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

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

January 22, 2018

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 develop a framework for evaluating this tradeoff and apply it to home care. Exploiting large-scale randomized experiments run by three state Medicaid programs, 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 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, Tami Gurley-Calvez, Seema Jayachandran, Brian Melzer, Matt Notowidigdo, Mike Powell, Emmanuel Saez, Diane Schanzenbach, Jesse Shapiro, Courtney Van Houtven, our department colleagues, and many seminar participants for helpful comments. We thank Jose Carreno, Vishal Kamat, Max Rong, and Zeyu Wang for excellent research assistance. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) through the Michigan Retirement Research Center (Grant #5 RRC08098401-05), funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of the Social Security Administration, any agency of the Federal Government, or the Michigan Retirement Research Center.

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). The vast majority of government spending on means-tested welfare programs is on in-kind health benefits,[1] and the recent Affordable Care Act increased such benefits substantially through expanded Medicaid eligibility and subsidies for health insurance. In domestic policy, foreign aid, and charitable giving, 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, medical care, and other goods.

Central to these debates is a tradeoff inherent to in-kind transfers. In-kind provision has a fundamental cost: Recipients would prefer 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 policy, little is known about their relative magnitudes across a wide range of important contexts.

In this paper, we develop a 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 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: health care. In the U.S. in 2013, government spending on Medicaid home care was $57 billion (Ng et al., 2016), and government spending on in-kind health-benefit programs as a whole was $1.1 trillion, more than six percent of GDP (Centers for Medicare

[1] In 2014, government spending on means-tested in-kind health benefits through Medicaid was $497 billion, over six times greater than government spending on any other welfare program, including the Supplemental Nutrition Assistance Program ($76 billion), the Earned Income Tax Credit ($68 billion), and Temporary Assistance for Needy Families ($32 billion).

and Medicaid Services, 2017). 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. 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 observation that an in-kind transfer has the same effect on a recipient's budget set as a (potentially non-linear) price subsidy, highlights three key determinants of the welfare consequences of in-kind provision.[2] The first 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. Using nationally representative data, we find that the demand for formal care is highly heterogeneous within benefit-eligible states. While 65 percent of people in benefit-eligible states do not consume any formal care, among those who do consume formal care, there is a long right tail. An individual at the 95th percentile receives 168 hours per week, which at the average hourly price of $15 (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 Medicaid "care plans" based on individual medical exams. This suggests that even extensively-"tagged" cash benefits (Akerlof, 1978) would leave significant risk uninsured.

The 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, nearly twice average consumption in the benefit-eligible population. This implies 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]The fundamental feature of in-kind transfers that we focus on is that they reduce the recipient's cost of consuming the good over some range of quantities. This feature is shared by a wide range of other policies, including vouchers, conditional cash transfers, benefit programs with ordeals (the benefit effectively subsidizes consumption of the ordeal), insurance policies with non-unitary "coinsurance rates," and commodity taxes and subsidies. The key tradeoff we analyze is likely central to the welfare consequences of these other types of policies as well.

The 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 can help 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 among those randomized to the in-kind benefit, and 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. 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; Braun et al., 2017) and implicit taxation 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. 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). But that even richly-tagged cash benefits would leave most of the risk uninsured means that providing home care in kind, although costly, might be the best way to insure the risk from chronic health 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.[3]

Our approach helps link the theoretical and empirical literatures on in-kind transfers, which, as Currie and Gahvari (2008) emphasize in their review, have been largely disconnected so far. The theoretical branch has investigated 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). Much of the empirical branch has focused on estimating the effects of in-kind provision on consumption.[4] Noting the equivalence of the effects of in-kind transfers and subsidies on recipients' choice sets allows us to utilize the well-developed theoretical and quantitative approaches for analyzing taxes and subsidies from the vast literature on optimal taxation following Mirrlees (1971). 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, health care.

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).[5] A key finding in this literature 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 an important 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

[3]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, e.g., Ettner, 1995; Van Houtven et al., 2013; Skira, 2015) and may alleviate the Samaritan's Dilemma (since informal care provision, by reducing labor supply and worsening health (Coe and Van Houtven, 2009; 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.

[4]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 food spending. Other work examines whether in-kind transfers reduce local prices (Cunha et al., 2011), are effective at self-targeting (Reeder, 1985; Cutler and Gruber, 1996; Jacoby, 1997), and reduce measured poverty rates (Smeeding, 1977). Our approach complements those of Finkelstein et al. (2015), who analyze the welfare effects of Medicaid health insurance coverage for prime-age adults.

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

are a key determinant of the welfare consequences of alternative benefit designs in contexts, like health care, in which demand for the benefit is highly heterogeneous.

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. The key feature of in-kind provision on which we focus, and which is shared by many other policies, 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. 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 of the world. But with small adjustments, the framework can be applied to questions of redistribution across different types of people as well.[6]

An in-kind benefit has the same effect on a recipient's budget set 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 recipient's budget set as a non-linear price subsidy of 100 percent on units up to the benefit limit and 0 percent on units above the limit. We focus on the case of a subsidy program with no quantity limit. We do this both for simplicity of exposition and because many Medicaid home care programs, including those in the Cash and Counseling experiments, do not appear to have binding benefit limits in practice. Appendix Section A.1.1 considers the case with a binding benefit limit.

The key considerations for in-kind provision can be seen in Figure 1. Figure 1 shows the

[6]To focus on the problem of insuring a risk, we ignore any second-best considerations that might arise from the interaction between the program and other distortions in the economy. The problem can therefore be viewed as that of a private insurer, which would not account for such effects, or of a government that can condition the insurance benefit on ability type. In the case of a government that cannot condition on ability, the optimal in-kind benefit in a government program would depend on the joint distribution of ability and demand for the good and on any effect of in-kind provision on tax system efficiency (see Atkinson and Stiglitz, 1976; Saez, 2002; Kaplow et al., 2008). We discuss how such considerations might affect optimal government home care benefits in Section 6.5.

value (equivalent variation) and efficiency cost of a price subsidy on a particular good in each of two states of the world with different levels of demand for the good. The subsidy has a moral hazard cost: It is worth less in each state than it costs the government or insurance company to provide due to the induced change in consumption, which is increasing in the compensated own-price elasticity of demand. The subsidy also has a targeting effect: It is worth more in states of the world in which demand for the good is greater. Relative to a cost-equivalent cash benefit, the subsidy redistributes toward states in which demand for the good is greater from states in which demand is smaller.

2.1 The benefit program and its budget constraint

An individual faces a risk that potentially affects prices, income, and preferences. The state of the world, θ, is drawn from the distribution, $G(\theta)$, with density $g(\theta)$. The planner knows $G(\theta)$ and how θ affects prices, income, and preferences but cannot verify which state has occurred ex-post.

Consider an idealized in-kind benefit program that potentially combines two elements: a cash benefit, b, and a linear subsidy, σ, on good K, x_K. The cash benefit and subsidy rate are common across all states of the world and are automatic in the sense that there are no take-up decisions; the individual receives the cash benefit and the subsidy on purchases of good K in all states. Expected program spending, B, is comprised of the cash benefit and spending on the subsidy on good K:

$$b + \sigma p_K^0 \int_\Theta x_K(\sigma, B; \theta) g(\theta) d\theta = B,$$

where p_K^0 is the (constant) before-subsidy (sellers') price of good K and $x_K(\sigma, B; \theta)$ is consumption of K in state θ as a function of the policy. For simplicity, we assume that the supply of every good is perfectly elastic. In this case, an increase in the subsidy reduces the individual's after-subsidy price of good K one-for-one (no incidence on supply), $p_K(\sigma) = (1-\sigma)p_K^0$, and has no effect on the prices of any other goods, $p_i(\sigma) = p_i^0$ for $i \neq X$, where p_i^0 is the price of good i without any benefit program.

Two special cases of this combined cash-plus-subsidy program are a pure cash-benefit program ($b = B, \sigma = 0$) and a pure subsidy program with no cash benefit ($b = 0, \sigma > 0$). A pure in-kind benefit program has a zero cash component and a full subsidy, ($b = 0, \sigma = 1$).

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

A budget-neutral shift toward in-kind benefits involves increasing the subsidy rate, σ, and at the same time decreasing the cash benefit in order to maintain the same program budget. The change in the cash benefit that maintains the same program budget in response to a

marginal increase in the subsidy rate is

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

The cash benefit falls by the increase in expected spending on the in-kind benefit (subsidy). Expected 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 good K in each state, $p_K^0 E_\Theta\left(x_K(\sigma,B;\theta)\right)$ ("mechanical effect"); and (ii) the increase in spending on the subsidy due to the induced change in consumption of good K in response to the shift in program benefits, $\sigma p_K^0 E_\Theta\left(\frac{\partial x_K(\sigma,B;\theta)}{\partial \sigma}\right)$ ("behavioral effect").[7]

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

In each state, the net ex-post value of the shift toward in-kind provision is the benefit of the increase in the subsidy on good 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 its after-subsidy price by p_K^0, $\frac{dp_K(\sigma)}{d\sigma}=-p_K^0$. The value, in units of the numeraire, of this reduction in the price of K to an individual of type θ is, by the envelope theorem (Roy's identity), proportional to consumption of K,

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

where $v(p,m;\theta)$ is the indirect utility function of an individual in state θ and $m(\sigma,B;\theta)=m_0(\theta)+b(\sigma,B)$ is benefit-inclusive income in state θ. 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, in units of the numeraire, of a budget-neutral marginal shift toward in-kind benefits of

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

[7]The "behavioral effect" includes income effects from the change in the cash benefit as well as income and substitution effects from the reduction in the net-of-subsidy price. Appendix Section A.1.2 analyzes the budget-neutral shift toward in-kind benefits in terms of income and substitution effects.

The marginal net value for an individual in 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), $p_K^0 \left[x_K(\sigma, B; \theta) - E_\Theta \left(x_K(\sigma, B; \theta)\right)\right]$, 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 (moral hazard cost), $\sigma p_K^0 E_\Theta \left(\frac{\partial x_K(\sigma, B; \theta)}{\partial \sigma}\right)$.

