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

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

October 12, 2017

Abstract

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

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

1 Introduction

In-kind transfers are a ubiquitous feature of government programs, private contracts, and charitable giving. In the U.S., government spending on in-kind health and education programs alone totals more than 12 percent of GDP (Currie and Gahvari, 2008). In-kind transfers are an important component of employment contracts, predominate in private health insurance contracts, and are thought to account for over 90 percent of charitable giving (Overseas Development Institute, 2015). In-kind transfers are also at the heart of a crucial debate about the relative desirability of benefit programs that are more universal and flexible versus more targeted and restrictive.[1]

Central to this debate is a tradeoff inherent to in-kind transfers. In-kind provision has a fundamental cost: Recipients would prefer cost-equivalent cash transfers. But this cost is linked to an important potential benefit: When information or other constraints preclude direct targeting, in-kind provision can better target desired recipients by leading certain people to take up more benefits than others (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988). In the context of insurance, for example, if someone's valuation of a particular in-kind benefit is higher in states of the world in which marginal utility is higher, in-kind provision can help concentrate benefits in those states and thereby better insure the risk. In such cases, there is a tradeoff between providing benefits that are more valuable to recipients (for which less restrictive cash-like benefits are best) and providing benefits that better target transfers to higher-marginal utility states (for which more restrictive in-kind benefits might be best). Although these costs and benefits are crucial determinants of optimal benefit design, little is known about their relative magnitudes across a wide range of important contexts.

In this paper, we develop a general framework for analyzing this key tradeoff of in-kind provision, and we apply it to the context of home care. Home care helps people with chronic health problems live at home instead of in nursing homes. It includes assistance with eating, dressing, and bathing, and it is provided by both professional caregivers ("formal care") and family and friends ("informal care"). Home care is an especially important and fruitful context in which to analyze the consequences of in-kind provision for three main reasons. First, it is one of the largest and fastest-growing components of what is likely the largest and fastest-growing type of in-kind transfer: in-kind health care. In the U.S. in 2015, spending on formal home care was $88 billion, and government spending on in-kind health-benefit

[1] In domestic policy, foreign aid, and private charitable giving, for example, there are active debates about the desirability of more flexible benefits (e.g., direct cash transfers and universal basic income programs) versus more restrictive in-kind transfers of food, housing (Collinson et al., 2015), medical care (Doty et al., 2010), and other goods.

programs totaled \$1.3 trillion, more than eight percent of GDP (Centers for Medicare and Medicaid Services, 2017). The recent Affordable Care Act increased such in-kind health benefits substantially through expanded Medicaid eligibility and subsidies for health insurance. Second, many states in the U.S. and countries in Europe have reformed their home care programs to make benefits more flexible and cash-like (National Conference of State Legislatures, 2007; Da Roit and Le Bihan, 2010). Fifteen state Medicaid programs allow recipients to use benefits to pay informal caregivers or buy equipment for their homes (Doty et al., 2010), and early versions of the bill that became the Affordable Care Act included a long-term care insurance program that would have paid cash benefits.[2] Third, the Cash and Counseling experiments—large-scale experiments in which participants were randomized to either Medicaid's traditional in-kind home care benefit or a near-cash benefit—enable us to credibly identify the main inputs of our framework.

The theory, which is based on the simple insight that an in-kind transfer is equivalent to a (potentially non-linear) price subsidy, highlights three key determinants of the welfare consequences of in-kind provision. One is heterogeneity in demand for the good within benefit-eligible states of the world. The greater is this heterogeneity, the greater the targeting effects of in-kind provision. We find that the demand for formal care is highly heterogeneous within benefit-eligible states (having two or more activities of daily living limitations). While 62 percent of people in benefit-eligible states do not consume any formal care, among those who do there is a long right tail. The 95th percentile is 168 hours per week, which at the average hourly price of \$15 per hour (Genworth Financial, 2005) amounts to about \$131,000 per year. Moreover, we find that the demand for formal care is highly heterogeneous even conditional on an extensive set of personal, household, and family characteristics, including the results of a detailed medical exam. This suggests that even extensively-"tagged" cash benefits (Akerlof, 1978) would leave significant risk uninsured.

A second key determinant of the welfare consequences of in-kind provision is its moral hazard effect, the extent to which it 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 exogenous variation in home care benefits from the Cash and Counseling experiments, we find that in-kind provision increases formal care consumption substantially. Our estimates imply that among people consuming any formal care, in-kind provision increases formal care consumption by 25 hours per week, twice average consumption among the benefit-eligible population. Our estimates also imply that many recipients value their in-kind benefit far below its cost. A recipient of the average in-kind transfer in the Cash and Counseling experiments, for example, would value it at just 28 percent of its cost.

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

A third key determinant of the welfare consequences of in-kind provision is the covariance between benefits and the marginal utility of income. If in-kind provision differentially reduces benefit take-up in relatively low-marginal utility states of the world, it helps insure the risk. On the extensive margin program take-up decision, we estimate that only 4 to 16 percent of those eligible for Medicaid home care take up benefits. Compared to the average eligible individual, those who take up have much greater demand for formal care, are sicker, and have fewer potential informal caregivers. Using the Cash and Counseling experiments, we find that in-kind provision concentrates benefits substantially on the intensive margin as well. The variance in benefits is seven times greater for those randomized to the in-kind benefit. Among the top five percent of formal care users in each of the randomized groups, the average benefit is four times greater in the in-kind group. Together these results indicate that in-kind provision sharply concentrates benefits on a small fraction of the eligible population that has a greater demand for formal home care, is sicker, and has worse informal care options than the average eligible. To the extent that such states of the world 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 much less valuable to recipients, but flexible cash-like benefits leave most of the risk uninsured. This raises the question: Does the targeting benefit of in-kind provision exceed the moral hazard cost? We combine our reduced-form estimates with a structural model to quantify these costs and benefits in a stylized expected utility framework. We find that under a wide range of assumptions the optimal contract involves a large in-kind component and delivers substantial welfare gains over cash-benefit contracts, despite the large moral hazard cost.

A large literature analyzes several known barriers to private, voluntary long-term care insurance, with two of the most important being adverse selection (Finkelstein and McGarry, 2006; Hendren, 2013) and the implicit taxation of private insurance by Medicaid (Pauly, 1990; Brown and Finkelstein, 2008) (see Brown and Finkelstein, 2011, for a review). We complement and extend this literature by estimating the importance of two barriers to any long-term care insurance, whether private or government, voluntary or mandatory: hard-to-verify heterogeneity and moral hazard. Our findings reveal a fundamental dilemma for benefit design. The large moral hazard cost of in-kind provision means that many recipients would be significantly better off *ex post* with a cost-equivalent cash transfer.[3] But that even richly-tagged cash benefits leave most of the risk uninsured means that providing home care in kind might be an unfortunate but unavoidable cost of insuring the risk from chronic health

[3]It also means that a large "moral hazard tax" plagues most long-term care insurance contracts and raises the effective loads to consumers above existing estimates (e.g., Brown and Finkelstein, 2007; Friedberg et al., 2014).

problems. Especially when combined with the other potential advantages of providing home care in kind, our findings raise concerns about the many recent reforms that make long-term care benefits more flexible and cash-like.[4]

Our approach links the theoretical and empirical literatures on in-kind transfers, which have been largely disconnected so far (see Currie and Gahvari, 2008, for a review). The theoretical branch investigates a variety of potential advantages of in-kind transfers, including improving targeting efficiency (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988), increasing the efficiency of the tax system (Munro, 1992), and reducing moral hazard in the context of the Samaritan's Dilemma (Bruce and Waldman, 1991). The empirical branch has mostly focused on estimating in-kind provision's effects on consumption.[5] Noting the equivalence of an in-kind transfer and a subsidy allows us to utilize the well-developed theoretical and quantitative approaches for analyzing subsidies from the vast literature on optimal taxation. Our analysis of home care sheds new light on the costs and benefits of in-kind provision in an important instance of the largest class of in-kind benefits, in-kind health care benefits.

Our work also contributes to the literature that studies the targeting of benefit programs with incomplete take up, including disability insurance (Low and Pistaferri, 2015; Deshpande and Li, 2017), Medicaid (Cutler and Gruber, 1996), housing assistance (Reeder, 1985), and Supplemental Security Income (Benitez-Silva et al., 2004) (see Currie, 2006, for a review). A related literature investigates the targeting effects of fees (Greaney et al., 2016), ordeals (Atalas et al., 2016), subsidized prices (Cohen and Dupas, 2010), and delegating authority over the distribution of benefits to local leaders (Atalas et al., 2012; Basurto et al., 2017). A key finding in these literatures is that in many benefit programs, only a small fraction of the eligible population takes up benefits. While low take-up can be undesirable in some contexts, our analysis suggests that low take-up of home care benefits improves risk sharing. Our work complements and extends these literatures by providing a simple framework for analyzing a key tradeoff of program features that lead to incomplete take up: they may improve targeting at the expense of reducing the value of benefits. Our framework is well-suited to analyzing programs with not only binary extensive margin take-up decisions but continuous intensive margin take-up decisions as well. Intensive margin take-up decisions are a key determinant

[4]To the extent that providing home care in kind reduces informal care, it likely improves tax system efficiency (since informal care provision appears to reduce labor supply and wages Ettner, 1995; McGarry, 2006; Van Houtven et al., 2013) and may alleviate the Samaritan's Dilemma (since informal care provision, by reducing labor supply and worsening health (Do et al., 2015), may increase reliance on means-tested transfers in the future). In addition to these other potential benefits of in-kind provision, a full welfare analysis must account for any differences in administrative and other costs of different benefit types as well.

[5]Moffitt (1989), Whitmore (2002), Hoynes and Whitmore Schanzenbach (2009), and Hastings and Shapiro (2017), for example, analyze the effects of in-kind food transfers on consumption. A rare exception is Cunha et al. (2011) who find that in-kind provision of food transfers reduced food prices in Mexican villages.

of the welfare consequences of alternative benefit designs in contexts in which demand for the benefit is highly heterogeneous, such as in-kind health care benefits.

2 Theory

This section develops a theoretical framework for analyzing a central tradeoff for in-kind provision: in-kind provision can improve targeting at the expense of distorting consumption and being less valuable to recipients than a cost-equivalent cash transfer. In order to guide our analysis of home care insurance, we focus on the problem of insuring a risk, where the goal is to target high-marginal utility states. But with small adjustments, the framework can be applied to questions of redistribution across different types of people as well.

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

An in-kind benefit can be modeled as a (potentially non-linear) price subsidy. Many in-kind benefit programs, such as food stamps, offer individuals up to a fixed quantity of the good at no charge. When resale is not possible, this has the same effect on a participating individual's budget constraint as a non-linear price subsidy of 100 percent on units up to the benefit limit and 0 percent on units above the limit.[6] In this section we focus on the case of a subsidy program with no quantity limit. We do this both for simplicity of exposition and because in many states, including the states that ran the Cash and Counseling experiments, the Medicaid home care program does not appear to have binding benefit limits in practice. The results are easily extended to cases with benefit limits.

The key considerations for in-kind provision can be seen in Figure 1. Figure 1 shows the values (in terms of equivalent variations) and efficiency costs of a price subsidy on a particular

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

good for each of two people with different levels of demand for K. The price subsidy is worth less to each person than it costs the government or insurance company to provide due to the induced change in consumption. The size of this change is increasing in the compensated own-price elasticity of demand. The price subsidy is worth more to people who consume more of the subsidized good, so, relative to a cost-equivalent cash benefit, the subsidy redistributes toward people who consume more of the good from people who consume less of the good.

2.1 The benefit program and its budget constraint

Suppose an individual faces a risk with potential loss θ. The planner knows the distribution of the potential losses, $F(\theta)$, but cannot verify which state has occurred ex-post.[7]

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

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

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

where $p_K(\sigma)$ is the subsidy-exclusive price of K (the sellers' price) and $x_K(\sigma;\theta)$ is the consumption of K in state θ as a function of the subsidy rate.

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

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

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

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

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

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

The cash benefit must fall by the increase in average per-state spending on the in-kind benefit (subsidy). Average spending on the subsidy is the sum of two terms: (i) the mechanical increase in spending on the subsidy due to the increase in the subsidy rate, holding fixed consumption of K in each realized state, $E_{\Theta}\left(x_K(\sigma;\theta)p_K^0\right)$ ("mechanical effect"); and (ii) the increase in spending on the subsidy due to the induced change in consumption of K in response to the shift in program benefits, $E_{\Theta}\left((\sigma p_K^0)\frac{dx_K(\sigma;\theta)}{d\sigma}\right)$ ("behavioral effect").[8]

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

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

$$\frac{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial p_K(\sigma)} \frac{dp_K(\sigma)}{d\sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}} = x_K(\sigma;\theta)p_K^0,$$

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

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

(in units of income) of a budget-neutral marginal shift toward in-kind benefits of

$$\frac{dV(\sigma;\theta)}{d\sigma} \equiv \frac{\frac{dv(p(\sigma),m(\sigma,B);\theta)}{d\sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}} = \frac{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial p_K}\frac{dp_K(\sigma)}{d\sigma} + \frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}\frac{\partial b(\sigma,B)}{\partial \sigma}}{\frac{\partial v(p(\sigma),m(\sigma,B);\theta)}{\partial m}}$$
$$= x_K(\sigma;\theta)p_K^0 - \int_\Theta \left[x_K(\sigma;\theta)p_K^0 + (\sigma p_K^0)\frac{dx_K(\sigma;\theta)}{d\sigma}\right] f(\theta)d\theta$$
$$= \left[x_K(\sigma;\theta)p_K^0 - E_\Theta\left(x_K(\sigma;\theta)p_K^0\right)\right] - (\sigma p_K^0)E_\Theta\left(\frac{dx_K(\sigma;\theta)}{d\sigma}\right). \quad (1)$$

The marginal net value for an individual with realized state θ of a budget-neutral marginal shift in benefits toward in-kind benefits is the net benefit of the resulting redistribution to that state (redistribution benefit), $\left[x_K(\sigma;\theta)p_K^0 - E_\Theta\left(x_K(\sigma;\theta)p_K^0\right)\right]$, which is greater for states with greater levels of demand for K, less the average marginal distortion cost from the induced change in consumption of K (distortion cost), $(\sigma p_K^0)E_\Theta\left(\frac{dx_K(\sigma;\theta)}{d\sigma}\right)$.

