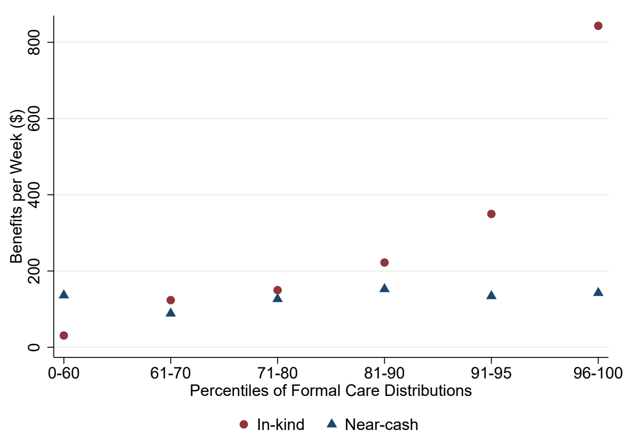

Figure 5: Targeting of In-Kind Versus Near-Cash Benefits

[Average transfers, in dollars per week, in the Arkansas Cash and Counseling experiment, separately for those randomized to the in-kind and near-cash benefit. For each benefit type, we rank those randomly assigned to that benefit by their formal care consumption. Then we calculate the average transfer, in dollars, received by people at different ranks of the distribution. Fifty-seven percent of those randomized to near-cash do not consume any formal care.]

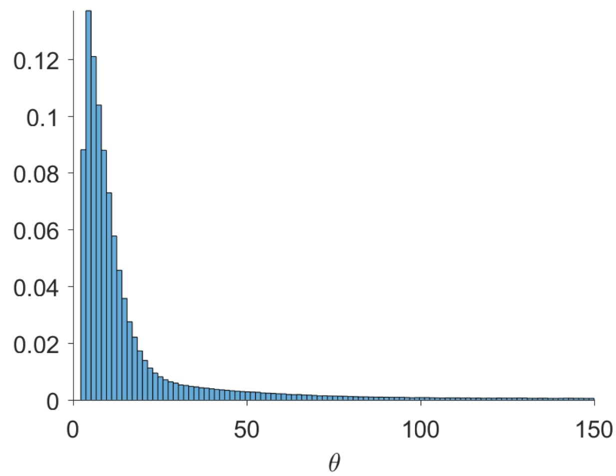

Figure 6: Distribution of Demand for Formal Care

[Simulated distribution of formal care satiation points, θ, in hours per week, among the non-institutionalized population aged 65 and older with two or more ADL limitations. The mean is 16.3 hours per week.]

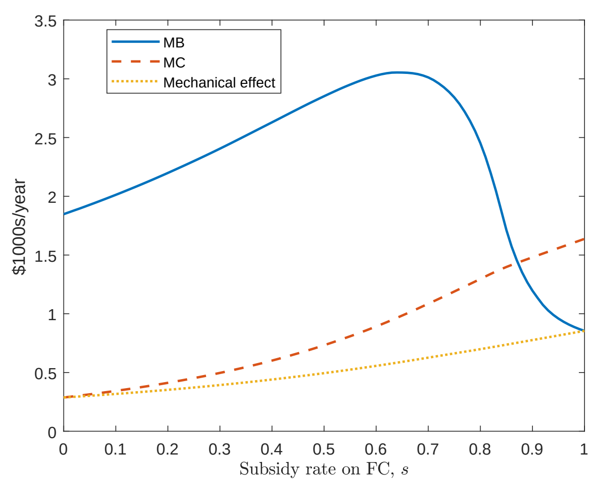

Figure 7: Marginal Benefit & Marginal Cost of Decrease in Price of Formal Care

[Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total spending on recipients. $s = 1$ corresponds to a pure in-kind benefit program, a 100 percent subsidy on formal care with no cash benefit. $s = 0$ corresponds to a pure cash benefit program, a 0 percent subsidy on formal care.]

Table 1: Average Formal Care Consumption by Treatment Group

	Near-cash	In-kind	Difference p-value
Overall	6.85	14.19	<0.01
Arkansas	6.29	10.76	<0.01
Florida	7.69	18.60	<0.01
New Jersey	7.01	16.10	<0.01

Means of formal care consumption in hours per week. "Near-cash" and "In-kind" groups are defined by randomized treatment assignment. P-values test for equality of means. Rows denote different samples.

Table 2: The Price Sensitivity of Demand for Formal Care

	(1)	(2)
Price	-1.78	-1.76
	(0.15)	(0.15)
Controls	No	Yes
First-stage F-statistic	1,139	1,144
Mean hours, in-kind	14.19	14.19
Observations	2,440	2,440

Dependent variable is formal care consumption in hours per week. Specifications are instrumental variables Tobits where formal care hours are censored at zero. Controls included in column (2) are indicators for sex, education level, race, self-rated health, five-year age bins, and state. Data are from the Cash and Counseling experiments. Robust standard errors reported.