Equation 1 shows that the shift toward in-kind provision has two key effects. It redistributes toward states with above-average demand for the good, and it distorts consumption of the good. The extent to which the individual gains ex post from a marginal shift toward greater in-kind provision is increasing in her level of demand for the good in that state and is 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

$$EU(\sigma, B) = \int_\Theta v\left(p(\sigma; \theta), m(\sigma, B; \theta); \theta\right) g(\theta) d\theta.$$

The partial derivative of expected utility with respect to the in-kind component σ, adjusting the cash component b to hold fixed total program spending B, is

$$\begin{aligned}\frac{\partial EU(\sigma, B)}{\partial \sigma} &= \int_\Theta \frac{dv(p(\sigma; \theta), m(\sigma, B; \theta); \theta)}{d\sigma} g(\theta) d\theta = \int_\Theta \lambda(\sigma, B; \theta) \frac{\partial V(\sigma, B; \theta)}{\partial \sigma} g(\theta) d\theta \\ &= p_K^0 Cov_\Theta \left[\lambda(\sigma, B; \theta),\ x_K(\sigma, B; \theta)\right] - \sigma p_K^0 E_\Theta[\lambda(\sigma, B; \theta)] E_\Theta \left(\frac{\partial x_K(\sigma, B; \theta)}{\partial \sigma}\right),\end{aligned}$$

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

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, the variance in marginal utility, and the correlation between demand for K and marginal utility—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.[8]

[8] Appendix Section A.1.3 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, $(b = B, \sigma = 0)$. Absent moral hazard, the optimal policy eliminates the covariance between marginal utility and the demand for K, since the in-kind benefit redistributes across states with different levels of demand for K at no efficiency cost. If

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. Roughly 15 percent of Americans over age 50 have at least one person helping them perform activities of daily living (ADL) such as bathing, eating, and dressing (Barczyk and Kredler, 2017). 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, 2017). Spending on formal home care was $88 billion in 2015 (Centers for Medicare and Medicaid Services, 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). 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, about $17,000 per beneficiary. This is about half of Medicaid's total spending on long-term care and about two-thirds of all spending on formal home care. Eligibility for Medicaid home care is determined by financial- and health-related criteria. An individual must have sufficiently low income and assets and must have at least two ADL limitations that are expected to last at least 90 days. The traditional Medicaid home care benefit is an in-kind benefit of formal home care from a Medicaid-approved agency.[9] 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 recipients 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).

the demand for K is heterogeneous across states and at least somewhat elastic, the optimal policy trades off the insurance benefit of increasing in-kind provision against the moral hazard 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 moral hazard cost and no targeting benefit. Appendix Section A.1.4 analyzes the first-best case in which θ is verifiable.

[9]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. Appendix Section A.2 discusses evidence on this, including Appendix Figures A.1–A.3 of the distributions of formal care consumption among Medicaid home care recipients in the three Cash and Counseling states. It also provides additional information about Medicaid home care.

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.[10] The main goal of the experiments was to test whether recipients could effectively manage their near-cash benefits and receive "enough" care. The results were almost uniformly positive. Members of the cash-benefit treatment group 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. This cost 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 (see Appendix Table A.2 for summary statistics), indicate that the demand for formal home care is highly heterogeneous within the eligible population. The population traditionally eligible for home care benefits is non-institutionalized people with two or more ADL limitations; the population eligible for Medicaid home care is the subset with low enough income and assets. Figure 2 shows the distribution of formal care consumption among non-institutionalized people aged 65 and older with two or more ADL limitations. Even within this group of people with severe chronic health problems, there is significant heterogeneity in demand for formal care. 65 percent do not consume any formal care, and among those who do there is a long right tail. Conditional on consuming any care, median consumption is 12 hours per week ($9,360 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

[10] Appendix A.2 contains more information about the Cash and Counseling experiments, including summary statistics of Cash and Counseling participants and balance tests which provide evidence of a valid randomization (Appendix Table A.1) as well as comparisons of Cash and Counseling participants to the broader populations of people eligible for home care benefits and people who take up Medicaid home care (Appendix Table A.2). The near-cash benefit was a cash budget that had to be spent on personal care services. This requirement seems unlikely to have been 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 the near-cash benefit could revert to the standard in-kind benefit at any time; those randomized to the in-kind benefit could not receive the near-cash benefit.

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.

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. Table 1 reports the fraction of the variation in formal care consumption among the benefit-eligible population and among Cash and Counseling participants that can be "explained" by observable characteristics. Even extensive sets of individual- and household-level variables—including information about demographics, health, insurance, health care consumption, potential and actual informal caregivers, living arrangements, and Medicaid care plans based on individual medical exams—never "explain" more than about 15 percent of the variation in formal care consumption out of sample. (See Appendix Table A.3 for more information about the variables and Appendix Section A.3 for details of the analysis.[11]) This suggests that extensively-tagged cash transfers would leave much of the heterogeneity in formal care and non-care consumption uninsured. The scope for tagging appears limited in this context, even ignoring any verification and moral hazard costs that would be involved.

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—high-cost states cannot be targeted directly. Such hard-to-verify heterogeneity in the costs of bad health introduces a potential targeting rationale for in-kind provision.

[11]Medicaid home care programs require recipients to complete a medical exam with a doctor or nurse, who creates a care plan meant to reflect the recipient's "need" for care. "Need" is meant to incorporate not just health problems but the availability of paid and unpaid caregivers as well (Dale et al., 2004). In principle, such case-by-case examinations by experts, which are commonly used in disability insurance programs, could be as close to a summary measure of demand or marginal utility on which an insurer might feasibly condition benefits. In practice, we find that Medicaid care plans have little power for predicting formal care consumption.

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

The welfare effects of in-kind provision depend on the sensitivity of demand to the composition of benefits, which, as shown in Section 2, 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.[12] The Cash and Counseling experiments have two major advantages. First, the randomization solves an especially difficult simultaneity problem: Many factors that shift the supply of formal care are also likely to shift the demand for formal care by changing the opportunity cost of informal care.[13] Second, the variation in the price of formal care spans the full range most relevant for policy, from zero to the market price.

The experimental results provide strong evidence that in-kind provision of home care has a large moral hazard cost. Table 2 shows that being randomized to in-kind benefits doubles average consumption of formal care from 7 to 14 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 24 to 55 percent) and increasing 95th-percentile consumption by 16 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 assigned to the near-cash benefit reverted to the traditional in-kind benefit and some people left Medicaid home care altogether—by using the randomized assignment as an instrument

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

[13] 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 likely shift the supply of formal care. But at the same time, changes in the minimum wage also likely change the opportunity cost of informal care-giving by changing the wage or employment prospects of some potential informal care-givers. This likely shifts the demand for formal care since formal and informal care are closely-related goods.

for the price each participant faced.[14] We estimate the system

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

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

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

where p_i is the price of formal care, $Cash_i$ is an indicator of whether the participant was randomized to the near-cash treatment, and X_i includes indicators for sex, education level, race, self-rated health at baseline, living alone at baseline, five-year age bins, and state. The key parameter of interest is β, the effect on formal care consumption of an increase in its net-of-subsidy price. 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 income.[15] 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 relationship, reported in Appendix Table A.4, is economically and statistically large: Being assigned to the in-kind benefit decreases the average price of formal care by approximately \$7.60, and the F-statistic exceeds 1,100. The instrumental variables estimate of β, presented in Table 3, 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 of -1.7. The conclusion that the demand for formal care is highly sensitive to its price holds in each of the three states and is robust to a wide range of alternative assumptions about the distribution of the error terms (Appendix Table A.5) and benefit limits (Appendix Table A.6). See Appendix Section A.4 for details and a discussion of the generalizability of the results to other populations and policies of interest.

The estimates imply that in-kind provision has a large moral hazard cost. Someone consuming 14 hours per week of formal care, the average among participants randomized to the

[14]Participants who receive the near-cash benefit or who leave Medicaid home care face the market price in their state. Participants who receive the in-kind benefit face a price of zero. In principle, care plans or maximum benefit rules could limit the amount of formal care that those receiving the in-kind benefit could consume free of charge and thereby raise the shadow price of formal care above zero. In practice, recipients of the traditional in-kind benefit, but not the near-cash benefit, had opportunities to increase their benefit amount, and a variety of evidence suggests that they were able to consume as much care as they wished free of charge. See Appendix A.2 for additional details and evidence. Appendix Section A.4 tests the robustness of the results to different assumptions about how binding care plans and maximum benefit rules might be.

[15]With non-zero income effects of demand for formal care, this parameter is sufficient for analyzing the Cash and Counseling experiments but not policies with different cash benefits. The Cash and Counseling experiments roughly held fixed Medicaid's spending on each participant of the experiments, a group whose average cost is much greater than that of the eligible population as a whole. 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 in the entire eligible population. With positive income effects of demand for formal care, our estimates will tend to understate the true moral hazard costs of these other policies of interest.

in-kind benefit, for example, would not consume any formal care without the subsidy and values the care she does receive at just 28 percent of its cost. This implication is 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 Counseling demonstrations. Better targeting, as shown in Section 2, means a greater covariance between benefits and marginal utility. 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, which, in many models, means higher marginal utility.[16]

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 who take up benefits (see Appendix A.2 for details). We estimate that only 4 to 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

[16] 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.

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 level of demand for formal care (by 15 hours per week, accounting for price differences), are sicker (64 vs. 47 percent have four or more ADL limitations), and appear to have worse informal care options (64 vs. 56 percent are unmarried and 37 vs. 27 percent live alone). Even conditional on the personal and family characteristics in Table 4, those who take up Medicaid home care have much greater demand for formal care, both on average and throughout the right tail of the distribution (see Appendix Table A.7). This suggests that the (self-)targeting of Medicaid take-up decisions might be hard to replicate even with tagged cash transfers that are conditioned on health and informal care-related characteristics. 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 lead to "perverse targeting" in which those most desperate for help are least likely to receive it, in practice take up of Medicaid home care is strongly increasing in proxies for marginal utility.

Unlike cash benefits, in-kind benefits can have important targeting effects on the intensive margin among those who take up benefits as well. We explore this possibility using the Cash and Counseling experiments. Whereas decisions about whether to take up Medicaid home care likely depend on a variety of factors in addition to the flexibility of the benefit, Cash and Counseling's experimental design isolates the effect of in-kind provision. It allows us to compare the distribution of benefits among those randomized to the in-kind benefit to both a counterfactual distribution of benefits that would have arisen under a hypothetical uniform cash benefit as well as the distribution of benefits among those randomized to the experimental near-cash benefit. Because the experimental near-cash benefits were based on individual medical exams, in principle they could be much more targeted than a hypothetical uniform cash benefit would be.

Figure 4a shows kernel density plots of the dollar cost of benefits received by participants randomized to each transfer type.[17] In-kind provision concentrates benefits substantially. The variance in benefits 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 receive no benefit is over three times larger in the in-kind group, 31 percent vs. 10 percent. The 99th percentile benefit of the near-cash group is exceeded by 17 percent of the benefits of the in-kind group. Figure 4b

[17]To isolate differences in the concentration of benefits, we scale up the benefits of the near-cash group, whose benefits were slightly smaller on average, to have the same mean as the benefits of the in-kind group. This works against our findings about benefits being more concentrated among those randomized to the in-kind benefit.

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 cash transfer equal to the per-participant average benefit in the in-kind group. In-kind provision significantly concentrates benefits on the intensive margin, even compared to Cash and Counseling's tagged near-cash transfer. 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 benefits 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, leaving those who consume more formal care fewer resources for non-care consumption. Appendix Table A.8, discussed in Appendix Section A.5, provides suggestive evidence that in-kind provision concentrates benefits on recipients who are sicker and have worse informal care options than the average recipient as well.

Taken as a whole, these results show that in-kind provision sharply concentrates benefits on a small fraction 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. When combined with the evidence that the potential for targeting these states directly using tagged cash benefits is quite limited, the targeting effects of in-kind provision appear 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 Benefit Versus Moral Hazard Cost

This section uses a stylized expected utility model to investigate the net welfare effects of the targeting benefit and moral hazard cost of the in-kind provision of home care benefits. The general approach is similar to those of the literatures on health spending risk and on optimal taxation, with small adjustments to match the home care setting.