Equation 1 shows that the shift toward in-kind provision has two key effects. It redistributes toward realized states with above-average demand for the good, and it distorts consumption of the good. The extent to which a consumer with a particular realized state gains from a marginal shift toward greater in-kind provision is increasing in that state's level of consumption of the good and decreasing in the average sensitivity of the demand for the good across all states.

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

Ex-ante expected utility is

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

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

$$\frac{dEU(\sigma)}{d\sigma} = \int_\Theta \frac{dv(p(\sigma), m(\sigma,B);\theta)}{d\sigma}f(\theta)d\theta = \int_\Theta \lambda(\sigma;\theta)\frac{dV(\sigma;\theta)}{d\sigma}f(\theta)d\theta = E_\Theta\left(\lambda(\sigma;\theta)\frac{dV(\sigma;\theta)}{d\sigma}\right)$$
$$= Cov_\Theta\left[\lambda(\sigma;\theta),\ x_K(\sigma;\theta)p_K^0\right] - (\sigma p_K^0)E_\Theta\left(\lambda(\sigma;\theta)\right)E_\Theta\left(\frac{dx_K(\sigma;\theta)}{d\sigma}\right), \quad (2)$$

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

Equation 2 shows the key roles of heterogeneity in the level of demand for K and the sensitivity of the demand for K to the composition of benefits in determining the welfare consequences of in-kind provision. The extent of heterogeneity in the demand for K and the extent to which it is correlated with marginal utility determine the targeting benefit of in-kind provision. The greater is the covariance across states in marginal utility and the demand for K, the greater is the targeting benefit of in-kind provision. The sensitivity of the demand for K to the composition of benefits determines the distortion cost of in-kind provision. The greater is the sensitivity of the demand for K to the composition of benefits, the greater is the distortion cost of in-kind provision.[9]

This analysis reveals three key determinants of the welfare effects of in-kind provision. The first is heterogeneity within benefit-eligible states in the demand for K. This determines the extent to which in-kind provision concentrates benefits in certain eligible states and not others. The second is the sensitivity of the demand for K to the composition of benefits. This determines the moral hazard cost of in-kind provision and the value to recipients of the in-kind benefit. The third is the covariance across states in the demand for K and marginal utility. This covariance—which is increasing in the variance in the demand for K and in marginal utility as well as the correlation between the two—determines the targeting benefit of in-kind provision. In the following sections, we investigate these key determinants of the welfare effects of in-kind provision in the context of home care insurance.

3 Home Care Risk and Insurance

Chronic health problems are the source of one of the most important risks people face over the life cycle. Most people will at some point in their lives develop severe health problems that limit their ability to perform activities of daily living (ADL) such as bathing, eating, dressing, and managing their household without significant, time-intensive assistance. The amount and costs of the assistance any one person will need are highly variable. Roughly 15 percent of Americans over age 50 have at least one person helping them due to ADL limitations. The vast majority of those receiving help (87 percent) live in the community (the rest live in care-giving facilities, mainly nursing homes), and 74 percent of all care hours occur in private homes (Barczyk and Kredler, 2016). Spending on formal home care was $88 billion in 2015 (Reaves and Musumeci, 2015; Centers for Medicare and Medicaid Services,

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

2017), 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 services (Arno et al., 1999).[10] Despite the magnitude of this risk, just 10 percent of people 65 and older own private long-term care insurance, and 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, over half of its total spending on long-term care (Kaiser Commission on Medicaid and the Uninsured, 2016). Eligibility for Medicaid home care benefits 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 persist 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 medical examination, though in the specific cases we analyze there does not appear to be a binding upper limit (see Appendix D). But 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). These programs tend to allow people to spend their benefits on a wide range of personal care goods and services, including assistive 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, for example, 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- vs. less- flexible benefits, and an important source of evidence in our paper, is the Cash and Counseling demonstrations. These were large-scale experiments run by Medicaid programs in Arkansas, Florida, and New Jersey that began in 1998. Participants were randomly assigned to either the traditional in-kind home care benefit or a near-cash benefit.[11] The main goal of the experiments was to test whether recipients could effectively manage their cash benefits and receive "enough" care. The results were almost uniformly positive. Members of the cash-benefit treatment group

[10] Arno et al. (1999) estimate that in 1997, the economic value of informal care alone was more than $196 billion dollars.

[11] Appendix B contains more information about Medicaid home care and the Cash and Counseling experiments, including summary statistics of Cash and Counseling participants and balance tests which provide evidence of a valid randomization. The near-cash benefit was a cash budget that had to be spent on personal care services. This requirement was unlikely to be binding in practice since the vast majority of participants had been receiving enough informal care at baseline to more than exhaust their benefit. Participants randomized to near-cash could revert to the standard in-kind benefit at any time; those randomized to the in-kind benefit could not choose to join the near-cash treatment.

reported greater satisfaction with their care (Foster et al., 2003) and with their lives as a whole (Brown et al., 2007) and had similar, if not better, health outcomes (Lepidus Carlson et al., 2007). In the official final report on the experiments, Brown et al. (2007) conclude that the near-cash transfer had overwhelmingly positive effects on recipients.

That recipients prefer more flexible transfers is an important cost of providing home care in kind that must be weighed against any benefits. But despite the rich evidence from the Cash and Counseling experiments, little is known about the potential benefits of in-kind provision, whether for Medicaid home care or for other programs more generally (Currie and Gahvari, 2008). A potential benefit likely to be important in many contexts, including home care, is better targeting.

The targeting effects of in-kind provision, as discussed in Section 2, are increasing in the heterogeneity in demand for the good within the eligible population. Data from the National Long Term Care Survey (NLTCS), a nationally representative survey of Americans 65 and older, indicate that the demand for formal home care is highly heterogeneous within the eligible population.[12] Figure 2 shows the distribution of formal care consumption in the population eligible for home care benefits (people with two or more ADL limitations living in the community). Even within this group of people with severe chronic health problems, there is significant heterogeneity in demand for formal care. 62 percent do not consume any formal care, but among those who do there is a long right tail. Conditional on consuming any care, median consumption is 14 hours per week ($11,000 per year at the average market price) and the 95th percentile is 168 hours per week ($131,000 per year). Such heterogeneity within the benefit-eligible population means that in-kind provision likely has important targeting effects, concentrating benefits on those with the greatest demand for formal care. The heterogeneity also suggests that a cash benefit would leave significant risk uninsured, since heterogeneity in spending on formal care translates into heterogeneity in the resources available for non-care consumption.[13]

Whether the targeting effects of in-kind provision can be achieved more directly by "tagging" cash transfers (Akerlof, 1978) depends on how well the heterogeneity in the demand for formal care can be predicted by verifiable, and ideally immutable, characteristics. Appendix C shows that the vast majority of the variation in formal care consumption cannot be explained by even an extensive set of individual and household characteristics, including care plans based on detailed medical exams. The share of the variance that even non-parametric,

[12]Basic summary statistics of the 65 or older population with at least two ADL limitations are presented separately for those who take up Medicaid home care and those who do not in Table 4.

[13]Of course, some of the heterogeneity in the cross section reflects permanent or predictable heterogeneity across different types of people rather than heterogeneity in realizations of a common risk faced by ex-ante identical individuals.

machine-learning models can predict out of sample never exceeds about 20 percent. A key consequence of this unexplained variation is that even extensively-tagged cash transfers—and even ignoring the verification and moral hazard costs of using certain tags—would leave much of the heterogeneity in formal care and non-care consumption uninsured.

A likely cause of the unexplained variation in the demand for formal care is hard-to-verify heterogeneity in both health problems and the costs of coping with a given set of health problems. Among people with the same severe chronic health problems, for example, the cost of coping with those problems is likely to be much greater for those who do not have good informal care options. But it may be difficult for insurers, whether private insurers or government programs, to condition benefits on such differences. To the extent that the cost of coping with bad health varies widely within states of the world that insurers cannot easily distinguish from one another—as suggested by the likely difficulties of verifying differences in health and coping costs and by the substantial residual variation in formal care consumption conditional on even large sets of characteristics—benefits cannot target high-cost states directly. Hard-to-verify heterogeneity in the costs of bad health therefore appears to be a significant barrier to insuring the risk from chronic health problems, even with mandatory government programs. Such heterogeneity introduces a potential targeting rationale for in-kind provision. Whether any such targeting benefits of in-kind provision outweigh the cost of the reduced value of benefits to recipients depends, as discussed in Section 2, on the sensitivity of demand for the good to the composition of benefits and the covariance across states of the world between marginal utility and demand for the good. We now turn to investigating these objects in the context of formal home care.

4 Moral Hazard Effects of In-Kind Provision and the Value of In-Kind Benefits

As shown in Section 2, the welfare effects of in-kind provision depend on the sensitivity of demand to the composition of benefits, which determines the moral hazard cost of in-kind provision and how much recipients value the in-kind benefit. We use the Cash and Counseling experiments to estimate the slope of this demand curve.[14] The Cash and Counseling experiments have two major advantages for estimating the slope of this demand curve. First, the randomization solves an especially difficult simultaneity problem, since many factors that

[14]Previous research on the Cash and Counseling demonstrations has focused on *paid* and *unpaid* home care, whether provided by professionals or family and friends, rather than on formal care (provided by professionals). For example, Carlson et al. (2007) and Brown et al. (2007) compare hours of paid care, unpaid care, and total hours of care across the in-kind and near-cash groups. To our knowledge, we are the first to use the Cash and Counseling experiments to estimate the slope of demand for formal care.

shift the supply of formal care are also likely to shift the demand for formal care by changing the opportunity cost of informal care.[15] Second, the variation in the composition of benefits spans the full range most relevant for policy, from a pure in-kind benefit to a near-cash benefit.

The experimental results provide strong evidence 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 consumption of formal care from 7.1 to 14.8 hours per week. Figure 3 shows that in-kind provision increases formal care consumption throughout the distribution, more than doubling the fraction of people who consume formal care (from 21 to 55 percent) and increasing 95th-percentile consumption by 15 hours per week.

We estimate the sensitivity of the demand for formal care to the composition of benefits 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 left Medicaid home care altogether and some people assigned to the near-cash benefit reverted to the traditional in-kind benefit—by using the randomized assignment as an instrument for the price each participant faced.[16] 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 gender, education level, race, self-rated health, 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. Absent income effects of demand for formal care, β is sufficient for analyzing counterfactual policies that affect the relative price of formal care, regardless of any effects they might have on

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

[16]Participants in the near-cash group or who leave Medicaid home care face the market price in their state. Participants in the in-kind group face a price of zero. Although quantity limits from care plans or maximum benefit rules could in principle raise the shadow price of formal care above zero, a variety of evidence suggests that the marginal value of Medicaid formal care is zero for most recipients of the traditional in-kind benefit (see Appendix D). Appendix D also tests the robustness of the results to alternative assumptions about the marginal value of Medicaid formal care.

income.[17] As a baseline, we assume that (ε_i, ν_i) are jointly normal and estimate this system using an instrumental variables Tobit specification.

The first-stage results are presented in Table 2, and the instrumental variables estimate of β is presented in Table 3. As one would expect given the nature of the experiment, the first stage relationship is economically and statistically large (being assigned to the in-kind benefit decreases the average price of formal care by \$8.84, and the F-statistic exceeds 1,000) and adding control variables has little effect on any of the estimates. The instrumental variables estimate implies that a one-dollar increase in the hourly price of formal care reduces consumption by 1.8 hours per week. Evaluated at the sample means, this implies an elasticity near -1.2. As discussed in Appendix D, the conclusion that the demand for formal care is quite sensitive to its price is robust to a wide range of alternative assumptions about the distribution of the error terms and benefit limits, and it holds in each of the three states.[18]

The estimated slope of demand implies that in-kind provision has a large moral hazard cost. Someone consuming the average amount of formal care among participants randomized to in-kind benefits (14 hours per week), for example, would not consume any formal care without the subsidy and values the care she does receive at only 28 percent of its cost. Someone consuming 25 hours of formal care per week values that care at approximately 51 percent of its cost. These implications of the estimated demand curve are qualitatively consistent with earlier research documenting negative effects of being assigned to the in-kind benefit on satisfaction with one's care and life as a whole (Foster et al., 2003; Brown et al., 2007).