Table 3: Targeting of Medicaid Home Care

	(1) Take-up = 0	(2) Take-up = 1	(3) Difference p-value
Fraction of eligibles who do vs. do not take up, under different definitions of eligibility			
Income eligible, < 2 cars	0.95	0.05	
Income eligible, no cars	0.90	0.10	
Restrictive income, no cars	0.81	0.19	
Summary Statistics			
Level of formal care demand	8.30	20.82	<0.01
Age	80.01	80.82	0.45
Four or more ADLs	0.46	0.66	<0.01
Health fair or poor	0.69	0.78	0.12
Female	0.70	0.72	0.66
Lives alone	0.29	0.39	0.12
Unmarried	0.59	0.67	0.19
Has children	0.75	0.78	0.73
Household income, monthly	847.95	675.56	0.01

Means for people who did (column 2) versus did not (column 1) take up Medicaid home care. "Difference p-value" tests the equality of means across groups. Take-up rates based on non-institutionalized individuals aged 65 and older with two or more ADL limitations who meet different sets of financial-related eligibility criteria. "Income eligible" is based on the income thresholds each state uses to determine eligibility. "Restrictive income" applies the most stringent (lowest) income limit to all states to try to estimate an upper bound on take up. Number of cars is an important determinant of eligibility for Medicaid home care. Summary statistics by take-up decision are for those who meet the "Income eligible, <2 cars" criteria. This sample has 448 individuals. The level of formal care demand, in hours per week, uses our estimate of price sensitivity to simulate each individual's hours of formal care if she faced a price of $18.50 per hour, the maximum in the data. The alternative to health fair or poor is health good or excellent. Data from the 1999 NLTCS.

Table 4: Welfare Analysis and Key Robustness Tests

	(1) Baseline	(2) $\beta = 5$	(3) Drop $\theta > 50$	(4) $\theta/2$	(5) Log utility	(6) State-dependent utility
Optimal subsidy, s^*	0.87	0.83	0.57	0.74	0.35	0.60
EV gain over pure-cash policy	\$6,416	\$5,086	\$1,683	\$2,554	\$133	\$1,505
E(ex post value)/E(cost)	0.49	0.59	0.86	0.74	0.80	0.68
Corr(marg. util., formal care)	0.89	0.86	0.72	0.88	0.92	0.54

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 5. Column 3 truncates the right tail of the risk by dropping all values of θ greater than 50. Column 4 divides every θ by 2, reducing the variance of the risk to one-fourth its baseline value. Column 5 sets the coefficient of relative risk aversion to one (log utility), whereas the baseline coefficient of relative risk aversion is three. Column 6 reduces relative marginal utility in higher-demand states to an extent designed to match the upper end of the most relevant estimates of state dependence in utility, those of Finkelstein et al. (2013).

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 of the world, but may or may not take up the in-kind benefit. Any in-kind benefit the

[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.

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 targeting benefit of an increase in the benefit limit is increasing in the covariance between

[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.

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 benefits on personal care services. The idea was that the remaining 10 percent could be

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

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.

Since selection could significantly bias such levels comparisons, we pursue a differences-

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

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 highly robust to plausible changes in the model. Column 2 shows that even a price sensitivity

[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.

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.

Appendix Figures and Tables

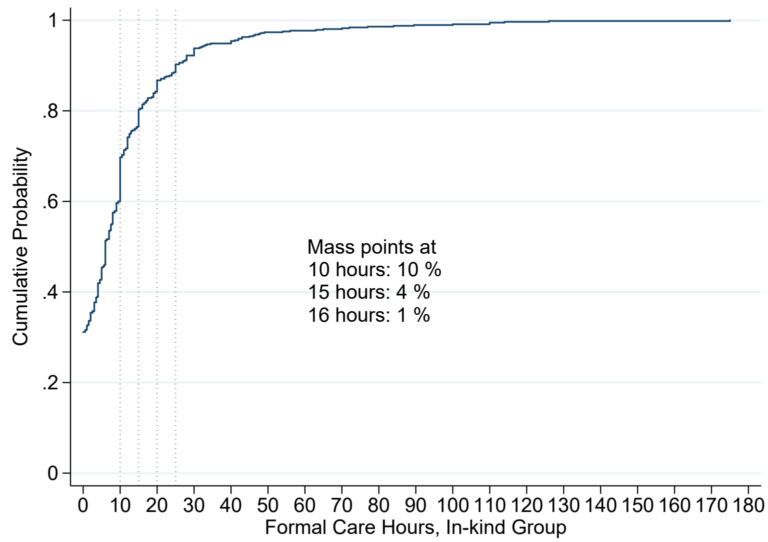

Figure E.1: CDF of Formal Care Consumption in Cash and Counseling, Arkansas

[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.]