6.1 Model

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

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

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

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

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

This utility function is motivated by key evidence from our setting. It produces a simple function for the demand for formal care that is consistent with the sensitivity of formal care consumption to its price, with many people who need assistance not consuming any formal care, and with people becoming satiated at finite levels of formal care consumption.[18] 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 (see Appendix A.6.1). It has an intuitive interpretation: Utility is decreasing in any unmet, residual health needs, $(\theta - F)$, the size of which is decreasing in formal care consumption, F, and increasing in the level of demand for formal care, θ. 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.

We analyze idealized mixed in-kind/cash-benefit policies with a linear subsidy rate σ and a cash benefit b. Take up is automatic and there are no participation costs. Following standard practice for Medicaid home care and private long-term care insurance, we focus on programs

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

that limit eligibility to people with two or more ADL limitations. For each candidate subsidy rate $\hat{\sigma}$, we find the cash benefit $b(\hat{\sigma}, B_{IK})$ that holds fixed expected program spending at the expected spending on a pure in-kind benefit (a 100 percent subsidy with no cash benefit), B_{IK}. Policies with smaller subsidy rates have larger cash benefits. We measure the welfare effect of a policy $\hat{\sigma}$ as its ex ante equivalent variation, $EV(\hat{\sigma})$, the extra income the individual would have to receive to make her as well off, in expected utility, $EU(\sigma, B)$, under a budget-equivalent pure-cash policy as she is under the policy in question:

$$EU(\sigma = 0, b = B_{IK} + EV(\hat{\sigma})) = EU(\sigma = \hat{\sigma}, b = b(\hat{\sigma}, B_{IK})).^{19}$$

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 states of the world in which the individual is eligible for home care benefits, $G(\theta)$. For β, we use our main estimate from the Cash and Counseling experiment. For $G(\theta)$, we use our estimate of β to convert the observed joint distribution of formal care consumption and formal care prices in the NLTCS into a distribution of the level of demand for formal care, $G(\theta)$.[20] For the main analysis, which takes as given the standard eligibility criteria and conditions on the individual being eligible for home care benefits ex post, our sample is non-institutionalized individuals aged 65 and older with two or more ADL limitations. For the tags analysis, we estimate separate $G(\theta)$ distributions for different sub-samples defined by their tagged characteristics (e.g., for people with different numbers of ADL limitations).

Estimating $G(\theta)$ would be entirely straightforward were it not for people who consume no formal care when facing a positive price. For such people, 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. Since 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. As a baseline, 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 θ values of people who consume no formal care when facing a positive

[19] Expected utility is $EU(\sigma, B) = \int_{\Theta} \max\{\bar{u}, v(p(\sigma), m + b(\sigma, B); \theta)\} f(\theta) d\theta$, where $\bar{u}$ is the utility floor guaranteed by separate means-tested programs, whether government programs or private charity. That we do not include spending by the means-tested consumption floor program in the budget of the home care program tends to reduce the value of insurance and the targeting benefit of in-kind provision, since part of what insurance does is displace spending by the means-tested program. In the context of long-term care, this implicit taxation of private insurance by government means-tested programs is quite large (Brown and Finkelstein, 2008).

[20] This follows the common practices of using the observed cross-sectional distribution to proxy for the (unobservable) counterfactual distribution facing any one individual (e.g., Finkelstein et al., 2015) and of treating all ex-post heterogeneity as the outcome of an exogenous process (e.g., the vast majority of the large literature on optimal taxation; see Keane, 2011, for a review).

price. See Appendix A.6.2 for details.

Figure 6 presents our main estimate of the density of the level of demand for formal care, $g(\theta)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits a long right tail (the median is 8.1 hours per week, the mean is 15.8, and the 90th percentile is about 48) and that most of the mass is at low values (about 59 percent percent of θ values are less than 10 hours per week).

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

With these parameters, the risk within the set of states of the world traditionally eligible for home care benefits 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 exceed the expected cost of the first-best program by \$9,377 (136 percent).

6.3 Welfare effects of in-kind provision

Figure 7 shows the equivalent variation of the mixed in-kind and cash benefit program as a function of the in-kind component, the subsidy rate σ. The optimal subsidy rate is 88 percent, close to that of a pure in-kind program, and it 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 program under an alternative pure-cash benefit program, the cash benefit would have to exceed the expected cost of the optimal program by \$5,528 (80 percent). Though not optimal, a pure in-kind benefit program with a 100 percent subsidy and no cash benefit also improves substantially on the pure-cash program. In order to make the individual as well off as she is with the pure in-kind benefit program under an alternative pure-cash benefit program, the cash benefit would have to exceed the expected cost of the in-kind program by \$4,421 (64 percent).

Table 5 shows clearly the key tradeoff of in-kind provision. The optimal subsidy reduces expected non-care consumption due both to moral hazard and, to a lesser extent, to crowd-

ing out transfers from the consumption floor. The moral hazard cost is large: Formal care consumption is 2.4 times greater than it is in the absence of the program, and the expected ex-post equivalent variation of the optimal benefit is only 48 percent of its cost. Despite this, it is optimal to subsidize formal care at a high rate because doing so provides valuable insurance. The standard deviation of annual non-care consumption is 4.5 times greater under the pure-cash program than under the optimal program, \$5,610 versus \$1,237, and the optimal program is quite effective in targeting states of the world with greater marginal utility. The correlation between an individual's marginal utility in the absence of any program and her ex-post equivalent variation of benefits under the optimal program is 0.84. The gain from in-kind provision comes from the large transfers it makes to the rarely-occurring states with the greatest demand for care and lowest non-care consumption. Ex post the individual values the optimal benefit at least as much as the cash benefit of the cost-equivalent pure-cash program only 16 percent of the time. This might help explain why many countries and U.S. states have made home care benefits more cash-like: Making benefits more cash-like helps most recipients 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 smaller benefits in states of the world with high demand for formal care, which likely worsens ex-ante insurance and welfare.

The other columns of the table show how different assumptions about the key ingredients of the model affect the results. The conclusions are highly robust to changes in the price sensitivity of demand, β, the distribution of the level of demand, $G(\theta)$, and state-dependence in the utility function.[21] The only plausible specification in which the optimal subsidy is not large is one that combines relatively low risk aversion together with a relatively generous consumption floor. But this reflects the fundamental undesirability of *any* insurance—including a first-best contract—in situations in which means-tested programs are sufficiently attractive, not any undesirability of in-kind provision per se (see Brown and Finkelstein, 2008, for a related result). A government program that internalized spending by the providers of the consumption floor would wish to subsidize formal care even in this case. These columns also shed light on 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 and in the generosity of alternative insurance arrangements, such as any consumption floor or means-tested programs. It is increasing in risk aversion and in the extent to which any state-dependence in utility increases marginal utility in states with greater demand for formal care.

Although the optimal policy significantly improves upon a pure-cash policy, it achieves only 59 percent of the incremental value over a pure-cash benefit that the hypothetical first-

[21] Appendix A.6.3 has more information about the analysis of state-dependent utility, and Appendix A.6.4 discusses these and other results in more detail. The results are robust to state dependence in utility greater than the most relevant estimates, by Finkelstein et al. (2013), likely imply.

best policy does. This 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 more extensive tagging

This section extends the analysis to the case in which certain groups of states of the world can be verifiably distinguished and so offered different benefits. We estimate the gains from catering benefits to different groups of states of the world defined by whether the individual lives alone or by the number of ADL limitations the individual has (2–4, 5, and 6), the two strongest predictors of formal care consumption uncovered in Appendix Table A.9.[22] The procedure is the same as that in the last section, except that we estimate different θ distributions for each verifiably distinguishable group of states and allow the program to offer a different benefit to each group. Appendix Figures A.5 and A.6 show the θ distributions by whether the individual lives alone and by the number of the individual's ADL limitations, respectively. Appendix Section A.6.2 provides additional details.

The ex-ante welfare gains from using tags to target high-marginal utility states, reported in Table 6, are quite small. The 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 gain from optimally tagging benefits based on the number of ADL limitations someone has is smaller still. The fundamental reason for tags' ineffectiveness in insuring this risk is that much of the heterogeneity in demand for formal care occurs within rather than across states that can be distinguished on the basis of 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 ADL limitations" tag. These results are similar to those of Mankiw and Weinzierl (2010), who use height as a tag for income taxation.[23] Better data would presumably enable an insurer to improve at least somewhat on these results. But the small gains from tagging with two of the strongest predictors of formal care consumption and the limited ability, as discussed in Section 3, of even extensive sets of characteristics to predict formal care consumption suggest that the scope for tagging home care benefits is quite limited. In-kind benefits appear to play an important role in targeting benefits to high-marginal utility states of the world.

[22] We are limited in the number of partitions into which we can divide the state space by the size of the NLTCS sample, given the data-hungry non-parametric estimation of the distribution of demand for formal care. We chose the partitions to maximize the across-partition heterogeneity in the demand for formal care.

[23] 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."

6.5 Likely effects of omitted factors

This analysis is highly stylized and leaves out several potentially relevant factors. It does not include any potential benefits of in-kind provision other than targeting, whereas in-kind provision of home care likely improves tax system efficiency, alleviates the Samaritan's dilemma, and has paternalistic benefits.[24] Nor 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 costs and benefits outside our analysis would be to increase the relative attractiveness of in-kind provision.[25] The analysis abstracts from take-up decisions and costs. Patterns of take-up decisions among those eligible for Medicaid home care appear to be consistent with "good targeting" on average, as discussed in Section 5, but there are surely take-up and non-take-up choices that are ex-ante inefficient. The analysis focuses on home care and does not explicitly model substitution with nursing home care. This was done for simplicity given that such substitution appears to be small (Kemper, 1988; Grabowski and Gruber, 2007). Finally, it ignores any effects on the welfare of potential care-givers. Whether the actual and potential suppliers of informal care would prefer that home care benefits be provided in kind or in cash depends on the fundamental determinants of informal care (e.g., the relative importance of feelings of altruism and guilt). This is an interesting topic for future research.

7 Conclusion

We develop a framework 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 the targeting benefit of in-kind provision of home care appears to exceed its large moral hazard cost. Although targeting with in-kind transfers

[24] In-kind provision of home care seems likely to improve tax system efficiency and alleviate the Samaritan's dilemma by reducing informal care and increasing the labor supply of potential informal caregivers (Ettner, 1995; Van Houtven et al., 2013; Skira, 2015). In-kind provision of home care may have significant paternalistic benefits as well, given the severe cognitive health problems from which some recipients suffer. This was one of the main concerns with more flexible benefits that the Cash and Counseling experiments aimed to evaluate.

[25] 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 Medicaid 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.

is costly, the main alternative—using (more extensively-)tagged cash benefits—appears to be much less effective in this context. These conclusions are fundamentally driven by the substantial 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.

Our results have important policy implications. Several recent policy reforms and proposals make restrictive in-kind benefits more flexible and cash-like. A major impetus for these reforms is the view that recipients would much prefer cost-equivalent cash transfers, a view that is consistent with our analysis of the particular case of Medicaid home care. But such reforms typically also 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 more difficult with flexible benefits, as our analysis suggests is the case in home care, the gain from increasing the value of the benefit to recipients must be weighed against the cost of any worsening of targeting.