5 Targeting Effects of In-Kind Provision

We investigate the targeting effects of in-kind provision of home care using both nationally representative data from the NLTCS and the experimental variation in the Cash and Coun-

[17]With non-zero income effects of demand for formal care, this parameter is appropriate for analyzing the Cash and Counseling experiments but not policies with different cash benefits. The Cash and Counseling experiments roughly hold fixed spending on each participant of the experiments—a group whose average in-kind benefit is much greater than the average among the population of eligibles. Cash and Counseling's near-cash benefits were therefore on average greater than those under the main policy counterfactual we have in mind, which, as discussed in Section 2, holds fixed total spending on the program. With positive income effects of demand for formal care, our estimates will tend to understate the true moral hazard costs of the policies of interest.

[18]In Appendix D we also discuss the generalizability of these results to other populations and policies of interest. There are two key issues that tend to offset one another. First, people whose demand was more sensitive to price had a greater incentive to participate in the experiment. Second, the nature of the experiment—especially its unexpected occurrence and uncertain duration—likely reduced the size of the responses relative to those that would be expected under an anticipated, permanent change in policies. In light of the possible issues with generalizability, in our welfare analysis (Section 6), we test the robustness of our results to a wide range of values of the slope of demand.

seling demonstrations. As discussed in Section 2, better targeting means a higher covariance between benefits and marginal utility. Because we only observe one state of the world for any individual (the actual realization of uncertainty), we investigate the cross-sectional distribution in the population. Observed differences therefore reflect permanent heterogeneity in addition to differences in the realization of uncertainty. Since marginal utility is not observable, we summarize the relationship between benefits received and various observable characteristics likely to be associated with marginal utility. We focus on three sets of characteristics that both empirical evidence and theoretical reasoning suggest are closely linked to marginal utility in our context: formal care consumption, proxies for informal care costs, and health. The greater is someone's formal care consumption and the worse are someone's informal care options and health, the greater are the costs of coping with bad health. Greater costs of coping with bad health leave fewer resources for non-care consumption. In many models, this means higher marginal utility.[19]

In-kind transfers can have targeting effects on both the extensive and intensive margins. On the extensive margin, if taking up benefits is costly, people with relatively low demand might not join the program. This concentrates benefits on those who do join. We investigate take-up of Medicaid home care benefits among the eligible population using nationally representative data from the NLTCS. Take up of Medicaid home care reflects the combined effects of not only in-kind provision but also other features of Medicaid home care, including awareness of the program, hassle costs of taking up benefits, and stigma of participating. The first three rows of Table 4 show estimates of the fraction of people eligible for Medicaid home care that take up benefits (see Appendix B for details). We find that only 4–16 percent of those eligible take up benefits. The low take-up rate implies a significant concentration of benefits within the eligible population: Benefits per recipient are between 6 and 24 times greater than they would be under a hypothetical program with the same budget and 100 percent take up.

Whether the targeting induced by extensive-margin take-up decisions improves insurance depends on whether take-up is greater in higher-marginal utility states. The next several rows of Table 4 compare the characteristics of those who do versus do not take up benefits. People who take up have a much greater demand for price-adjusted formal care, are sicker,

[19] Although spending on formal care is far from the only cost of bad health, high formal care consumption seems likely to be the best indicator of high marginal utility in this context. That many private long-term care insurance contracts subsidize the consumption of formal care is suggestive revealed-preference evidence that formal care consumption is positively related to marginal utility. Moreover, many models of formal care consumption, including the standard model of health risks in which health spending is equivalent to a wealth shock, predict a (usually strong) positive link between formal care consumption and marginal utility. Formal care consumption likely reflects the combined influence of health, informal care options, and other determinants of coping costs. Differences in formal care consumption are not offset by differences in informal care. In the Cash and Counseling experiments, the correlation between formal and informal care hours is roughly zero.

and appear to have worse informal care options.[20] Even controlling for an extensive set of personal and household characteristics, including measures of health and proxies for informal care costs, those who take up Medicaid home care consume much more formal care (see Appendix Table G.1). These results are consistent with in-kind provision and other aspects of Medicaid home care affecting take-up decisions in a way that targets relatively high-marginal utility states. Although in principle awareness, stigma, or other factors could have led to "perverse targeting" in which those most desperate for help were least likely to receive it, in practice take-up is strongly increasing in the demand for formal care.

Unlike cash benefits, in-kind benefits can have important targeting effects on the intensive margin among those who take up benefits as well. We investigate the targeting effects of in-kind provision on the intensive margin using the Cash and Counseling experiments. Whereas decisions to take up Medicaid home care could reflect a variety of factors, the experimental design isolates the effect of in-kind provision. As a result, there are stronger reasons to expect in-kind provision to target people with a greater demand for formal care, since factors such as take-up costs and awareness are less important conditional on participating in the program. The key question we use the experiments to address is therefore slightly different: How do the targeting effects of in-kind provision compare to the targeting effects of Cash and Counseling's tagged near-cash benefit (based on medical reviews)?

Figure 4a shows kernel density plots of benefits for members of the in-kind and near-cash groups.[21] In-kind provision concentrates benefits substantially. The variance in benefits received is 7 times greater in the in-kind group, with much larger fractions of very low and very high benefits. The fraction of people who received no benefit is over three times larger in the in-kind group (31 percent vs. 10 percent), and 17 percent of the in-kind group received benefits whose cost is at least as great as the 99th percentile benefit of the near-cash group. Figure 4b plots differences in benefits between the in-kind group and either the near-cash group or a hypothetical pure-cash benefit group (each of whom receives an identical, untagged cash transfer equal to the per-participant average benefit in the in-kind group). Even compared to Cash and Counseling's tagged near-cash transfer, in-kind

[20]For this table, we have adjusted each individual's formal care consumption for differences in prices in order to isolate differences in the level of demand. In particular, we use our estimated price sensitivity of demand to simulate each individual's consumption if she faced a price of $18.50 per hour, the highest price in the data.

[21]The figure is based on data from Arkansas because it is the only state with information on care plan hours, which is necessary to estimate the size of the near-cash benefits. The near-cash benefit is the product of care plan hours and the hourly price of care. The in-kind benefit amount is the product of hours of care received and the hourly price of care. Because our interest is in the concentration of benefits, we scale up the near-cash group's benefit to have the same mean as the in-kind group (the near-cash group's average benefit was slightly smaller than that for the in-kind group). In practice, this leads our reported measures of differences in concentration to be smaller than those calculated with the near-cash group's unscaled cost data.

provision significantly concentrates benefits on the intensive margin. The relative lack of targeting by the tagged near-cash benefit reinforces the evidence discussed in Section 3 that the vast majority of the variation in the demand for formal care cannot be predicted by even extensive sets of individual and household characteristics.

Figure 5 shows average benefits among those randomized to the in-kind and near-cash groups by percentile of the distribution of formal care consumption. Because formal care consumption is highly concentrated even among participants of the Cash and Counseling experiment, in-kind benefits are highly concentrated as well. Whereas the average in-kind benefit is $133 per week, those between the 91st and 95th percentiles of the formal care distribution receive an average of $350 per week and those above the 95th percentile receive an average of $843 per week—almost 7 times the average benefit. The tagged near-cash benefits, by contrast, are roughly constant throughout the formal care distribution. People receive similar benefits regardless of how much formal care they consume, leaving those who consume more formal care with fewer resources available for non-care consumption. Appendix E provides suggestive evidence that in-kind provision targets on the intensive margin states in which people are sicker and have worse informal care options as well.

Taken as a whole, these results show that in-kind provision sharply concentrates benefits on a small subset of benefit-eligible states in which people are sicker, have worse informal care options, and have a greater demand for formal care. These results are consistent with in-kind provision having a large insurance benefit. The much greater targeting effects of the in-kind benefit than of the tagged near-cash benefit, especially when combined with the evidence discussed in Section 3, suggests that the potential for targeting these states directly using tagged cash benefits is quite limited. The targeting effects of in-kind provision are therefore unlikely to be achievable with alternative, less costly means of targeting. This raises the question of whether the targeting benefit of in-kind provision outweighs the moral hazard cost, the question to which we now turn.

6 Welfare Effects of In-Kind Provision: Targeting Versus Moral Hazard

6.1 Model

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

$$A + pF = m,$$

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

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

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

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

This utility function is motivated by key evidence from our setting. It produces a simple function for the demand for formal care that is consistent with the sensitivity of formal care consumption to its price and that people become satiated at finite levels of formal care consumption.[22] This utility function also has several appealing features. It nests as a special case the widely-used model in which health spending is equivalent to a wealth shock.[23] It implies that the demand for formal care is linear in its price within the range of prices in which the individual is not at a corner. It has an intuitive interpretation: Utility is decreasing in any unmet, residual health needs, $(\alpha - F)$, the size of which is decreasing in formal care consumption, F, and increasing in the level of demand for formal care, α. This captures the idea that certain health problems are costly for people to cope with on their own. Marginal utility of income depends on the demand for formal care mainly through the budget constraint: Greater spending on formal care means lower non-care consumption.

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

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

6.2 Baseline parameter values

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

Figure 6 presents our main estimate of the distribution of the level of demand for formal care, $F(\alpha)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits a long right tail (the mean far exceeds the median) and that most of the mass is at low values.

The remaining parameters take standard values. We follow most of the literature on health spending risks and use a constant relative risk aversion utility function, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ (e.g., Brown and Finkelstein, 2008; De Nardi et al., 2010; Ameriks et al., 2011). In our model, the argument c is "net consumption," non-care consumption net of any residual coping costs, $c = A - \frac{(\alpha-F)^2}{2\beta}$. We follow Brown and Finkelstein (2008) and many others in taking as a baseline value a coefficient of relative risk aversion, γ, of three. Income before transfers, m, is \$15,000 per year. The distribution of before-subsidy prices of formal care is the empirical distribution observed in the NLTCS. People who cannot achieve net consumption of at least $\bar{c}$ =\$5,000 per year receive transfers that enable them to enjoy net consumption of \$5,000 per year (a consumption floor). This is meant to approximate the important means-tested programs, Medicaid and Supplemental Security Income.

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

6.3 Welfare effects of in-kind provision

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

Figure 7 summarizes the key results. It shows the equivalent variation of the mixed in-kind and cash benefit policy as a function of the in-kind component, the subsidy rate σ. The optimal subsidy rate is 88 percent, close to a pure in-kind program (under which the after-subsidy price is zero). The optimal subsidy increases welfare substantially relative to a pure-cash benefit program. In order to make the individual as well off as she is with the optimal policy under an alternative pure-cash benefit program, the cash benefit would have to be about 80 percent greater than the average cost of the in-kind program. Figure 7 also shows, however, that the optimal subsidy is significantly less valuable than the hypothetical first-best policy. The optimal in-kind subsidy achieves 59 percent of the incremental value over a pure-cash benefit that the first-best policy does.

Table 5 shows a variety of outcomes in several versions of the model. The purpose of this table is to provide intuition for and assess the robustness of the key results. The first column of the table shows results for the baseline specification just discussed. The key tradeoff involved in increasing the in-kind component of the benefit can be seen clearly by comparing the average level of and dispersion in non-care consumption under the optimal subsidy program and under the cost-equivalent pure-cash benefit program. The optimal subsidy reduces average non-care consumption due to the consumption distortion (and to a lesser extent due to foregone transfers from the consumption floor), but it also greatly reduces the dispersion of non-care consumption, as measured by the standard deviation. Under the pure-cash program the standard deviation of annual non-care consumption is 4.5 times greater than under the optimal program, $5,610 vs. $1,237.

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

The other columns of the table test the robustness of the results to making different assumptions about the key ingredients of the model. The price sensitivity of demand for formal care must be quite large—over 10 times larger than we estimate based on evidence from the Cash and Counseling experiment—in order to overturn the conclusion that the optimal subsidy is large. Even if the distribution of partially-identified α values is in the "worst-case" configuration (i.e., each α_i equal to the maximum value consistent with observed behavior), the optimal subsidy rate is still 86 percent. The utility function must exhibit strong state dependence of just the right kind—greatly decreasing the marginal utility in states with high demand for formal care in just the right way—in order to overcome the fact that, holding other resources constant, greater formal care consumption leads to lower non-care consumption. Although the right tail of the distribution of demand for formal care is an important determinant of the targeting benefit and so the optimal subsidy, the optimal subsidy remains large even when the right tail of the distribution is chopped off or when all of the α values are scaled down. If states in which a person consumes more than 50 hours per week of care are dropped, the optimal subsidy is 59 percent. If all of the α values are cut in half, the optimal subsidy is 75 percent. Finally, a combination of relatively low risk aversion together with a relatively generous consumption floor can overturn the optimality of a large subsidy on formal care, although this reflects the undesirability of *any* insurance—including a first-best contract—in situations in which means-tested programs are sufficiently attractive rather than any undesirability of in-kind benefits per se. Appendix F discusses these robustness tests in more detail.

The alternative specifications also provide information about the key factors driving the results. As expected, the net benefit of subsidizing formal care is decreasing in the price sensitivity of demand for formal care. When demand for formal care is completely inelastic ($\beta = 0$), a 100 percent subsidy achieves the first best.[24] The targeting benefit of in-kind provision is increasing in risk aversion and decreasing in the generosity of alternative insurance arrangements, such as any consumption floor or means-tested programs. The targeting benefit of subsidizing formal care is increasing in the extent to which there is state-dependent utility in which marginal utility is greater in states with greater demand for formal care (above and beyond the effects operating through the budget constraint or residual coping costs). If such state-dependence is strong enough, it is optimal to more than fully subsidize formal care (columns 8 and 11).