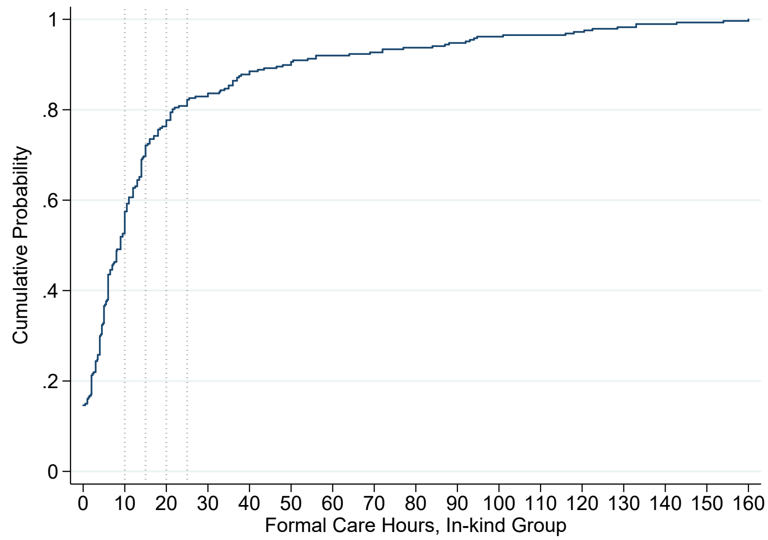

Figure E.2: CDF of Formal Care Consumption in Cash and Counseling, Florida

[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.]

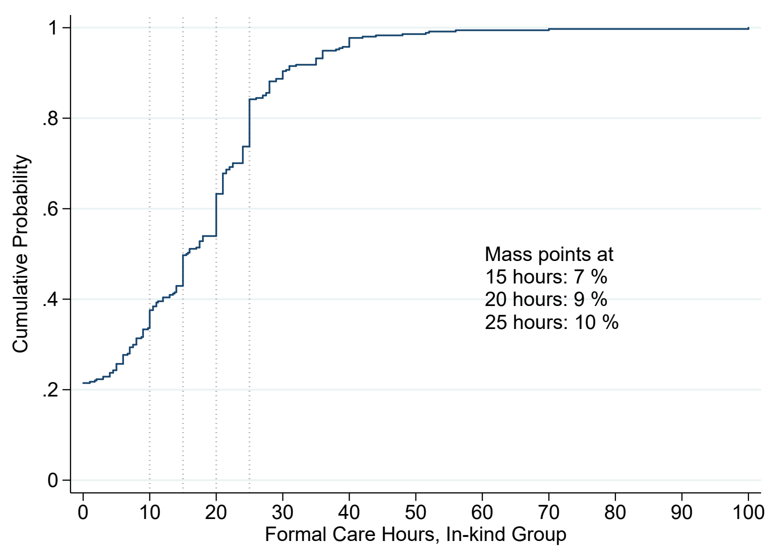

Figure E.3: CDF of Formal Care Consumption in Cash and Counseling, New Jersey

[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.]

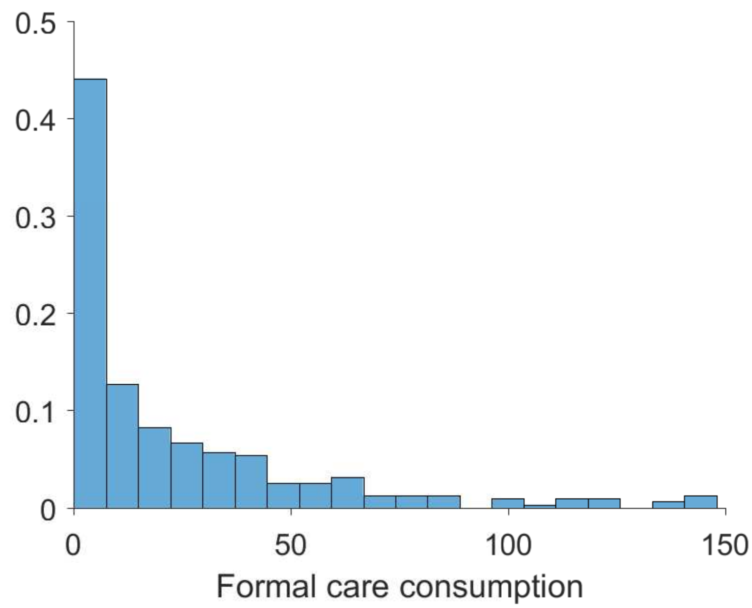

Figure E.4: Distribution of Formal Care Consumption in Benefit-Eligible Population

[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)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	90th Quantile	95th Quantile	99th Quantile	OLS	90th Quantile	95th Quantile	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.