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

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

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

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

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

Atkinson, A. B. and J. E. Stiglitz (1976). The design of tax structure: direct versus indirect taxation. *Journal of Public Economics 6*(1–2), 55–75.

Barczyk, D. and M. Kredler (2017). Evaluating long-term care policy options, taking the family seriously. *Review of Economic Studies*, rdx036.

Basurto, P. M., P. Dupas, and J. Robinson (2017). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi. *NBER Working Paper 23383*.

Benitez-Silva, H., M. Buchinsky, and J. Rust (2004). How large are the classification errors in the Social Security disability award process? *NBER Working Paper 10219*.

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

Braun, R. A., K. A. Kopecky, and T. Koreshkova (2017, March). Old, frail, and uninsured: Accounting for puzzles in the U.S. long-term care insurance market. FRB Atlanta Working Paper 2017-3, Federal Reserve Bank of Atlanta.

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. Technical report.

Coe, N. B. and C. H. Van Houtven (2009). Caring for mom and neglecting yourself? The health effects of caring for an elderly parent. *Health Economics 18*(9), 991–1010.

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

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

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

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

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

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.

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

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

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

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

Finkelstein, A. and 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.

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

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.

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

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

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

Kaplow, L. et al. (2008). Optimal policy with heterogeneous preferences. *The BE Journal of Economic Analysis & Policy 8*(1), 1–30.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature 49*(4), 961–1075.

Kemper, P. (1988). The evaluation of the National Long Term Care Demonstration: Overview of the findings. *Health Services Research 23*(1), 161.

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

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

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

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

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.

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

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.

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

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

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

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

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

Saez, E. (2002). The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *Journal of Public Economics 83*(2), 217–230.

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

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

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

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

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

Viscusi, W. K. and W. N. Evans (1990). Utility functions that depend on health status: estimates and economic implications. *American Economic Review 80*(3), 353–374.

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

Tables and Figures

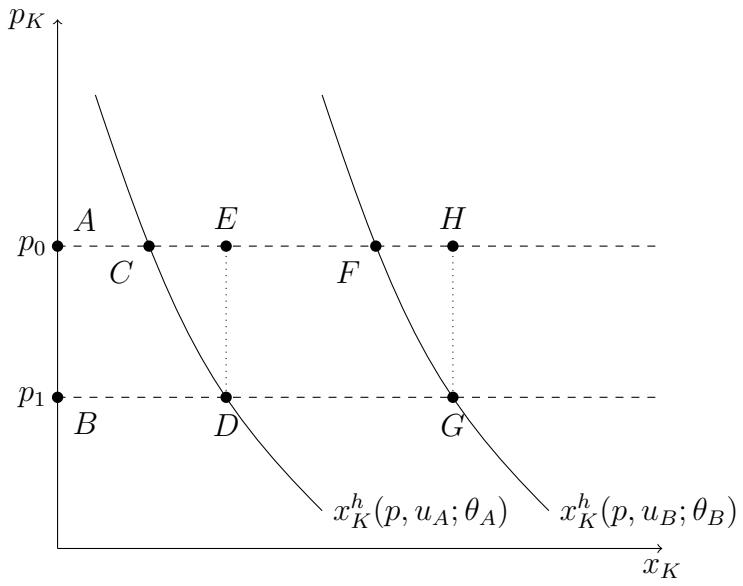
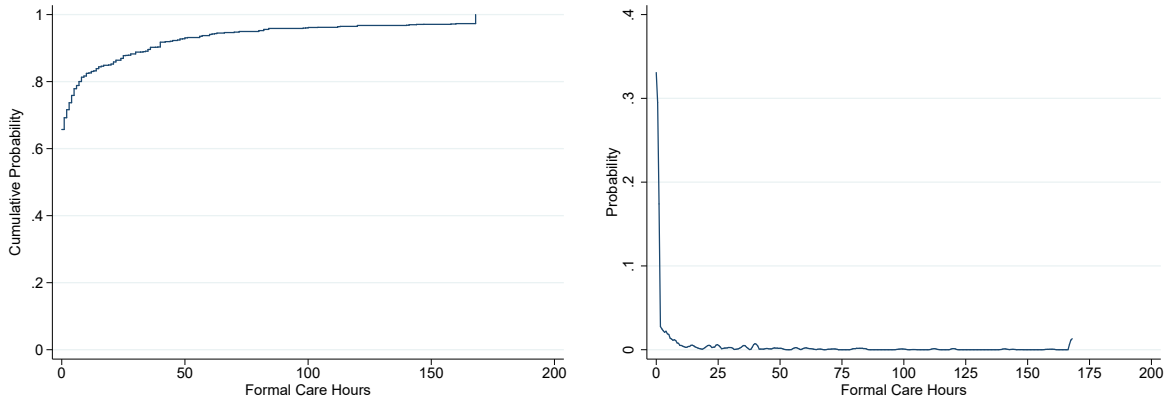


Figure 1: Equivalent variations and excess burdens of a subsidy

[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. It is greater in the high-demand state B, in which it is the area bounded by the vertices $ABGF$, than in the low-demand state A, in which it is the area bounded by the vertices $ABDC$). The excess burdens of the subsidy are independent of the level of demand; they depend only on the slope. The excess burden in state A is the area bounded by the vertices CDE. The excess burden in state B is the area bounded by the vertices FGH.]



(a) CDF (b) PDF

Figure 2: Distribution of Formal Care Consumption in the Benefit-Eligibile Population
[Consumption of formal home care, in hours per week, among the non-institutionalized population aged 65 and older with two or more ADL limitations. Data from the 1999 National Long-Term Care Survey. 65 percent do not consume any formal care. Conditional on consuming formal care, median consumption is 12 hours per week, the 75th percentile is 40 hours per week, the 90th percentile is 120 hours per week, and the 95th and 99th percentiles are 168 hours per week (around-the-clock care). One individual reported consuming more than 168 hours of care per week and has been omitted from the figures.]

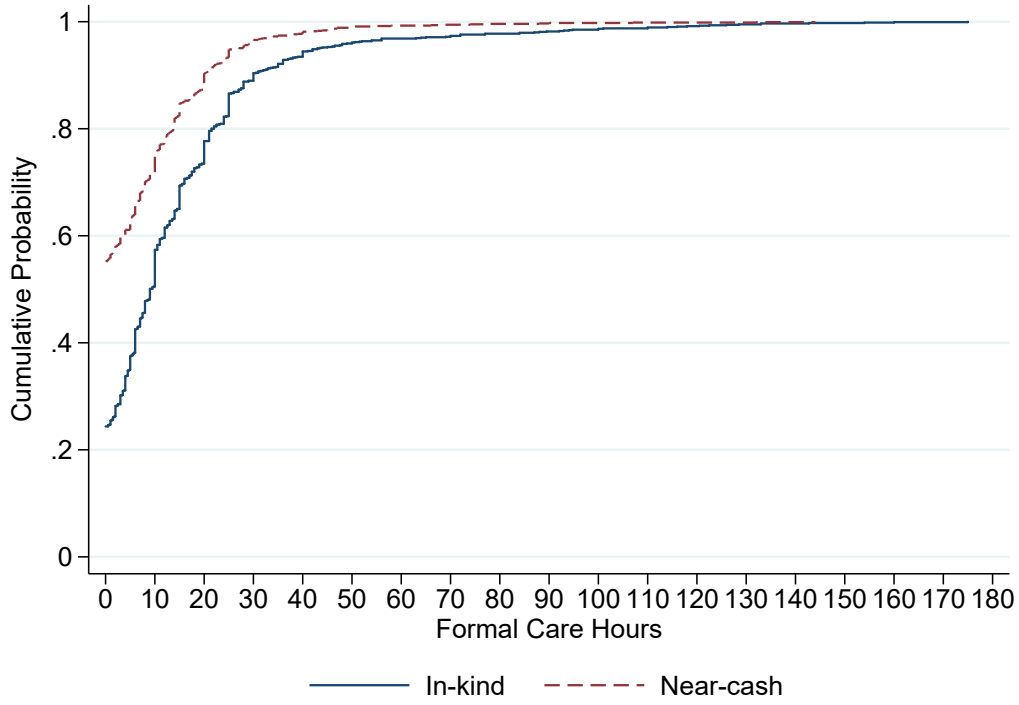
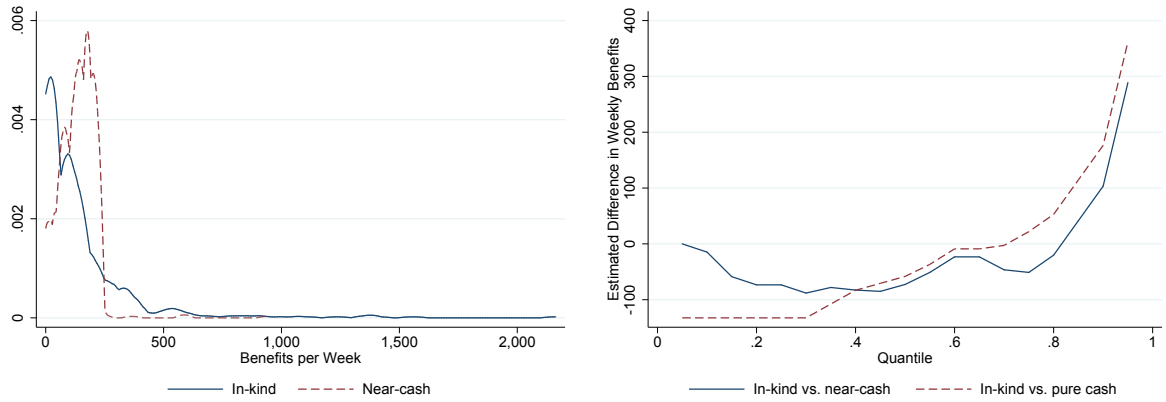


Figure 3: CDFs of Formal Care Consumption by Randomized Benefit Assignment
[Formal home care consumption in hours per week among participants randomly assigned to in-kind vs. near-cash benefits. Data from Cash and Counseling follow-up survey.]



(a) PDFs of Benefits, Excluding Zeros

(b) Differences of Benefits

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

[Distributions and differences of benefits in the Arkansas Cash and Counseling experiment. Benefits are measured in cost per week at market prices. Groups are based on each individual's randomized assignment. Panel (a) plots kernel density estimates. Panel (b) shows the excess of the in-kind benefit over either the near-cash benefit or a hypothetical uniform cash benefit of $133 per week, the average benefit of those assigned to the in-kind benefit. We limit the sample to Arkansas, the only state with information on care plan hours, since care plan hours are needed to estimate the near-cash benefit. The near-cash benefit is the product of care plan hours and the hourly price of care. The in-kind benefit cost is the product of hours of care received and the hourly price of care. We scale up the near-cash group's benefit to have the same mean as the in-kind group's (the near-cash group's average benefit was slightly smaller) to isolate differences in the concentration of benefits, not their average size.]