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

6.4 Welfare effects of tags

This section extends the analysis to the case in which different groups of people, defined by their verifiable characteristics, can be offered different benefits. We estimate the gains from catering benefits to different groups of people defined by whether they live alone and the number of activities of daily living limitations they have (2–4, 5, and 6), the two strongest

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

predictors of formal care consumption uncovered in Section 5.[25] The procedure is the same as that in the last section, except that we estimate different α distributions for different groups of people and allow the program to offer different benefits to people in different groups. Figures of the α distributions of each group are reported in Appendix F.

Table 6 shows that the ex-ante welfare gain from using tags to target high-marginal utility states is quite small. The incremental welfare gain from optimally tagging a pure-cash benefit based on whether someone lives alone is \$227, just 4 percent of the gain from an optimal un-tagged mixed benefit. The incremental welfare gain from optimally tagging benefits based on the number of activities of daily living limitations someone has is even smaller. The fundamental reason for tags' ineffectiveness in insuring this risk is that much of the heterogeneity in demand for formal care occurs within rather than across states that can be distinguished on the basis of their verifiable characteristics; the correlation between marginal utility in the absence of any program and the optimal tagged cash benefit is just 0.20 with the "lives alone" tag and 0.05 with the "number of activities of daily living limitations" tag. These results are consistent with those of Mankiw and Weinzierl (2010) on the effects of using height as a tag for optimal income taxation.[26] Although different combinations of observable characteristics could potentially improve on those we have analyzed here, both the small gains from tags based on two of the strongest predictors of formal care consumption and the limited extent to which observable characteristics predict formal care consumption (as discussed in Section 5) suggest that the scope for tags is limited and reinforce the conclusion that in-kind benefits have an important role to play in terms of targeting benefits to high-marginal utility types.

6.5 Discussion of results

This analysis is subject to several caveats. It assumes that people's decisions about consumption are rational. This ignores possible paternalistic rationales for in-kind transfers, which could be important in the case of home care given the cognitive health problems from which some recipients suffer. Such considerations, which were one of the main motivations for the Cash and Counseling experiments, seem likely to increase the value of in-kind as opposed to cash transfers in this context. The analysis also abstracts from any costs of taking up. This is done to focus on the core tradeoff at the heart of in-kind provision, but it is important

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

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

to note that many in-kind programs have low take-up rates, whether from low knowledge about or high costs of taking up the programs. It assumes that all ex-post heterogeneity is the outcome of an exogenous process. This feature, which is shared by the vast majority of the large literature on optimal taxation, rules out ex-ante moral hazard (effects of policies on the distribution of ex-post types), which tends to increase the net value of insurance or redistribution. It focuses only on home care and does not explicitly model substitution across other types of care. This was done for simplicity given that there appears to be little substitution across different types of long-term care (Grabowski and Gruber, 2007; Kemper, 1988). Finally, it focuses on singles in order to avoid the many complexities involved in modeling couples, including any financial risk sharing and utility consequences of different caregiving arrangements. An analysis of couples and extended families is an interesting topic for future work.

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

7 Conclusion

We develop a general approach for analyzing a central tradeoff inherent to in-kind provision—in-kind provision can improve the targeting of benefits at the cost of being less valuable to recipients—and apply it to home care. Despite the ubiquity of in-kind transfers and the centrality of this tradeoff for their welfare effects, little is known about the magnitude of these key costs and benefits. We find that for home care, the targeting benefit of in-kind provision appears to exceed its large moral hazard cost. The main factor driving this result is the significant, hard-to-verify heterogeneity in the demand for formal care—whether from hard-to-verify differences in underlying health or in the costs of coping with a given set of health problems—which implies significant heterogeneity in non-care consumption and so, in many models, in marginal utility.

In interpreting our results, it is important to keep in mind two caveats that are a high priority for future work. First, the magnitude of the targeting benefit depends crucially on the utility function, particularly any (state-)dependence of utility on the unverifiable heterogeneity that motivates in-kind provision in the first place. This caveat, common to all questions about the optimal design of insurance programs, is extremely important. This concern is mitigated somewhat by the fact that our comparisons are within bad-health states (two or more ADL limitations) and that key features of private and government long-term care insurance contracts are suggestive about the net, average effects of any such state-dependence. But improvements to our knowledge about this would clearly be helpful. Second, we focus only on the targeting benefit and moral hazard cost of providing formal care in kind. We do not consider other potential benefits, including improving tax system efficiency, alleviating the Samaritan's dilemma, and paternalistic benefits, all of which might be important in this context. Neither do we consider any differences in the costs of administering or taking up different types of benefits. In the particular case of the Cash and Counseling near-cash benefit, any such differences in costs seem likely to be second order, so it seems likely that the net effect of the other costs and benefits outside our analysis would be to increase the relative attractiveness of in-kind provision of formal care.[27] Whether and the extent to which they do is an important topic for future work.

Our results have important policy implications. Several recent policy reforms and proposals have made or propose to make benefits that have traditionally been restrictive in-kind benefits more flexible and cash-like. This is true not only in the case of home care but in many other areas as well.[28] A major impetus for these proposals is the view that recipients would much prefer cost-equivalent cash transfers, a view that is consistent with our analysis of the particular case of Medicaid home care benefits. But another consequence of such reform proposals that tends to receive less attention is that they tend to systematically change the distribution of benefits received by different groups of people. To the extent that achieving a good targeting of benefits in any particular context is difficult without in-kind provision, as our analysis suggests is the case in home care, any gain from increasing the value of the benefit to recipients must be weighed against any resulting reduction in targeting efficiency.

[27] The Cash and Counseling near-cash benefit seems about as likely to involve larger as smaller costs than the traditional in-kind benefit. It requires the same medical exam to create a care plan (which, given the evidence that care plans are not binding for the traditional in-kind benefit, is higher-stakes for the near-cash benefit). It requires counseling that participants might value less than cost (otherwise the requirement would not be necessary). And it requires that recipients track and document their spending and that the program monitor this spending. These aspects of the Cash and Counseling near-cash benefit may make it an exception to the general rule that more flexible, cash-like benefits tend to be less costly to administer and take up than in-kind benefits. Of course, any such cost differences are central to the welfare effects of different types of benefits.

[28] The European Commission, for example, recently suggested that policymakers "always ask the question, 'Why not cash?"' U.N. Secretary-General Ban Ki-moon has argued that "cash-based programming should be the preferred and default method of support." The desirability of a universal basic income is the subject of an active debate in many countries. Switzerland held a referendum in 2016 on whether to have a universal basic income. (It was rejected with a 77 percent majority.) Pilot programs for universal basic income guarantees are planned in Finland and Canada. GiveDirectly is providing a universal basic income to thousands of people in dozens of villages in Kenya.

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Carlson, B. L., 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.

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

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

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

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

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

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

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

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

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

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

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

Do, Y. K., E. C. Norton, S. C. Stearns, and C. H. Van Houtven (2015). Informal care and caregiver's health. *Health Economics 24*, 224–237.

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

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

Finkelstein, A. and K. McGarry (2006, September). Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review 96*(4), 938–958.

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

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

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

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

Greaney, B. P., J. P. Kaboski, and E. Van Leemput (2016). Can self-help groups really be "self-help"? *Review of Economic Studies 83*(4), 1614–1644.

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

Hendren, N. (2013, 09). Private information and insurance rejections. *Econometrica 81*(5), 1713–1762.

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

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

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

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

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

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

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

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

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

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

McGarry, K. M. (2006). Does caregiving affect work? In D. A. Wise and N. Yashiro (Eds.), *Health Care Issues in the United States and Japan*, pp. 87–106. University of Chicago Press.

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

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

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

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

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

Overseas Development Institute (2015, September). Doing cash differently: How cash transfers can transform humanitarian aid. Report of the High Level Panel on Humanitarian Cash Transfers. Technical report.

Pauly, M. V. (1990, February). The rational nonpurchase of long-term-care insurance. *Journal of Political Economy 98*(1), 153–168.

Reaves, E. L. and M. Musumeci (2015, December). Medicaid and long-term services and supports: A primer. Technical report, Kaiser Family Foundation.

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

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

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.

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

Tables and Figures

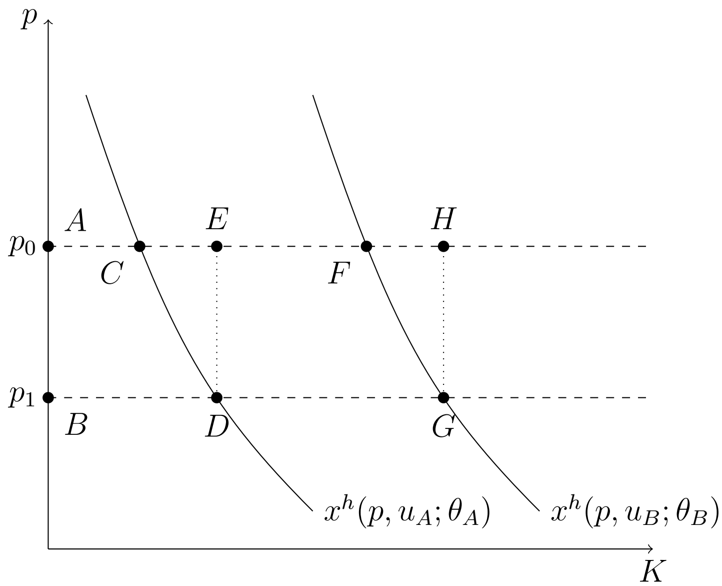

Figure 1: Equivalent variations and excess burdens of a subsidy

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

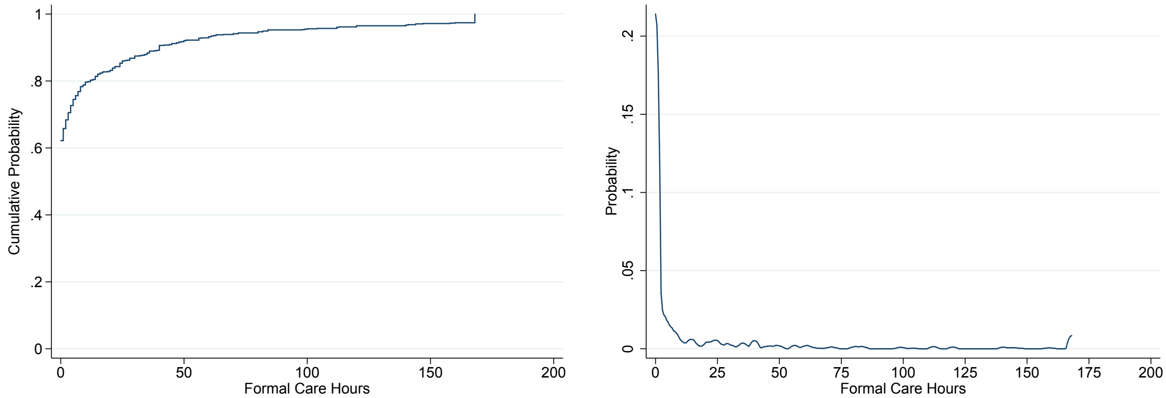

(a) CDF

(b) PDF

Figure 2: Distribution of Formal Care, NLTCS

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

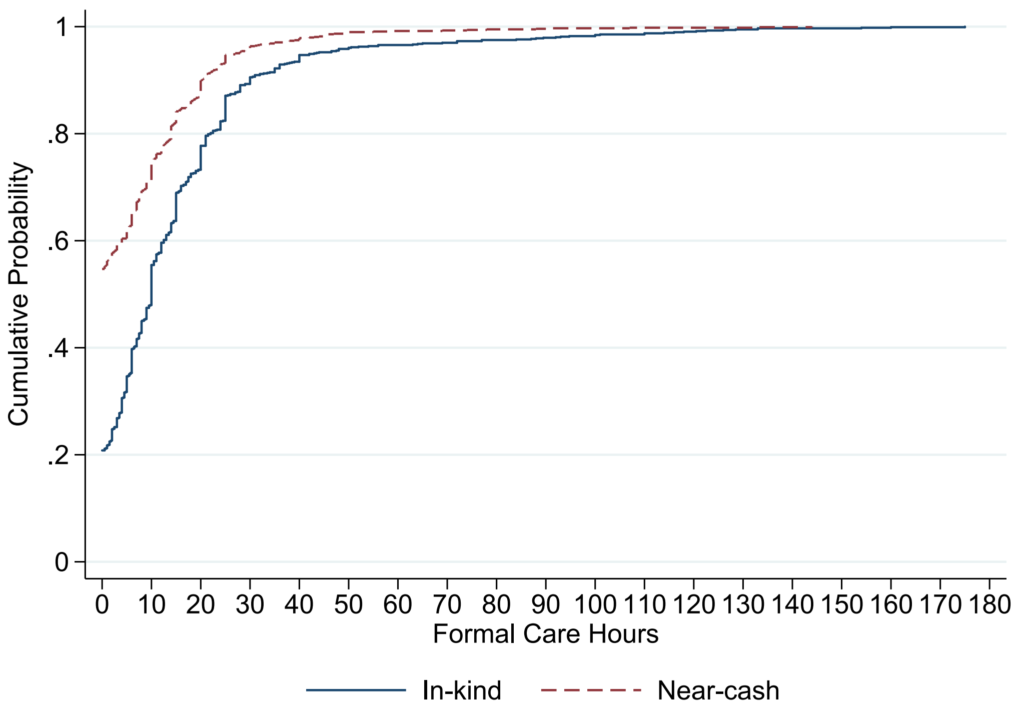

Figure 3: CDFs of Formal Care by Randomized Assignment
[Data from the Cash and Counseling follow-up survey. Hours of formal home care per week.]