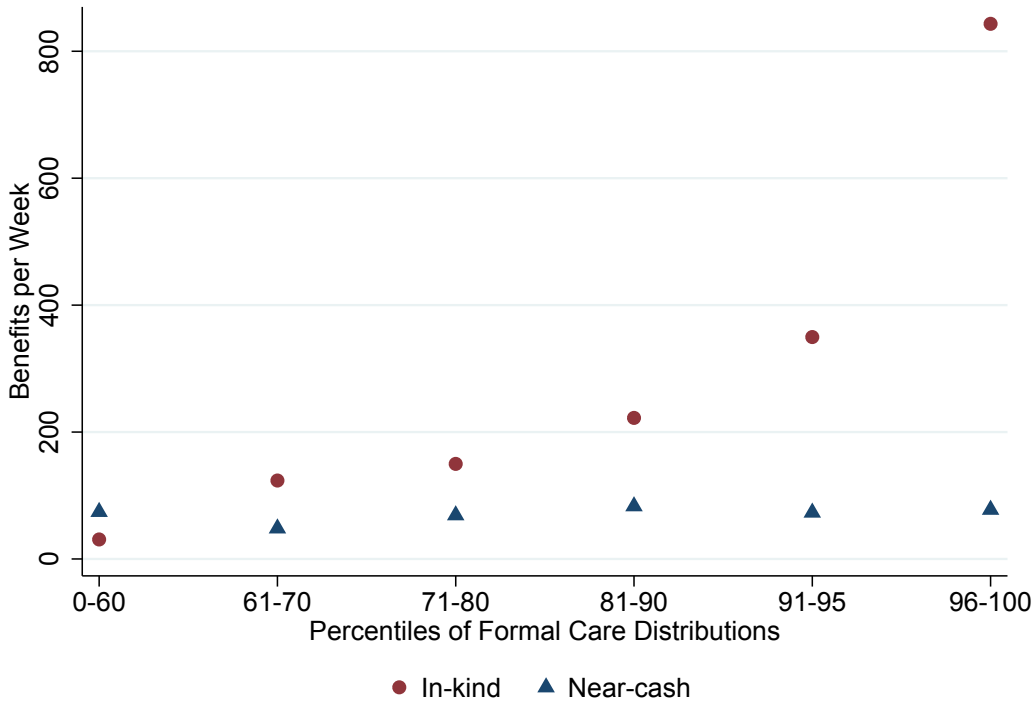


Figure 5: Targeting Effects of In-Kind Provision on the Intensive Margin

[Average benefit costs per week in the Arkansas Cash and Counseling experiment, separately for those randomized to the in-kind and near-cash benefit. 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.]

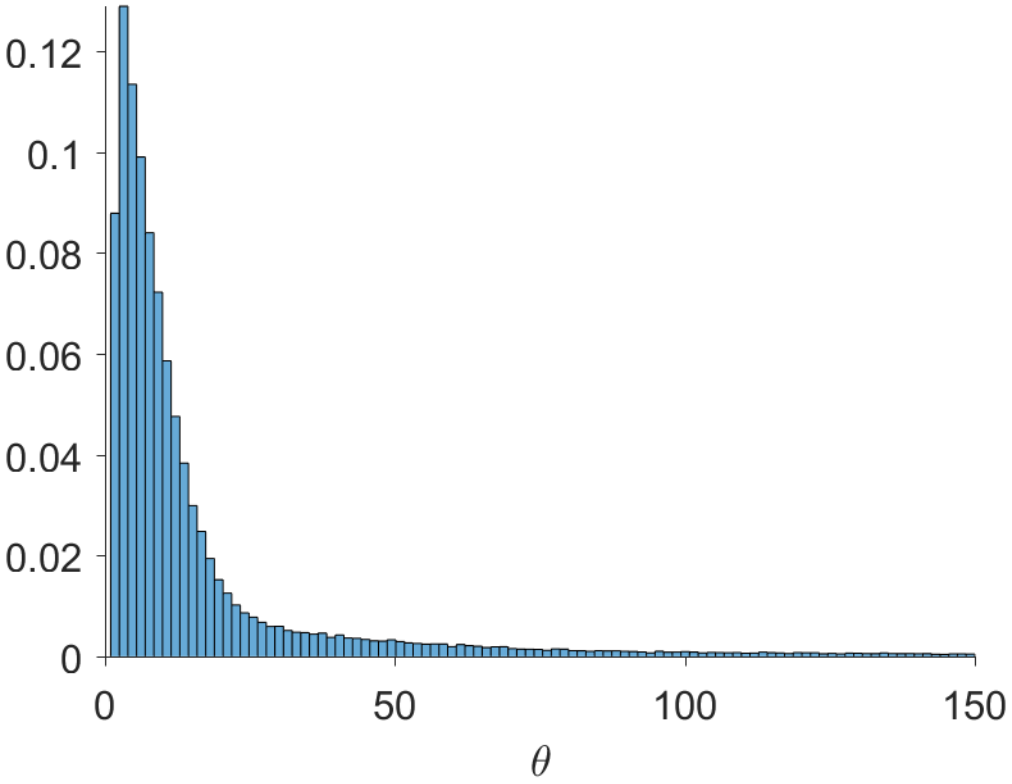


Figure 6: Distribution of the demand for formal care

[Simulated distribution of formal care satiation points, θ, in hours per week, among the non-institutionalized population aged 65 and older with two or more ADL limitations. The mean is 21 hours per week.]

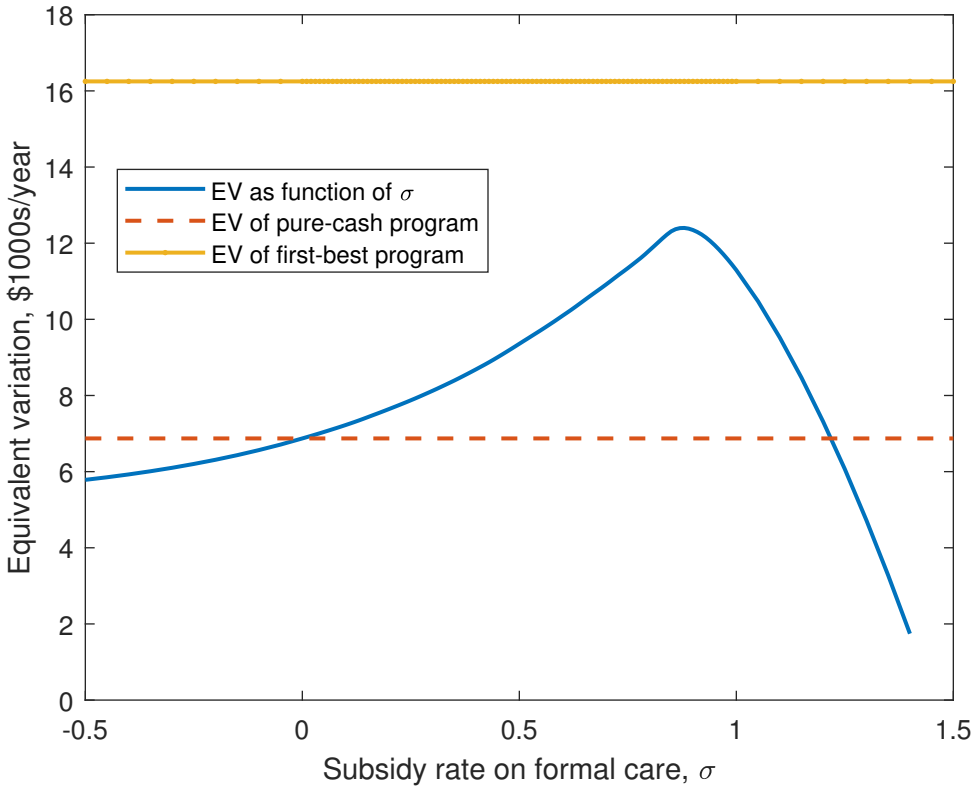


Figure 7: Equivalent variation of mixed cash/in-kind program as function of subsidy rate, σ [Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total program spending. $\sigma = 1$ corresponds to a pure in-kind benefit program, a 100 percent subsidy on formal care with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

Table 1: Predicting Formal Care Consumption

	(1) OLS in-sample	(2) OLS out-of-sample	(3) Machine learning out-of-sample
NLTCS:			
Health controls	0.141	0.072	0.080
	(0.020)	(0.032)	(0.028)
Add informal care	0.168	0.080	0.085
	(0.020)	(0.037)	(0.028)
Add income	0.185	0.086	0.091
	(0.021)	(0.038)	(0.026)
Add interactions with ADLs	0.217	0.043	
	(0.027)	(0.055)	
Make all categorical except income	0.280	-0.063	
	(0.026)	(0.091)	
All categorical, interact with unmarried	0.353	-0.206	
	(0.030)	(0.144)	
All categorical, interact with ADLs	0.718	-1.798	
	(0.038)	(1.046)	
All health, informal care, and income			0.130
			(0.031)
Cash and Counseling:			
Care plan	0.005	-0.009	-0.010
	(0.004)	(0.011)	(0.016)
Health, informal care, demographics	0.191	0.132	0.146
	(0.026)	(0.028)	(0.030)
Add care plan	0.200	0.133	0.154
	(0.026)	(0.029)	(0.028)

Means and standard deviations (in parentheses) of R^2 statistics based on 500 random splits into training and testing subsamples. Rows denote different sets of variables included as predictors. Columns denote the prediction model used (OLS or machine learning) and whether it is an in-sample or out-of-sample R^2.
NLTCS: Sample is non-institutionalized individuals aged 65 and older with two or more ADL limitations. All rows include controls for the price of formal home care in the individual's state (minimum, maximum, and 25th, 50th, and 75th percentiles of price distribution) and indicators for Medicaid and Medicaid home care use. "Health controls" include age, number of ADLs, self-rated health, and sex. "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. "Add interactions with ADLs" includes full set of variables plus interactions of each variable with the number of ADLs. "Make all categorical except income" includes indicator variables for each value of each variable except income and prices. "All categorical, interact with unmarried" includes interactions with an indicator for being unmarried. "All categorical, interact with ADLs" includes interactions of each variable with the number of ADLs. "All health, informal care, and income" includes all 946 variables in the NLTCS related to health, informal care options, or income. See Appendix Table A.3 for further information. This row cannot be estimated with OLS because there are more variables than observations.
Cash and Counseling: Sample is participants in Arkansas aged 65 and older. "Care plan" includes care plan hours at baseline and twelve months. "Health, informal care, demographics" includes controls for health, informal care, demographics, and the price of care, but not care plans. The final row includes both the care plan variables and the controls for health, informal care, demographics, and the price of care.