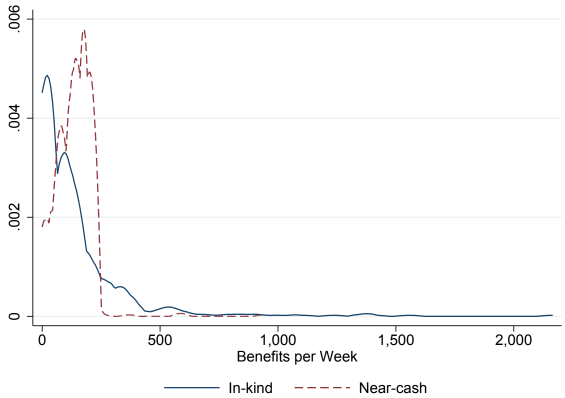

(a) Distributions of Benefits

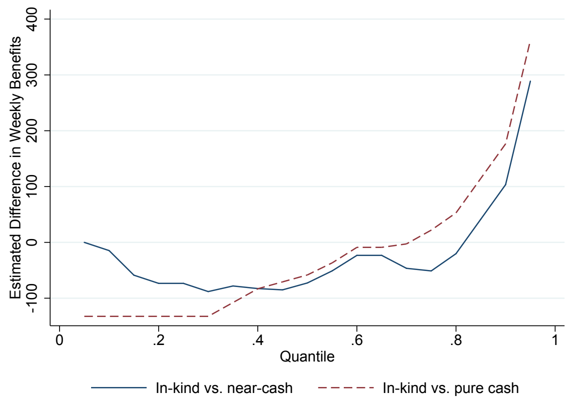

(b) Differences of Benefits

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

[Data from the Arkansas Cash and Counseling experiment. Benefits are costs per week (in dollars, at market prices). Kernel density estimates shown for each group on the left. Groups are based upon each individual's randomized assignment. A pure cash program is a hypothetical transfer that provides everyone in the Arkansas Cash and Counseling experiment the average dollar benefit of the in-kind group, approximately $133 per week. The figure to the right shows benefits for the near-cash or pure cash transfer subtracted from the in-kind transfer.]

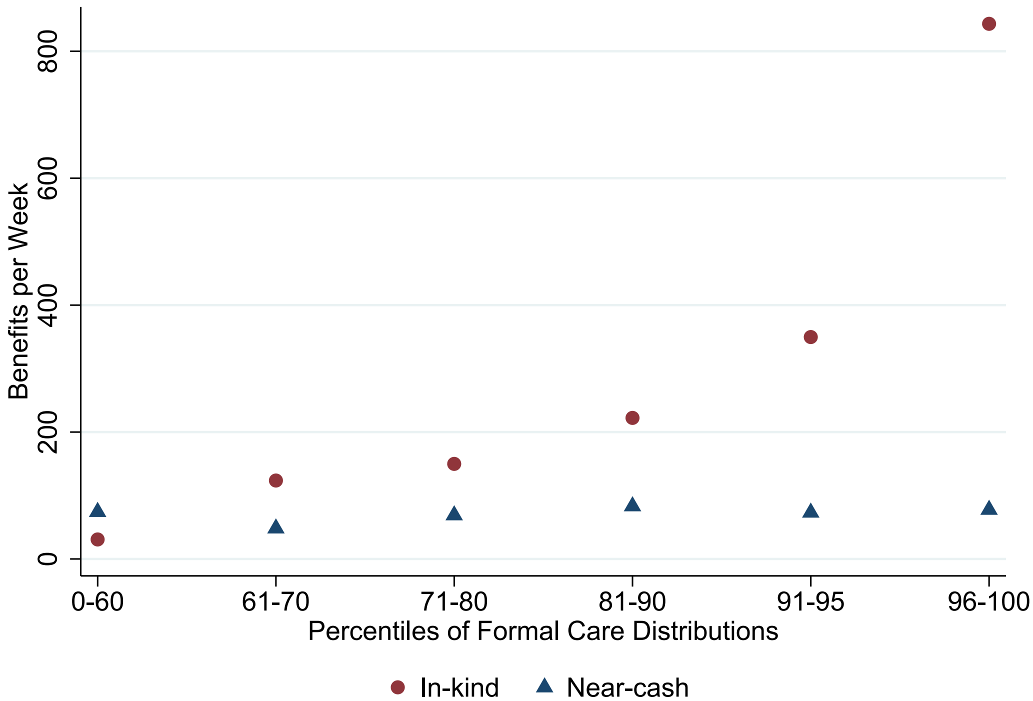

Figure 5: Targeting in the Arkansas Cash and Counseling Experiment

[Data from the Arkansas Cash and Counseling experiment. Average benefits per week shown separately for those randomized to the in-kind and near-cash groups. Within groups, individuals are ranked by their use of formal care at follow-up to determine their percentiles. 57 percent of those randomized to near-cash do not consume any formal care.]

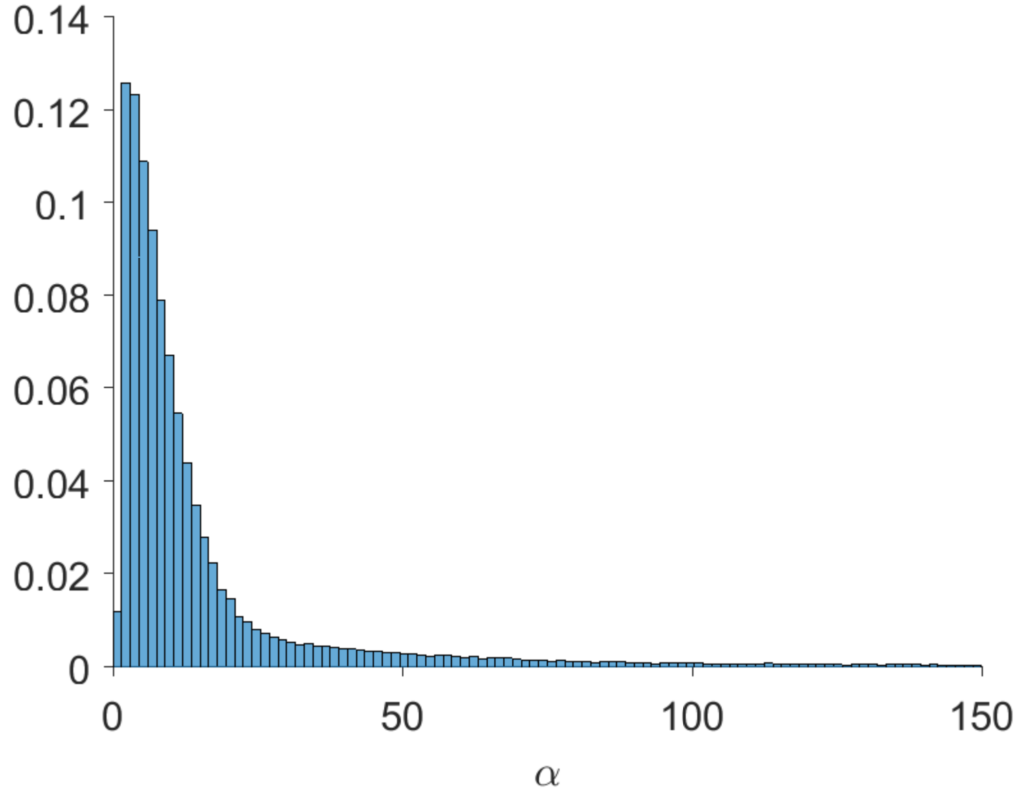

Figure 6: Distribution of the demand for formal care
[Simulated distribution of formal care satiation points, α, in hours per week. The population is people age 65 and older with at least two activities of daily living limitations. The mean is 21 hours per week.]

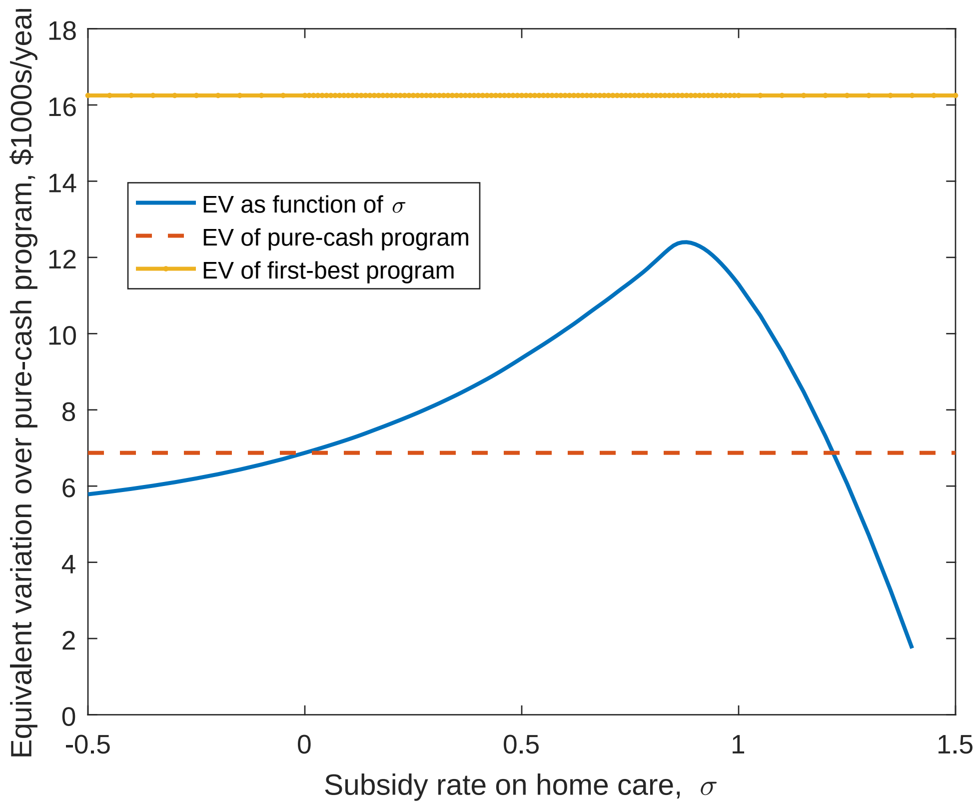

Figure 7: Equivalent variation of mixed cash/in-kind program as function of subsidy rate, σ
[Equivalent variation of mixed cash/in-kind program as function of the subsidy rate, σ. Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total program spending. $\sigma = 1$ corresponds to a pure in-kind benefit program (a 100 percent subsidy on formal care with no cash benefit). $\sigma = 0$ corresponds to a pure cash benefit program (a 0 percent subsidy on formal care).]

Table 1: Average Hours of Formal Care by Treatment Group

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

Means for formal care hours per week. Near-cash indicates group was randomized to the near-cash transer; in-kind indicates group was randomized to traditional Medicaid home care. P-value for test of equality across groups shown in last column. Rows denote different samples.

Table 2: Price Sensitivity of Demand for Formal Care, First Stage Estimates

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

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

Table 3: The Price Sensitivity of Demand for Formal Care

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

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

Table 4: Take-up 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	0.96	0.04	
Income eligible, no cars	0.90	0.10	
Restrictive income, no cars	0.84	0.16	
Summary Statistics			
Level of formal care demand	8.16	17.73	0.02
Age	82.06	82.76	0.53
Four or more ADLs	0.47	0.62	0.02
Health fair or poor	0.63	0.78	0.01
Female	0.71	0.76	0.43
Unmarried	0.58	0.73	0.02
Has children	0.76	0.75	0.84
Income	700.27	572.33	0.02

Means presented separately for those who had not taken up Medicaid home care (column (1)) and those who had (column (2)). "Difference p-value" tests the equality of means across groups. Income eligible is based upon the income thresholds each state uses to determine eligibility. Restrictive income uses the lowest income limit and applies it to all states to provide an upper bound on takeup. The number of cars is an important determinant of asset eligibility for Medicaid home care. Data from the 1999 NLTCS. Summary statistics by take-up decision are for those who met the "Income eligible, <2 cars" criteria. This leads to a sample of 481 individuals. The alternative to health fair or poor is health good or excellent.

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

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

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

Table 6: Tags analysis. Average formal care consumption, in hours per week, is estimated in the NLTCS. The sample consists of people age 65 and older with at least two activities of daily living limitations. Subsidy rates are constrained to be no smaller than -0.5 (a 50 percent tax) and no greater than 1.5 (a 150 percent subsidy, under which individuals are paid 50 percent of the market price to consume units of formal care). "Corr(marg. utility, tagged pure-cash benefit)" is the correlation between marginal utility in the absence of any policy and the optimal tagged pure-cash benefits.

Appendices

A Theory Appendix

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

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

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

The first-order condition, which holds with equality at an interior optimum, σ^*,[29] is

$$\frac{dEU(\sigma^*)}{d\sigma} = \int_{\Theta} \frac{dv(p(\sigma^*), m(\sigma^*, B); \theta)}{d\sigma} f(\theta) d\theta = E_{\Theta}\left(\lambda(\sigma^*; \theta) \frac{dV(\sigma^*; \theta)}{d\sigma}\right) = 0$$

$$\iff Cov_{\Theta}\left[\lambda(\sigma^*; \theta) x_K(\sigma^*; \theta) p_K^0\right] = (\sigma^* p_K^0) E_{\Theta}\left(\lambda(\sigma^*; \theta)\right) E_{\Theta}\left(\frac{dx_K(\sigma^*; \theta)}{d\sigma}\right). \quad (3)$$

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

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

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

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

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

of Ls.