Table 2: Average Formal Care Consumption by Treatment Group

	(1) Near-cash	(2) In-kind	(3) Difference p-value
Overall	6.85	14.19	0.00
Arkansas	6.29	10.76	0.00
Florida	7.69	18.60	0.00
New Jersey	7.01	16.10	0.00

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

Table 3: The Price Sensitivity of Demand for Formal Care

	(1)	(2)
Price	-1.80***	-1.77***
	(0.15)	(0.15)
Controls	No	Yes
Mean hours, in-kind	14.19	14.19
Observations	2,440	2,440

Dependent variable is formal care consumption in hours per week. Specifications are instrumental variables Tobits where formal care hours are censored at zero. Controls included in column 2 are indicators for sex, education level, race, self-rated health at baseline, living alone at baseline, five-year age bins, and state. Data are from the Cash and Counseling experiments. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table 4: Targeting of Medicaid Home Care

	(1) Take-up = 0	(2) Take-up = 1	(3) Difference p-value
Fraction of eligibles who do vs. do not take up, under different definitions of eligibility			
Income eligible, < 2 cars	0.96	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	7.78	22.83	0.00
Age	79.92	80.19	0.80
Four or more ADLs	0.47	0.64	0.01
Health fair or poor	0.68	0.78	0.07
Female	0.70	0.72	0.65
Lives alone	0.27	0.37	0.10
Unmarried	0.56	0.64	0.20
Has children	0.76	0.78	0.74
Income	682.14	560.17	0.02

Means for people who did (column 2) vs. did not (column 1) take up Medicaid home care. "Difference p-value" tests the equality of means across groups. Take-up rates based on non-institutionalized individuals aged 65 and older with two or more ADL limitations who meet different sets of financial-related eligibility criteria. Income eligible is based on the income thresholds each state uses to determine eligibility. Restrictive income applies the lowest income limit to all states to try to estimate an upper bound on take up. Number of cars is an important determinant of eligibility for Medicaid home care. Summary statistics by take-up decision are for those who meet the "Income eligible, <2 cars" criteria. This sample has 481 individuals. The level of formal care demand, in hours per week, uses our estimate of price sensitivity to simulate each individual's consumption if she faced a price of $18.50 per hour, the maximum in the data. The alternative to health fair or poor is health good or excellent. Data from the 1999 NLTCS.

Table 5: Welfare Analysis and Robustness

	(1)	(2)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
		β			Low θ's		State-dependent utility						θ distribution	
	Baseline	0	25	50	0	Max	in-dec	in-inc	out-dec	out-inc	$\gamma = 1$	$\bar{c}$ =\$2.5k	Drop θ >50	$\theta/2$
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 EV over total cost	0.48	0.65	0.88	-	0.44	0.51	0.02	-	0.63	0.26	-	0.52	0.85	0.73
$E(q_{FC}$|optimal subsidy)	14.01	15.80	2.84	0.00	11.67	17.73	20.31	3.78	11.55	17.30	3.78	14.15	5.37	4.97
$E(q_{FC}$|pure-cash policy)	5.72	13.38	0.02	0.00	5.51	6.26	5.72	5.72	5.72	5.72	5.72	5.72	2.33	2.18
Targeting benefit														
Corr(marg. utility, EV)	0.84	0.90	0.24	-0.68	0.90	0.81	0.11	0.19	0.64	0.85	-0.94	0.81	0.74	0.81
E(1(subsidy $\succ$ cash pol.))	0.16	0.24	0.04	0.00	0.18	0.14	0.14	0.83	0.15	0.15	0.83	0.16	0.18	0.13

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 EV over cost" is the total ex-post equivalent 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, EV)" is the correlation between marginal utility in the absence of any policy and the ex-post equivalent variation of benefits under the optimal subsidy. "E(1(subsidy $\succ$ cash pol.))" is the fraction of people who prefer the optimal subsidy to the pure-cash policy benefit ex post. Column 1 presents results under the baseline assumptions. Columns 2–5 vary the value of β away from the baseline estimate of 1.8. Columns 6 and 7 vary the values of the θ's corresponding to people who consume no formal care when facing a positive price, which are only partially-identified. Column 6 sets these θ's to zero. Column 7 sets each of these θ's to the maximum values consistent with each individual's choice to consume no formal care. Columns 8–11 use different models of state-dependent utility in which $\mu(\theta)$ is linear in θ and varies by a factor of 100 over the range of θ, $\max_\theta\{\mu(\theta)\}/\min_\theta\{\mu(\theta)\} = 100$. 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, $\mu(\theta)$ is decreasing, and in columns 9 and 11, $\mu(\theta)$ is increasing. See Appendix A.6.3 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 θ that exceed 50 hours per week. Column 15 cuts every θ value in half.

Table 6: Welfare Effects of Tagging Benefits

	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		

Average formal care consumption, in hours per week, is estimated in the NLTCS. The sample consists of non-institutionalized individuals aged 65 and older with two or more ADL 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.1 Theory Appendix

A.1.1 In-kind benefits with maximum benefit limits

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

As discussed in Section 2, in-kind transfers have the same effect on recipients' choice sets as potentially nonlinear subsidies. The simplest type of in-kind transfer with a binding maximum benefit limit allows recipients to consume up to a certain, finite amount of the good free of charge and does not subsidize consumption beyond that limit. Many food transfer programs, for example, have this structure. Provided that resale is not possible, this has the same effect on recipients' choice sets as a piecewise-linear subsidy schedule with a 100 percent marginal subsidy rate on the first μ units of consumption and a 0 percent marginal subsidy rate on any additional units of consumption.[26]

Consider a benefit program that combines a cash benefit, b, with a 100 percent subsidy on the first μ units of consumption of good K and no subsidy on additional consumption beyond μ. The individual automatically receives the cash benefit, regardless of the state of the world, but may or may not take up the in-kind benefit. Any in-kind benefit the individual receives cannot be resold. Expected program spending, B, is the sum of the cash benefit and expected spending on the in-kind benefit:

$$b + p_K^0 \int_\Theta TU(\mu, B; \theta) \min\{\mu, x_K(\mu, B; \theta)\} g(\theta) d\theta = B,$$

where $TU(\mu, B; \theta) \in \{0, 1\}$ is an indicator of whether the individual takes up in-kind benefits in state θ under the policy (μ, B).

Consider a budget-neutral shift toward in-kind provision. This increases the maximum benefit limit, μ, while decreasing the cash benefit to maintain the same expected spending. The change in the cash benefit that maintains the same program budget in response to a

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

marginal increase in the maximum benefit limit is

$$\frac{\partial b(\mu, B)}{\partial \mu} = -p_K^0 \int_\Theta \left(TU(\mu, B; \theta) \frac{\partial \min\{\mu, x_K(\mu, B; \theta)\}}{\partial \mu} + \min\{\mu, x_K(\mu, B; \theta)\} \frac{\partial TU(\mu, B; \theta)}{\partial \mu} \right) g(\theta) d\theta.$$

The cash benefit falls by the increase in expected spending on the in-kind benefit, which is comprised of the increase in benefits among those who take up and any increase in take up. If take up is unaffected by this particular marginal shift toward in-kind provision, this equation simplifies to

$$\frac{\partial b(\mu, B)}{\partial \mu} = -p_K^0 \int_\Theta TU(\mu, B; \theta) \frac{\partial \min\{\mu, x_K(\mu, B; \theta)\}}{\partial \mu} g(\theta) d\theta.$$

The reduction in the cash benefit needed to hold expected spending fixed in response to a marginal increase in the maximum benefit limit is increasing in the fraction of people who take up benefits and in the expected increase in the amount of benefits taken up by those who take up benefits.

The marginal ex-ante welfare gain from the shift toward in-kind provision is

$$\begin{aligned} \frac{\partial EU(\mu, B)}{\partial \mu} &= \int_\Theta \frac{dv(p(\mu; \theta), m(\mu, B; \theta); \theta)}{d\mu} g(\theta) d\theta \\ &= \int_\Theta \lambda(\mu, B; \theta) \left(V(\mu, B; \theta) + \frac{\partial b(\mu, B)}{\partial \mu} \right) g(\theta) d\theta. \end{aligned}$$

The first term in the parentheses, $V(\mu, B; \theta)$, is the ex-post value in state θ of the reduction in the price of the μth unit of good K from the market price, p_K^0, to zero. This value depends on the individual's demand for K in state θ. In inframarginal states in which the individual would consume at least the maximum benefit limit even if she could resell the in-kind benefit at the market price, the value of the reduction in the price of the μth unit of good K equals p_K^0, the full marginal cost of supplying the greater benefit in this state. In "extra-marginal" states in which the individual would consume strictly less than the maximum benefit limit if she could resell the in-kind benefit at the market price, the value of the reduction in the price of the μth unit of good K is strictly less than p_K^0. The value can be expressed as $V(\mu, B; \theta) \equiv TU(\mu, B; \theta) MRS_{K,A}(\mu, B; \theta)$, where $MRS_{K,A}(\mu, B; \theta)$ is the marginal rate of substitution, the marginal value of K in terms of A, at the chosen allocation under policy (μ, B).

The marginal welfare gain from the shift toward in-kind provision can be rewritten as

$$\frac{\partial EU(\mu, B)}{\partial \mu} = Cov_{\Theta}\left[\lambda(\sigma, B; \theta),\ V(\mu, B; \theta)\right] - E_{\Theta}[\lambda(\sigma, B; \theta)]E_{\Theta}\left\{TU(\mu, B; \theta)\left[p_K^0 \frac{\partial \min\{\mu, x_K(\mu, B; \theta)\}}{\partial \mu} - MRS_{K,A}(\mu, B; \theta)\right]\right\}.$$

The first term is the targeting benefit, the covariance between the marginal utility of income and the value of the increase in the maximum benefit limit. The second term is the consumption distortion, the excess of the cost over the expected value of the increase in in-kind benefits.

Although this analysis considers a different counterfactual than that in the main text, the same core tradeoff of in-kind provision arises and the same considerations apply. The targeting benefit of in-kind provision is increasing in the variance in marginal utility, the variance in the value of increasing the benefit limit, and the correlation between marginal utility and the value of increasing the benefit limit. The value of increasing the benefit limit is closely related, though not identical, to the level of demand for the good. The distortion cost is increasing in the extent to which in-kind provision leads people to consume more of the good than they would when facing the market price.

Unlike the case in the main text of increasing a linear subsidy rate with no maximum benefit limit, the case of increasing a maximum benefit limit on a 100 percent subsidy has a targeting effect only if take up is incomplete, $TU(\mu, B; \theta) < 1$, or if the individual reaches satiation below the benefit limit in at least some states of the world. If neither of these conditions is met, the increase in the maximum benefit limit causes a one-for-one increase in expected in-kind benefits received, and the cash benefit must fall by the full cost of this increase in in-kind benefits. In this case, the shift toward in-kind provision simply eliminates part of the choice set in each state without adding any new options. Absent optimization failures by recipients or external costs or benefits of trade in good K, this shift toward in-kind provision weakly reduces welfare.

A.1.2 Analysis of a budget-neutral shift toward in-kind provision in terms of income and substitution effects

Expected program spending is

$$B = b + \sigma p_K^0 \int_{\Theta} x_K(p(\sigma; \theta), m(\sigma, B; \theta); \theta) g(\theta) d\theta,$$

where $x_K(p(\sigma; \theta), m(\sigma, B; \theta); \theta)$ is Marshallian demand for good K. Totally differentiating expected program spending with respect to the subsidy rate σ, holding fixed total expected

spending B, gives

$$0 = \frac{\partial b(\sigma, B)}{\partial \sigma} + p_K^0 E_\Theta \left(x_K(p, m; \theta) \right) + \sigma p_K^0 E_\Theta \left(\frac{\partial x_K(p, m; \theta)}{\partial p_k} (-p_k^0) + \frac{\partial x_K(p, m; \theta)}{\partial m} \frac{\partial b(\sigma, B)}{\partial \sigma} \right).$$

The second term is the mechanical effect of the increase in the subsidy rate on expected program spending. The third term is the behavioral effect, which is the sum of the effects of the reduction in the net-of-subsidy price and the induced change in the cash benefit necessary to hold expected program spending fixed. The effect of the reduction in the net-of-subsidy price can be further decomposed into income and (compensated) substitution effects:

$$0 = \frac{\partial b(\sigma, B)}{\partial \sigma} + p_K^0 E_\Theta \left(x_K(p, m; \theta) \right) + \\ \sigma p_K^0 E_\Theta \left(-p_k^0 \frac{\partial x_K^h(p, u; \theta)}{\partial p_k} + p_k^0 x_k \frac{\partial x_K(p, m; \theta)}{\partial m} + \frac{\partial x_K(p, m; \theta)}{\partial m} \frac{\partial b(\sigma, B)}{\partial \sigma} \right),$$

where $x_K^h(p, u)$ is Hicksian demand for good K. The change in the cash benefit necessary to hold expected spending fixed in response to a marginal increase in the subsidy rate is

$$\frac{\partial b(\sigma, B)}{\partial \sigma} = \frac{-p_K^0 E_\Theta \left(x_K(p, m; \theta) \right) + \sigma (p_K^0)^2 E_\Theta \left(\frac{\partial x_K^h(p, u; \theta)}{\partial p_k} - x_k \frac{\partial x_K(p, m; \theta)}{\partial m} \right)}{1 + \sigma p_K^0 E_\Theta \left(\frac{\partial x_K(p, m; \theta)}{\partial m} \right)}.$$

The numerator is the sum of the mechanical and behavioral effects of the reduction in the net-of-subsidy price. The denominator scales these effects to account for the fact that the induced change in the cash benefit itself affects consumption of x_K and so the cost of the subidy component of the program.

A.1.3 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 subject to expected program spending being B:

$$\max_\sigma EU(\sigma, B) = \int_\Theta v(p(\sigma; \theta), m(\sigma, B; \theta); \theta) g(\theta) d\theta \text{ s.t. } b + \sigma p_K^0 \int_\Theta x_K(\sigma, B; \theta) g(\theta) d\theta = B.$$

The first-order condition, which holds with equality at an interior optimum, σ^*,[27] is

$$\frac{\partial EU(\sigma^*, B)}{\partial \sigma} = \int_\Theta \frac{dv(p(\sigma^*;\theta), m(\sigma^*, B;\theta);\theta)}{d\sigma} g(\theta)d\theta = E_\Theta\left(\lambda(\sigma^*, B;\theta)\frac{\partial V(\sigma^*, B;\theta)}{\partial \sigma}\right) = 0$$
$$\iff Cov_\Theta\left[\lambda(\sigma^*, B;\theta),\ x_K(\sigma^*, B;\theta)\right] = \sigma^* E_\Theta[\lambda(\sigma^*, B;\theta)]E_\Theta\left(\frac{\partial x_K(\sigma^*, B;\theta)}{\partial \sigma}\right). \quad (2)$$

The second version of Equation 2 shows that, at the margin at an optimum, the covariance between marginal utility and the level of demand for K must be the same sign as the mean marginal change in K due to the shift in benefit composition, i.e.,

$$\text{sign}\left(Cov_\Theta\left[\lambda(\sigma^*, B;\theta),\ x_K(\sigma^*, B;\theta)\right]\right) = \text{sign}\left(E_\Theta\left(\frac{\partial x_K(\sigma^*, B;\theta)}{\partial \sigma}\right)\right).$$

This is the classic insurance–moral hazard tradeoff. Absent moral hazard, i.e., if $E_\Theta\left(\frac{\partial x_K(\sigma^*, B;\theta)}{\partial \sigma}\right) = 0$, the optimal benefit fully eliminates the covariance between marginal utility and the demand for K, $Cov_\Theta\left[\lambda(\sigma^*, B;\theta),\ x_K(\sigma^*, B;\theta)\right] = 0$. More generally, the greater is the marginal moral hazard cost of shifting toward in-kind provision, the greater must be the marginal targeting benefit.

The first version of Equation 2 implies that, at the margin at an interior optimum, the benefit in some states of the world from shifting toward greater in-kind provision must be exactly offset by the cost in other states. Suppose there are two states, L and H. Then at an interior optimum, at the margin the planner optimally imposes

$$\frac{\left|\frac{\partial V(\sigma^*, B;\theta_L)}{\partial \sigma}\right|}{\frac{\partial V(\sigma^*, B;\theta_H)}{\partial \sigma}} = \frac{p_H \lambda_H}{(1-p_H)\lambda_L}$$

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

A.1.4 First best

In the first-best case in which the state of the world is verifiable, the planner can choose different (b, σ) benefit bundles for each state. The total derivative of indirect utility in state

[27] 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. Necessary conditions for a greater-than-100-percent subsidy to be feasible are that recipients cannot freely dispose of the good and that they eventually become satiated with the good.

θ with respect to the in-kind component of its benefit, σ, is

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

which is non-positive for all positive subsidy rates. When the state is verifiable, a pure cash contract is optimal, and the cash benefits in each state are chosen to equalize each state's marginal utility. With verifiable states, the planner can provide full insurance without distorting behavior, so there is no reason to introduce distortions.

A.2 Medicaid Home Care and the Cash and Counseling Demonstrations: Additional Background

A.2.1 Medicaid home care

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

Eligibility for Medicaid home care is determined by financial- and health-related criteria. An individual must have sufficiently low income and assets and must have at least two ADL limitations that are expected to last at least 90 days. Medicaid is financed jointly by the federal and state governments, and Medicaid policies vary somewhat across states. In most states, Medicaid provides home care primarily through two programs: the Medicaid Title XIX PCS optional State plan and the Medicaid 1915(c) HCBS waiver program. For the elderly, the means tests for Medicaid home care are often less restrictive than those for general Medicaid coverage. The majority of states provide coverage for individuals with incomes up to 300 percent of the monthly Supplemental Security Income (SSI) amount (LeBlanc et al., 2001). States with more restrictive income limits use 100 percent of the SSI amount.

In principle, the amount of Medicaid home care for which an individual qualifies is determined by a medical exam. The applicant's health care provider must submit a care plan that details the services deemed appropriate based on the applicant's health status. In many states, the amount of care people can receive is also limited by maximum benefit rules. Of the Cash and Counseling states, 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—while Florida had no limit. In practice, however, it appears that in many cases neither care plans nor maximum benefit rules are binding constraints on the benefits received by Medicaid

home care recipients.

In the Cash and Counseling data, formal care consumption, which is measured nine months after baseline, exceeds the number of hours in the individual's care plans at baseline and 12 months later for about 30 percent of Medicaid home care recipients. Although the mismatch in timing makes it possible that some of these individuals had a different care plan in operation when their consumption was measured, the strong correlation between care plan hours at baseline and 12 months later, 0.86, makes it highly unlikely that this can explain much of the excess of consumption over care plan hours. And if care plans were binding, it is not clear what incentive physicians might have to restrict care plan hours below what the recipient, their patient, would like. Physicians' professional norms and ethos emphasize acting as an agent of the patient, not Medicaid or other parties.

Maximum benefit limits do not appear to have been binding either. LeBlanc et al. (2001) survey Medicaid home care programs and discuss several explicit mechanisms for granting exceptions to the limits. For example, recipients in New Jersey, where the statutory limit was 25 hours per week, could with prior authorization receive up to 40 hours of care per week and with central office approval could receive as much care as "needed." Consistent with these or other mechanisms relaxing quantity limits, the distributions of formal care consumption 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 at the kink it creates in the budget constraint, since at a binding upper limit the price of formal care jumps sharply, from zero to the market price.

Appendix Figures A.1–A.3 present the distribution of formal care consumption among people randomized to the in-kind benefit in each of the three Cash and Counseling states. The distribution of formal care consumption in Arkansas, Appendix Figure A.1, shows no apparent signs of being influenced by the statutory limit of 16 hours per week. Nearly one-fifth of the sample consumed more than the statutory limit, and there is no apparent bunching at the statutory limit: Only 1 percent of recipients consume 16 hours per week, whereas 10 percent consume 10 hours per week and four percent consume 15 hours per week. The distribution of formal care consumption in New Jersey, Appendix Figure A.3, potentially shows some sign of being influenced by the statutory limit of 25 hours per week, as 10 percent of people consume the statutory limit. But this bunching is only slightly greater than that at other round-number amounts. For example, 7 percent of people consume 15 hours per week and 9 percent consume 20 hours per week. And about one-sixth of people consume more than the statutory limit. Of course, any test of bunching faces the limitation that measurement error lessens observed bunching. A useful feature of our context in this regard is that the tested-for kink in the budget constraint is quite sharp, increasing the price from zero to the market price. If benefit limits were binding, one would expect them to be highly salient, which might reduce attenuation from reporting error.

Take-up rates are notoriously difficult to estimate both for means-tested programs in general and for Medicaid in particular (U.S. Department of Health and Human Services, 1992; Currie, 2006). Eligibility rules often are complex, vary from state-to-state, and depend on household characteristics that are unobservable to the researcher. We estimate take-up rates of Medicaid home care by combining data from the NLTCS with information on the size of the 65-and-older population and administrative estimates of the number of Medicaid home care users from LeBlanc et al. (2001). We use the NLTCS to estimate the fraction of the elderly who are eligible for benefits, based on the eligibility criteria from Schneider et al. (1999). To be eligible, someone must have at least two ADL limitations and meet income and asset requirements. The main source of uncertainty in our estimated take-up rate is the incompleteness of the information on household assets in the NLTCS. Given this data limitation, we aim to bound the true eligibility rate. Our less restrictive eligibility threshold uses the income limits from Schneider et al. (1999) and limits eligibility to households with fewer than two cars. Our more restrictive eligibility threshold uses (much) more restrictive income and asset requirements than the actual limits in the vast majority of states: Household income must be no more than 100 percent of the SSI benefit and the household must have no cars (car value is one of the primary inputs to the asset tests). The more restrictive the eligibility definition, the greater the implied take-up rate among eligibles. Given that our more restrictive eligibility estimate likely understates eligibility substantially, the implied take-up rate of 16 percent likely exceeds the true take-up rate.

A.2.2 Cash and Counseling Demonstration

The Cash and Counseling experiments were large-scale experiments conducted by the Medicaid programs of Arkansas, Florida, and New Jersey in the late 1990s and early 2000s (for more details see Brown et al., 2007). Participants were enrolled beginning in 1998 in Arkansas, 1999 in New Jersey, and 2000 in Florida. In New Jersey and Florida, only individuals who were currently receiving Medicaid home care were eligible to participate in the demonstrations. Arkansas allowed a limited number of individuals who qualified for but were not receiving Medicaid home care to participate.[28] Both non-elderly and elderly individuals were enrolled and there was no screening on whether the individual had or would be able to find sources of care. Participants were given a baseline survey and then randomized to the traditional in-kind benefit or an experimental near-cash benefit, each with a 50 percent probability. Participants were surveyed 4–6 months after enrollment and again 9 months after enrollment. We use data from the baseline and 9-month follow-up surveys.