A.2 First best

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

$$\begin{aligned}\frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma} &= \lambda(\sigma;\theta)\left[x_K(\sigma;\theta)p_K^0 - x_K(\sigma;\theta)p_K^0 - (\sigma p_K^0)\frac{dx_K(\sigma;\theta)}{d\sigma}\right] \\ &= -\lambda(\sigma;\theta)(\sigma p_K^0)\frac{dx_K(\sigma;\theta)}{d\sigma},\end{aligned}$$

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

B Medicaid Home Care and the Cash and Counseling Demonstrations: Additional Background

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 by 2013. In addition to the growing number of participants, the fraction of Medicaid long-term care dollars that go to home care has risen from 18 percent in 1995 to 51 percent in 2014 (Kaiser Commission on Medicaid and the Uninsured, 2016).

An individual's eligibility for Medicaid home care is based upon a financial means tests and an assessment of her "need" for home care based on her health. Medicaid policies vary somewhat across states, at least in part, because Medicaid is financed jointly by the federal and state governments. 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). Those with more restrictive income limits use 100 percent of the SSI amount.

The amount of Medicaid home care for which an individual qualifies is determined by a medical exam. The applicant's health care provider must submit a care plan that details the services deemed appropriate based on the applicant's health status. Summaries of Medicaid-provided home care services are available in LeBlanc et al. (2001) and Kaiser Commission on Medicaid and the Uninsured (2011).

Estimating take-up rates for Medicaid home care, and means-tested programs more generally, is notoriously difficult (U.S. Department of Health and Human Services, 1992; Currie, 2006). Eligibility rules are complex, vary from state-to-state, and often depend upon household characteristics that are unobservable to the researcher. We use the NLTCS to estimate the fraction of the elderly who are eligible for benefits, based on the eligibility criteria from Schneider et al. (1999). We combine the fraction eligible with the size of the 65-and-older population and administrative estimates of the number of Medicaid home care users from LeBlanc et al. (2001). The main source of uncertainty in our estimated take-up rate is the incompleteness of the information on household assets in the NLTCS. In all cases, a person must have at least two activities of daily living limitations. The upper end of our range, 16 percent take-up, should be interpreted as an upper bound on the take-up rate because we imposed (much) more restrictive income and asset requirements than the actual limits in the vast majority of states. In particular, we imposed that the household earned no more than 100 percent of the SSI benefit and had no cars (car value is one of the primary inputs to the asset tests). Our least restrictive eligibility threshold uses the income limits from Schneider et al. (1999) and imposes that the household have fewer than two cars.

B.2 Cash and Counseling Demonstrations

In New Jersey and Florida, only individuals who were currently receiving Medicaid home care were eligible to participate in the demonstrations.[30] Arkansas allowed a limited number of individuals who qualified for but were not receiving Medicaid home care to participate. 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 enrolled beginning in 1998 in Arkansas, in 1999 in New Jersey, and in 2000 in Florida. Individuals who agreed to participate were given a baseline survey and then randomized to the in-kind or near-cash treatments, each with a 50 percent probability.[31] Participants were then surveyed 4-6 months after enrollment and again at 9 months after enrollment. We use data from the baseline and 9 month follow-up surveys.

The near-cash benefit was not the fully cashed-out cost of the individual's care plan. This was an artifact of 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. For example, in New Jersey, 10 percent of the value of the care plan was set aside to cover program costs. Counselors were available to participants to help them develop plans for spending the money, issue checks (to caregivers or for other services), handle paperwork associated with being an employer (e.g. payroll taxes), and maintain account records related to the demonstrations. Recipients had to submit receipts documenting that they spent their benefits on personal care services, though they were allowed to spend up to 10 percent of their allowance on services that could not be readily invoiced, e.g., payments to a neighbor for mowing the lawn.

[30] Our description of the experiments relies heavily upon (Brown et al., 2007).

[31] Those individuals in Arkansas who had not previously been enrolled in Medicaid home care had to verbally commit to seeking the in-kind benefits if they were not randomized to the near-cash benefit.

Appendix Table G.2 provides summary statistics on the Cash and Counseling participants. We restrict the sample to those who are at least 65 years of age and with nonmissing data on age, gender, race, education, and self-rated health. Not surprisingly, those who enrolled in the experiments are somewhat different from the broader population of Medicaid home care users. Comparing Appendix Table G.2 to Table 4, we see that those in the experiments were somewhat younger, used less formal care, and were slightly less likely to be married than Medicaid home care users in the NLTCS. Whether these differences are due to selection into the experiment or differences in the composition of Medicaid home care users across states (the NLTCS includes data from all states, not just the Cash and Counseling states: Arkansas, Florida, and New Jersey). Unfortunately, the sample sizes for Arkansas, Florida, and New Jersey in the NLTCS are not large enough to address this directly (11 observations). The internal and external validity of our analysis is discussed more in Section D.

Our final Cash and Counseling sample includes 1,946 individuals. At baseline, average formal care consumption ranged from 8 (Arkansas) to 16 (New Jersey) hours per week, and on average participants had two informal caregivers. The average age is in the upper 70s, the majority of participants are female, and the participants do not have high levels of education. Although non-negligible fractions of the treatment and control group attrited from the experiment before the nine-month follow-up survey (20 and 35 percent of treatment and control group members, respectively), of the 30 balance tests, none of the differences between treatment and control groups are statistically distinguishable from zero at the 5 percent level and only one is significant at the 10 percent level. This is fewer significant differences than would be expected to arise by chance without any differential attrition.

C Predicting Formal Care Consumption

In principle, insurance contracts could be written that depend on the observable characteristics of the individuals. If these characteristics were predictive of formal care consumption, then the distribution of formal care seen in Figure 2 would overstate the true risk faced by an individual. In this appendix, we investigate the extent to which observable characteristics are able to predict formal care use. We follow the standard approach from the predictive modeling literature: randomly split the sample into two subsamples, train the model on one subsample (the training sample), use the trained model to make predictions for the other subsample (the test sample), and assess predictive power in the test sample (R^2 in our application). We repeat that process 100 times and report the mean and standard deviation of the results. We implement two separate modeling approaches: OLS and machine learning. The advantage of the former is that it is familiar, transparent, and the regression output is easily interpretable; the advantage of the latter is that it likely represents the most advanced methods that insurers could use to predict the risk.

We begin with the NLTCS data and a simple set of variables related to an individual's health. We then add in controls for an individual's informal care options and her income.[32]

[32] Our NLTCS sample has 887 observations. The health variables include the individual's age, her number of ADLs, her self-rated health (from 1 =excellent to 4 =poor), and gender. Controls for informal care options include the number of children and whether the individual is married. The income control is a measure of the individual's total income.

Although reported values of some of these variables are easily manipulable by consumers and so not useful as tags, we include them in our analysis so that we can interpret our results as likely an upper bound on the fraction of variance in formal care use that is predictable.

We focus on the population of people who meet the traditional eligibility requirement for home care benefits, those requiring assistance with two or more ADL limitations. We adjust observed hours of formal care for differences in prices.[33]

Appendix Table G.3 reports the results. The first column presents the average R^2 based on predicting outcomes in the same data used to train the model. This in-sample R^2 corresponds to the R^2 reported in most regression analyses and will tend to overstate the true amount of variation explained due to over-fitting. The second column shows the average out-of-sample R^2. When we only use a set of controls for an individual's health, the in-sample fit is roughly twice as large as the out-of-sample fit (0.075 vs 0.034, respectively). In the final column of the table, we present results for the machine learning model. We implement random forest models that use five-fold cross-validation to tune the model. Dimensions that are optimized include the number of trees, tree depth, minimum leaf size, and number of variables to (randomly) sample at each split.[34] Although the machine learning algorithm allows for more flexible (nonparametric) relationships between the outcome and prediction variables, nonlinearities in the health controls do not appear to be particularly predictive of formal care use.

The next two rows of the table show that adding controls for informal care options and income do not substantively increase the model's predictive power. Neither the OLS nor machine learning model's average out-of-sample R^2 surpasses 0.05. In the next row, we see that treating all of the variables as categorical (e.g. creating separate indicator variables for each year of age) does increases the in-sample OLS R^2 to 0.200, but at the cost of out-of-sample prediction. The overfitting is severe enough that the out-of-sample R^2 becomes negative, -0.105.[35]

The next row uses the continuous versions of the variables, but adds in interactions between those variables and the number of ADLs the individual has. The following two rows use the categorical variables and add in interactions with either the person's number of ADLs or whether she is unmarried. In each case, the model has at best a modest ability to predict out-of-sample outcomes and at worst, has severely overfit the data to the point that the sample average would be provide better predictions out of sample (i.e. the R^2s are negative).

Machine learning can accommodate situations where there are more variables than observations. In the final row of the top panel of Appendix Table G.3, we use machine learning

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

[34]We have also experimented with two different aggregation methods, bagging and boosting. See James et al. (2013) for a discussion of machine learning techniques. These methods are becoming increasingly widespread in prediction problems in economics (Mullainathan and Spiess, 2017).

[35]We do not present results for the machine learning model in this and some subsequent rows because the machine learning algorithm is nonparametric. As a result, the machine learning model will not treat these rows as fundamentally different from the third row.

along with all of the variables related to health, informal care options, and income in the NLTCS (895 variables). The out-of-sample fraction of variance explained increases to 0.114. However, the vast majority of the variance is left unexplained. A quite modest amount of the variation in formal care use appears to be related to observable characteristics which could (potentially) serve as tags.

An alternative, costlier means of targeting is case-by-case examinations such as those commonly used by disability insurance programs. Medicaid home care programs require participants to have a medical exam with a physician or nurses who will create a care plan meant to reflect recipients "need for care. Interestingly, "need" is meant to include the availability of paid and unpaid caregivers (Dale et al., 2004). In principle, it should be as close to a summary measure of demand upon which an insurer could potentially contract as is feasible. In the bottom panel of Appendix Table G.3, we use data from the Arkansas Cash and Counseling experiment to assess how much variation in formal care use can be accounted for by an individuals care plan and demographics.[36] Among those using the in-kind benefit at the nine month follow up survey, care plan hours account for less than 2 percent of the variation in individuals formal care use.[37]

Taken together, it does not appear that much of the variation in formal care use is predictable with the types of observables that insurers would have access to. None of our analyses were able to predict even 12 percent of the variation in formal care use. As a result, it is not likely that even richly-tagged pure-cash transfers would prove as useful in targeting benefits to high marginal utility individuals as in-kind provision.

D 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. The price sensitivity of demand for care is important for other questions as well, including the extent to which private long-term care insurance contracts that subsidize formal care suffer from a "moral hazard tax." In this section, we address issues related to both the internal and external validity of our estimates of the price sensitivity of demand for formal care.

[36] Again, recall that Arkansas is the only state for which we have information on care plans.

[37] Care plans are updated every six months in Arkansas. We use two measures of care plan hours in the analysis. The first is care plan hours at baseline; the second is care plan hours twelve months after baseline. Because care use is measured at nine months, there is some question about differences in measured care plan hours and the individuals true care plan hours at that time. However, the lack of variation in care plans over the twelve month period—the correlation between an individuals care plan at baseline and follow up is 0.86—allays this concern to some degree. The health, informal care options, and demographic variables include self-rated health at baseline, age, gender, whether the participant needed help getting out of bed, bathing, or using the toilet, the number of unpaid caregivers at baseline, whether the participant was married, whether the participant lived alone at baseline, her educational status, and whether she was white.

D.1 Internal validity

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

Supply constraints are thought to have faced Medicaid home care recipients in Arkansas during the period of the Cash and Counseling experiment (Brown et al., 2007). These constraints apparently arose from some combination of Medicaid paying below-market prices and the local home care market being in disequilibrium around the time of the experiment. To the extent that such issues were important, ignoring them would tend to lead us to underestimate the true price sensitivity of demand. The simplest way to avoid this issue is to drop Arkansas from the analysis and instead focus on Florida and New Jersey.

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

Although in principle the combination of maximum benefit limits and care plan limits could limit the quantity of Medicaid home care available to recipients, in practice it does not appear that either one of these constraints significantly constrained consumption. On care plans, many recipients consume strictly less than their care plan hours, and it is not clear what incentive physicians may have to restrict hours. If anything, physicians' professional norms and ethos might lead them to act as an agent of the patient rather than Medicaid. Maximum benefit limits also appear to be less binding than might have been expected. LeBlanc et al. (2001) survey Medicaid home care programs and discuss several explicit mechanisms for granting exceptions to the limits. For example, in New Jersey, where the statutory limit was 25 hours per week, with prior authorization a recipient could receive between 26 and 40 hours of care per week and with central office approval a recipient could receive as much care as "needed." Consistent with these or other mechanisms relaxing quantity limits, the distributions of formal care hours among Cash and Counseling participants receiving traditional Medicaid home care do not exhibit much bunching around these limits. If the limits were binding, one would expect significant bunching because a binding limit causes a convex kink in the budget constraint between formal care and all other goods.[38] Appendix Figures G.1–G.3 present the CDFs of formal care hours for people randomized to the in-kind group

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

in each of the three Cash and Counseling states. In Arkansas (Appendix Figure G.1), there is no apparent bunching that would suggest that consumption was constrained by the state's limit. In addition to there not being a large mass point at 16 hours, nearly one-fifth of the sample consumed more care than the state's limit. In New Jersey (Appendix Figure G.3), there is bunching at certain points in the CDF of care hours, but this appears to be more of a function of rounding than any limits being imposed. The mass points at 15 and 20 hours (8 and 9 percent of the distribution, respectively) are similarly sized to the mass point at the statutory limit of 25 hours (11 percent).

In Appendix Table G.4, we present estimates of the price sensitivity of formal care for each state. The first row shows that the IV Tobit estimates range from -0.96 (Arkansas) to -2.79 (Florida). In the second row, we impose the upper bounds on care hours implied by the Arkansas and New Jersey limits. We censor observations above those cutoffs and use the IV Tobit to re-estimate the price sensitivity. The additional censoring reduces our estimated price sensitivity in Arkansas but increases it in New Jersey. (We exclude Florida since care hours are not limited there.) The differences across states are similar to those found with the standard IV Tobit.

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

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

We test the sensitivity of our results to a number of different distributional assumptions on ε_i. In each case, we continue to instrument for price as we did in the main analysis. These results can be found in Appendix Table G.5. As seen in columns (2) through (4), the estimated price sensitivity changes somewhat from one specification to the next but not dramatically so.

In the next four columns of Appendix Table G.5, we assume that everyone who is potentially at a corner solution has a marginal value of care of exactly p, the maximum consistent with their behavior. As seen in Figure 3, those in the cash group were more likely to consume zero hours of care than those in the in-kind group. In a Tobit model, this greater mass at the censoring point tends to reduce the (latent) mean of the care hours distribution for the cash group relative to the in-kind group. The 2SLS model does not have this feature and, as a result, tends to produce smaller mean differences between the cash and in-kind groups. In our setting, this translates into a smaller price sensitivity. Again, we instrument for the price of care with each participant's randomly assigned transfer type. Under these assumptions, we tend to find a price sensitivity around -1.

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

D.2 External validity

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

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

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

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

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

We use the Arkansas Cash and Counseling experiment to provide additional evidence on the targeting effects of in-kind provision. These results should be interpreted with caution because of likely non-random selection into the experiment. For example, consider the population of disabled, married individuals. Based on the results from the NLTCS, we would expect unmarried people to have greater demand for formal care than a person who is married. Because they chose to join the Cash and Counseling demonstration, the 15-20 percent of married individuals in the experiments (see Appendix Table G.2) are likely to have unusually low demand for formal care. This type of selection complicates targeting comparisons.

Because these levels comparisons are subject to selection bias, we instead pursue a differences-in-differences approach. In particular, we run regressions of the form

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

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 those with more of the characteristic X_i receive differentially greater transfers in the in-kind group than do those with lower values of X_i. For example, if X_i indicates having more disabilities, $\beta_3 > 0$ would indicate that those who are more disabled gain more from the in-kind program than do individuals who are less disabled, i.e., that the in-kind program targets those who are more disabled to a greater extent than the near-cash benefit. Note that the experiment's in-kind provision is being compared to the near-cash transfer, not a pure cash transfer. Because the near-cash transfer is tagged, it likely targets resources to the same set of eligibles targeted by the in-kind transfer. As a result, this analysis likely understates the degree to which in-kind provision targets particular demographic groups relative to a pure cash transfer.

Appendix Table G.6 reports the average effects estimated via OLS as well as impacts in the right tail of the distribution estimated via quantile regression. The right tail of the distribution is of particular importance because that is where there is the greatest scope for targeting to provide insurance value. If in-kind provision targets transfers, then the OLS estimates will reflect an average of the negative effects in one tail with the positive effects in the other tail. The quantile regression, however, will only reflect what is happening in the far right tail and could be more informative about which types of individuals are being targeted. Robust standard errors are reported for the OLS regressions while bootstrapped standard errors are presented for the quantile regressions.

Column (1) of Appendix Table G.6 reports that older and sicker individuals received differentially more transfers in the in-kind program. Measures of self-rated health, gender, and a proxy for availability of informal care do not appear to be associated with the type of transfer provision. Those who lived alone at the baseline receive differentially fewer resources in the in-kind group than those who did not live alone at the baseline. While living alone could signal having fewer informal care options, it could also signal being in better health since the individual is able to live alone. This latter interpretation appears to be more apt in our context because those who lived alone at baseline had lower costs than those who did not live alone ($107 per week for those who lived alone vs. $129 for those who did not live alone).

Columns (2) through (4) show results at the 90th, 95th, and 99th quantiles respectively. The point estimates in the second row suggest that more disabled participants received differentially more benefits from the in-kind transfer than under the near-cash transfer. This difference appears to be growing as we move further out into the tail of the distribution where the targeting benefits from in-kind provision could have their largest impacts. We find similar patterns for women and those with less access to informal care, the unmarried.

Columns (5) through (8) present the same analyses for the subset of participants in Arkansas who had not been in the Medicaid home care program at baseline. This group is more representative of the roughly 90 percent of eligibles who do not take up Medicaid home care. Again, we find that the Cash and Counseling in-kind program appears to target more resources towards those with worse health and fewer informal care options.

F Welfare Analysis: Further Details and Robustness

F.1 Optimal first-best insurance

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

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

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

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

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

F.2 Estimating the distribution of demand for formal care

As discussed in the text, we use the observed distribution of formal care consumption together with our estimate of the price sensitivity of demand for formal care to infer the latent distribution of the level of demand for formal care. We express the level of demand for formal care in terms of satiation points, α. The only tricky part of this calculation is that observed formal care consumption does not point-identify α for people consuming zero formal care, it only bounds it: $\alpha_i \leq \beta p_i$. We estimate the full α distribution, including the α's of people who consume zero formal care, in three steps.

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

Second, we convert each q^* to its corresponding α using the estimated price sensitivity of demand for formal care, $\alpha_i = q_i^*(p) + \beta p$. This adjusts (potentially latent) formal care consumption by our estimate of the impact of the price on consumption. Finally, we estimate the kernel density of the implied α distribution. Figure 6 shows the resulting α distribution. It is mostly just a rightward-shifted version of the observed distribution of formal care consumption, with adjustments for the censoring of people who consume no formal care.

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

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

F.3 State-dependent utility

As discussed in the text, any state-dependence in utility that is correlated with formal care consumption is centrally important for the value of in-kind provision, since it affects the value of redistribution across states with different levels of demand for formal care. State-dependence that increases the 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 the 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. Given the possibility that states with different demands for formal care might have systematically different utility functions, it is therefore important to test the robustness of the results to different possibilities about state-dependent utility.[39]

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

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

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

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

F.4 Robustness

This section provides additional intuition for and discussion of the robustness tests reported in Table 5 and discussed in the main text.

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

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

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

G Appendix Figures and Tables

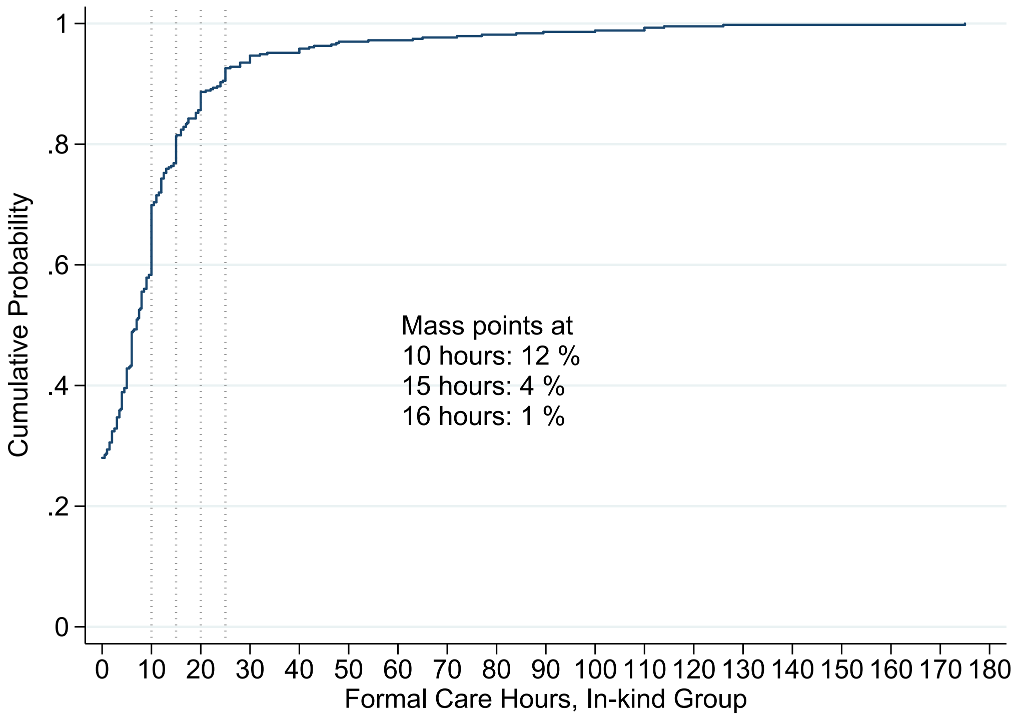

Figure G.1: CDF of Formal Care in Cash and Counseling States, Arkansas

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

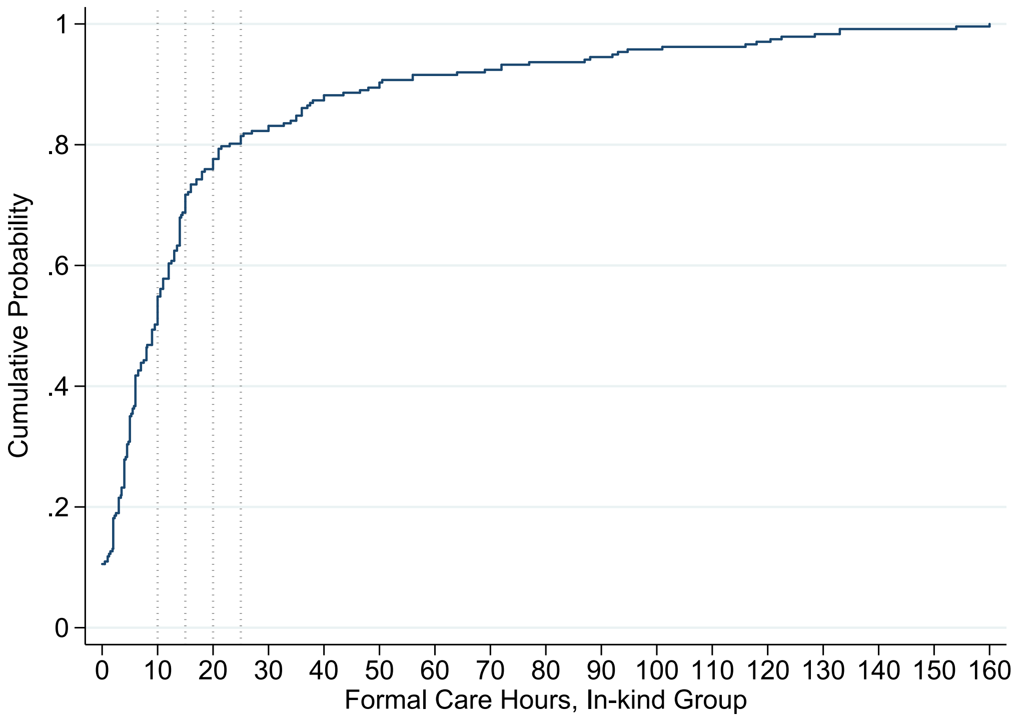

Figure G.2: CDF of Formal Care in Cash and Counseling States, Florida

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

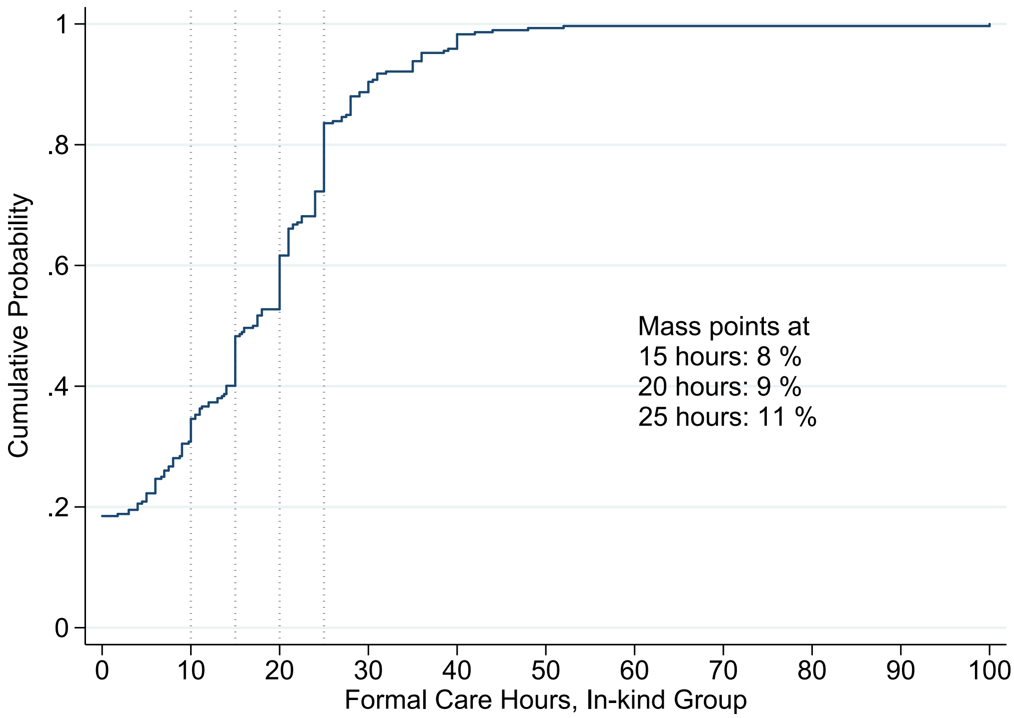

Figure G.3: CDF of Formal Care in Cash and Counseling States, New Jersey

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

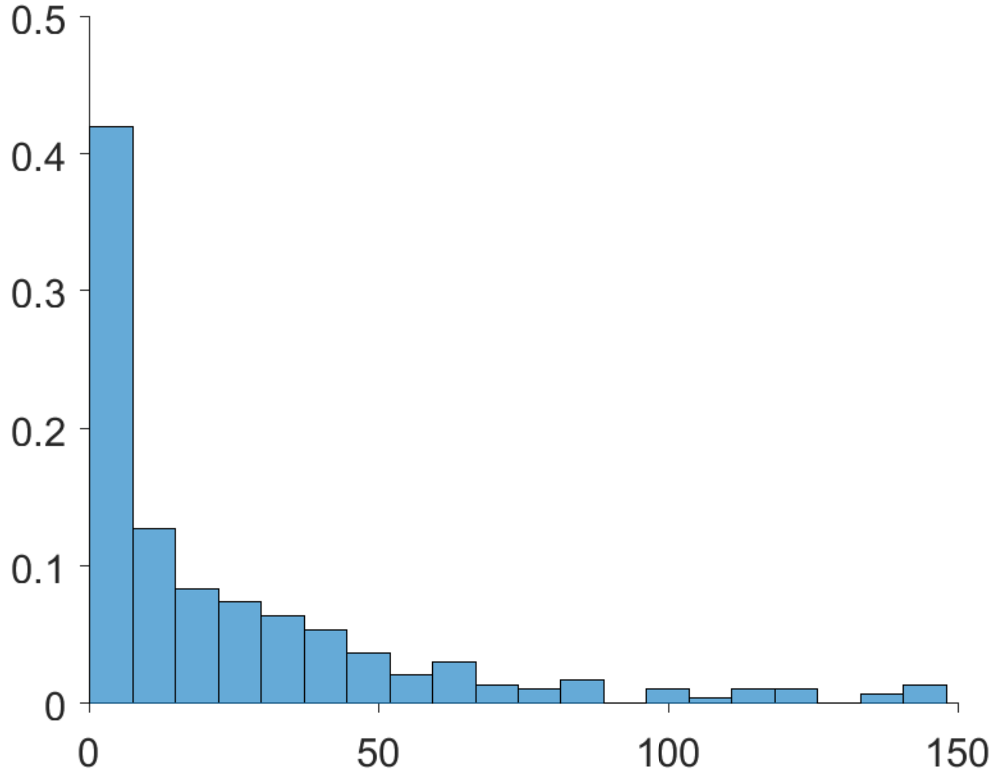

Figure G.4: Distribution of formal care consumption among people with two or more ADL limitations

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

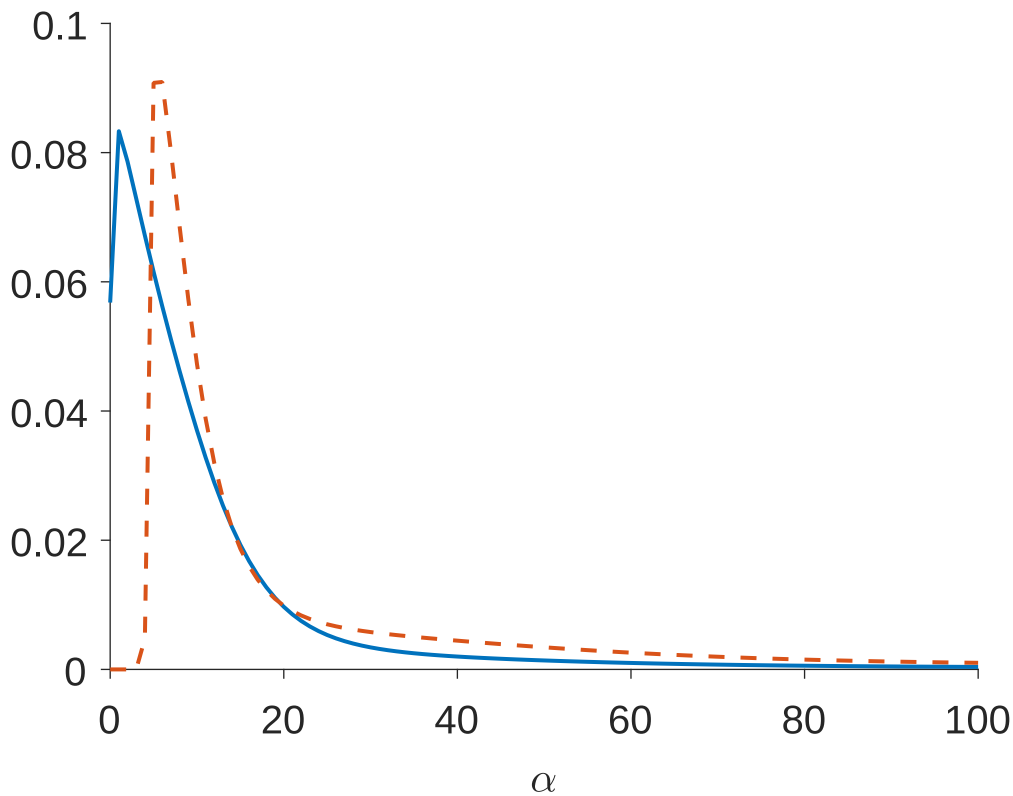

Figure G.5: Distribution of demand for formal care by whether someone lives alone

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

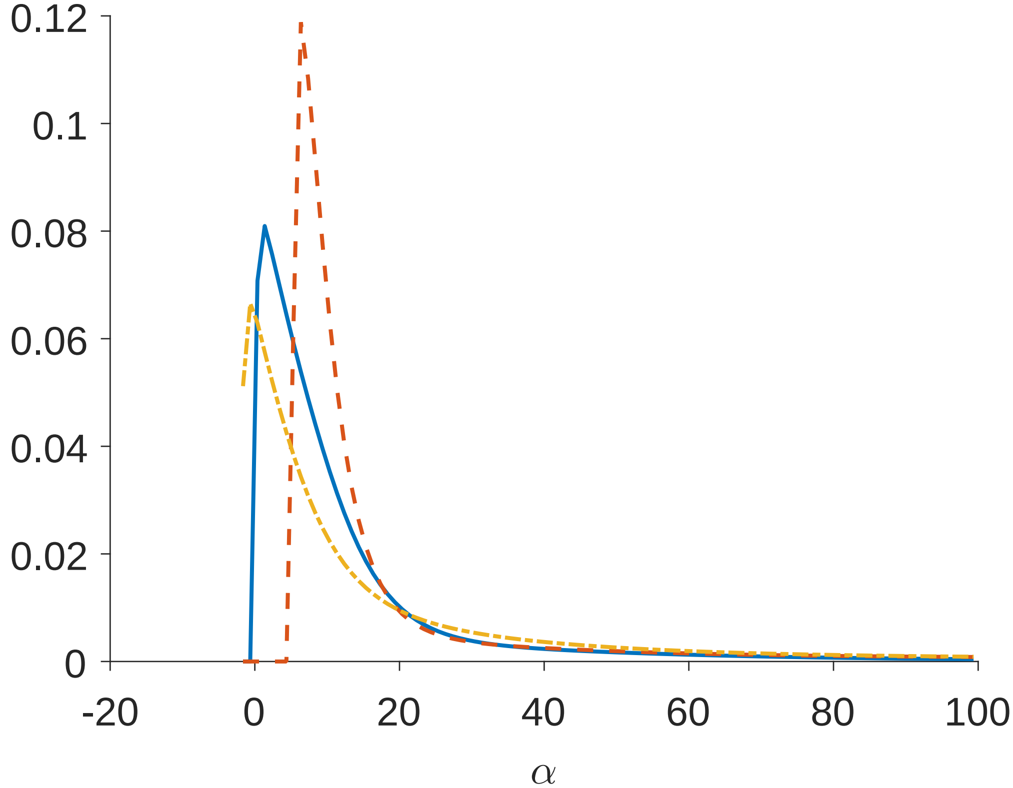

Figure G.6: Distribution of demand for formal care by number of ADL limitations

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

Table G.1: Predicting Price-Adjusted Formal Care Hours in the NLTCS

	(1) OLS	(2) 90th	(3) 95th	(4) 99th
Medicaid home care	9.59*	17.73	40.55	74.76*
	(5.44)	(26.49)	(30.23)	(41.65)
Age	0.40**	0.40	0.37	2.72*
	(0.17)	(0.62)	(0.96)	(1.43)
Four or more ADLs	10.73***	31.73	86.95***	16.75
	(3.25)	(19.49)	(26.58)	(30.43)
If health fair or poor	-1.85	-1.29	-2.20	33.25
	(3.56)	(12.58)	(16.99)	(21.15)
Female	1.26	-1.74	-2.68	-40.59
	(3.94)	(8.40)	(14.46)	(25.23)
Unmarried	13.58***	40.23***	60.88**	84.70***
	(3.45)	(12.06)	(26.20)	(30.44)
Has children	5.40	9.36	4.43	29.89
	(4.71)	(15.40)	(19.22)	(19.01)
Income	-0.00	-0.01	-0.00	-0.02
	(0.00)	(0.01)	(0.01)	(0.02)

Dependent variable is price-adjusted hours of formal care. Sample limited to those who are eligible for Medicaid home care (using "Income eligible, $<$ 2 cars" measure). The sample has 481 observations. Column (1) reports results from an OLS regression; Columns (2) - (4) present results from quantile regressions with the quantile specified in the column heading. Robust standard errors shown in column (1); bootstrapped standard errors shown in remaining columns. * $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table G.2: 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	2.20	2.11	0.30	1.95	2.04	0.45	2.04	2.11	0.59
Age	78.93	79.07	0.76	79.00	79.86	0.18	77.54	77.79	0.65
Male	0.17	0.17	0.92	0.18	0.21	0.45	0.18	0.22	0.15
White	0.62	0.64	0.37	0.67	0.71	0.28	0.50	0.56	0.12
Less than high school degree	0.67	0.66	0.93	0.35	0.38	0.48	0.66	0.65	0.79
High school degree	0.28	0.26	0.48	0.47	0.46	0.97	0.18	0.20	0.54
College degree or more	0.03	0.05	0.10	0.16	0.14	0.48	0.10	0.11	0.68
Health, baseline	3.19	3.22	0.51	3.14	3.06	0.26	3.19	3.16	0.65
Lives alone, baseline	0.32	0.31	0.67	0.25	0.31	0.14	0.33	0.38	0.20
Unmarried	0.85	0.85	0.95	0.77	0.81	0.20	0.79	0.76	0.25
Observations	567	569	.	303	291	.	368	355	.

Means presented by state and type of transfer. P-value is for test that means are the same across the cash and in-kind groups within the state. Formal care hours, Number unpaid caregiveers, Health, and Lives alone are presented for the baseline survey at time of randomization. Remaining variables are measured at the nine-month followup.

Table G.3: Predicting Formal Care Use

	(1) OLS in-sample	(2) OLS out-of-sample	(3) Machine learning out-of-sample
NLTCS:			
Health controls	0.075	0.034	0.030
	(0.016)	(0.033)	(0.025)
Add informal care	0.102	0.046	0.049
	(0.017)	(0.034)	(0.031)
Add income	0.106	0.040	0.054
	(0.018)	(0.036)	(0.035)
Make all categorical except income	0.200	-0.105	
	(0.020)	(0.073)	
Add interactions with # ADLs	0.123	0.010	
	(0.023)	(0.052)	
All categorical, interact with ADLs	0.562	-0.860	
	(0.042)	(0.381)	
All categorical, interact with unmarried	0.270	-0.224	
	(0.029)	(0.134)	
All related variables in dataset			0.115
			(0.035)
Cash and Counseling:			
Care plan, demographics	0.136	-0.046	0.213
	(0.027)	(0.082)	(0.157)

Each entry is the average R^2 from 1,000 sample splits. The standard deviation of the R^2 is presented below the mean for each entry. Rows denote different sets of variables included in the analysis. Columns denote the model used (OLS or machine learning) and whether it is an in-sample or out-of-sample R^2. "Health controls" include age, number of ADLs, self-rated health, and gender. "Add informal care" adds indicators for having children and being married (in addition to health controls). "Add income" adds a control for total income (in addition to health and informal care controls). The following four rows use all health, informal care, and income controls. "Make all categorical except income" creates indicator variables for each value of each variable (except income) and includes those in the regression. "Add interactions with # ADLs" includes full set of variables as well as interactions of each variable with the number of ADLs. "All categorical, interact with ADLs" uses the categorical variables and creates interactions of each variable with the number of ADLs. "All categorical, interact with unmarried" is the same as the previous row but interactions are with variable indicating if person unmarried. "All related variables in dataset" includes additional variables related to health, informal care options, or income. This row can not be estimated with OLS because there are more variables than observations. Row under *Cash and Counseling:* only uses data from the in-kind group in Arkansas. "Care plan, demographics" includes the measure of the individual's care plan at baseline and at twelve months as well as controls for health, informal care, and demographics from the baseline.

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

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

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

Table G.5: The Sensitivity of the Demand for Formal Care to the Composition of Benefits

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

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

Table G.6: Targeting in the Cash and Counseling Experiments, Arkansas

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

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