The near-cash benefit was slightly less than the cashed-out cost of the individual's care plan. This stemmed from a requirement that the experimental cash treatment be budget-

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

neutral, which meant that the costs of paying the counselors who helped treatment group members manage their care came out of the cash allowances. In New Jersey, for example, 10 percent of the value of the care plan was set aside to cover program costs. Counselors were available to help participants develop plans for spending their benefit, issue checks to caregivers and other service providers, handle paperwork associated with being an employer (e.g. payroll taxes), and maintain the necessary records. Recipients had to submit receipts documenting that they spent at least 90 percent of their benefits on personal care services. The idea was that the remaining 10 percent could be spent on services that could not be readily invoiced, like payments to a neighbor for mowing the lawn.

Appendix Table A.1 provides summary statistics on the Cash and Counseling participants and balance tests of the randomization, and Appendix Table A.2 compares Cash and Counseling participants to several populations of interest using the NLTCS. We restrict the sample to people who are at least 65 years of age and who have nonmissing data on age, sex, race, education, and self-rated health. Our final sample includes 2,470 individuals, of whom 30 are missing data on formal care consumption at follow-up, leaving us with 2,440 individuals for analyses that require this variable. At baseline, average formal care consumption ranges from 9 (Arkansas) to 16 (New Jersey) hours per week, and the average number of informal caregivers is two. The average age is in the upper 70s, the majority of participants are female, and education levels are low. Although non-negligible fractions of the treatment and control groups attrited from the experiment before the nine-month follow-up survey (20 and 35 percent, respectively), of the 30 balance tests, none of the differences between treatment and control groups are statistically significant at the 5 percent level and only one is significant at the 10 percent level—fewer than would be expected to arise by chance without any differential attrition.

Not surprisingly, participants in the experiments are somewhat different from the broader population of Medicaid home care users in the US. Appendix Table A.2 shows that compared to Medicaid home care users in the US, participants in the experiments are similar in terms of age (around 79 on average) and health status (about three-quarters self report fair or poor health), but they have lower formal care consumption (12 vs. 38 hours per week) and are less likely to be living alone (32 vs. 41 percent). The differences could arise from selection into the experiment, differences in the generosity of states' Medicaid home care programs, or from differences in the composition of Medicaid home care users across states. Unfortunately, the NLTCS has too few Medicaid home care recipients in Arkansas, Florida, and New Jersey to address this directly. We discuss issues related to the internal and external validity of our analysis in more detail in Section A.4.

A.3 Predicting Formal Care Consumption

The extent to which observable characteristics predict formal care consumption is informative about the extent to which the risk from chronic health problems could potentially be insured by directly-targeted, tagged cash benefits. We separately predict formal care consumption among the home care benefit-eligible population and among Cash and Counseling participants. We follow the standard approach from the predictive modeling literature: Randomly split the sample into two equal-sized subsamples, train the model on one subsample (the training sample), use the trained model to make predictions for the other subsample (the test sample), and assess predictive power in the test sample. We repeat this process 500 times and report the mean and standard deviation of the results. We implement two separate modeling approaches: OLS and machine learning. OLS is familiar and transparent and produces intuitive output. Machine learning methods are the state-of-the-art for prediction. We implement random forest models and use five-fold cross-validation. Dimensions that are optimized include the number of trees, tree depth, minimum leaf size, and number of variables to (randomly) sample at each split.[29]

Table 1 reports the results. The top panel shows the analysis of the home care benefit-eligible population, non-institutionalized people aged 65 and older with two or more ADL limitations, using the NLTCS. The NLTCS includes 887 benefit-eligible individuals and features an extensive set of individual- and household-level variables, including information about demographics, health, insurance, health care consumption, potential and actual informal caregivers, living arrangements, and formal care prices. It has especially rich measures of health, including both subjective self-rated health and many detailed objective measures. Appendix Table A.3 summarizes the key variables, of which there are 946. Yet despite the richness of the data, across the two techniques and the many specifications, observable characteristics never "explain" more than about 13 percent of the variation in formal care consumption among the benefit-eligible population out of sample, leaving the vast majority of the variation unexplained.[30]

The bottom panel of Table 1 reports the results of a similar analysis of participants of the Cash and Counseling experiments. This analysis differs from the previous one in two main ways. First, Cash and Counseling participants are a selected subset of the benefit-eligible population that is likely much more homogeneous, especially in terms of demand for formal care, than the benefit-eligible population as a whole. Second, the variables in

[29]We have also experimented with two different aggregation methods, bagging and boosting. These have little effect on the results. See James et al. (2013) for background on machine learning techniques and Mullainathan and Spiess (2017) for a discussion of the uses of machine learning in economics.

[30]Including enough variables in the model can produce an in-sample fit that is arbitrarily high. The in-sample R^2 gets as high as 0.72 in one specification. But in-sample R^2 can be a poor measure of predictive ability due to overfitting. Such overfitting is apparent in the out-of-sample fit of the richest models, which perform worse than a simple regression with a constant alone ($R^2 < 0$).

the Cash and Counseling data, while less extensive than those in the NLTCS, include a variable of particular policy interest: Medicaid care plan hours. As discussed in Section 3, Medicaid home care programs require each participant to complete a medical exam with a doctor or nurse, who creates a care plan meant to reflect the recipient's "need" for care. In principle, this measure of "need," which is supposed to reflect not just health problems but the availability of paid and unpaid caregivers as well (Dale et al., 2004), could be as close to a summary measure of demand or marginal utility on which an insurer might feasibly condition benefits. It is an alternative, costly way to condition benefits on individual circumstances beyond using readily-observable characteristics.

This analysis focuses on Cash and Counseling participants in Arkansas, the only state for which we have information on care plans. We use both available measures of care plan hours: care plan hours at baseline and twelve months later. Care plans are updated every six months in Arkansas. Because formal care consumption is measured nine months after baseline, we do not know which, if either, of the care plans is in operation at that time. Fortunately, care plans are highly persistent over the twelve month period—the correlation between an individual's care plan at baseline and follow up is 0.86—which mitigates this concern.

The results show that the vast majority of the variation in formal care consumption among Cash and Counseling participants remains unexplained by care plans or any other variable in the data. The out-of-sample R^2 never exceeds about 0.15. Care plan hours account for little of the variation in formal care consumption, either the raw, unconditional variation or the residual variation when the other variables are included.

The NLTCS analysis suggests that even extensively-tagged cash transfers would leave much of the variation in formal care consumption among the disabled, elderly population uninsured. The Cash and Counseling analysis suggests that even Medicaid's costly way of assessing an individual's "needs" improves little on predictions based on more readily-available characteristics. Having more or better data would presumably improve the predictions at least somewhat. But the NLTCS already has extremely rich data on key health- and informal care-related factors; it is not clear what feasible measures might improve on these. More important, these analyses ignore the moral hazard and verification costs that would be involved in using many of these variables as tags. In practice these factors would limit the net value of using tagged cash benefits still further. Taken as a whole, these results suggest that the scope for insuring home care risk with directly-targeted, tagged cash transfers is limited.

A.4 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 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 estimate of the price sensitivity of demand for formal care.

A.4.1 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 Section 3 and Appendix Section A.2, the amount of Medicaid home care that someone can consume is determined by a care plan written by their physician. If physicians, whether in an effort to be "good agents" of Medicaid or for other reasons, prescribe care plans whose hours fall short of their patient's satiation point, then Medicaid home care recipients may not be able to reach satiation. Although maximum benefit limits and care plans do not appear to have constrained consumption in our context, as discussed in Appendix Section A.2, we assess the robustness of the estimated price sensitivity to different assumptions about how binding these might

have been.

Appendix Table A.6 shows estimates of the price sensitivity of demand separately for each state. The first row shows that the IV Tobit estimates range from -1.04 (Arkansas) to -2.78 (Florida). In the second row, we impose the upper bounds on care hours implied by the Arkansas and New Jersey benefit limits. We censor observations above those cutoffs and use the IV Tobit to re-estimate the price sensitivity. The additional censoring reduces our estimated price sensitivity in Arkansas but increases it in New Jersey. (We exclude Florida since care hours are not limited there.) The differences across states are similar to those found with the standard IV Tobit. Because average care consumption varies somewhat across states, it is also useful to consider the percentage changes implied by the coefficients. A one-dollar increase in the price of formal care is estimated to increase formal care consumption by 10 percent in Arkansas, 10 percent in New Jersey, and 15 percent in Florida.

Generally, the results are consistent with the concern that quantity constraints—whether from supply constraints in Arkansas or statutory limits in Arkansas and New Jersey—might be biasing our price sensitivity estimates towards zero. The state without limits (Florida) consistently displays greater price sensitivity than the other states. This suggests that our estimate will tend to understate the true price sensitivity.

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

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

far greater than any of our estimates can overturn the result that the optimal subsidy on formal care in the model in Section 6 is significantly greater than zero.

A.4.2 External validity

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

Appendix Table A.2 compares Cash and Counseling participants to various representative samples of Americans from the NLTCS. As discussed in Section 3, Cash and Counseling participants are similar to the broader population of Medicaid home care recipients in terms of age (around 79 on average) and health status (about three-quarters self report fair or poor health), but they have lower formal care consumption (12 vs. 38 hours per week) and are less likely to be living alone (32 vs. 41 percent). These differences are consistent with negative selection on demand for formal care of Medicaid home care recipients into the Cash and Counseling experiments. This is unsurprising given that the gain from a more flexible benefit is decreasing in the demand for care. Compared to the broader population of people eligible for home care benefits, whether financially eligible for Medicaid home care or not, Cash and Counseling participants are in worse health (79 vs. 70 percent self report fair or poor health), are more likely to be female (82 vs. 66 percent), and are more likely to be unmarried (81 vs. 55 percent). These differences are consistent with the strong selection into Medicaid home care among the eligible population of those who are sicker and who have worse informal care options, as shown in Table 4, overcoming any selection into the Cash and Counseling experiments among Medicaid home care recipients of those who are healthier and who have better informal care options.

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

There are two key issues that tend to offset each other. First, people whose demand was more sensitive to the composition of benefits had a greater incentive to participate in the experiment. It is therefore natural to expect that participants were more sensitive to the price of formal care than the broader population of Medicaid home care recipients in the Cash and Counseling states. This tends to increase our estimate of the price sensitivity of

demand for formal care relative to what we would expect to find among the population of recipients of Medicaid home care.

Second, the nature of the experiment—especially its unexpected occurrence and uncertain duration—likely reduced the sensitivity of demand to the composition of benefits relative to its likely value under an anticipated, permanent change in policies. Care-giving arrangements, for which people often make important investments like moving or adjusting their labor supply, likely depend on both the past history of policies and expectations about future policies. People arrange their lives in order to make the best of the opportunities available to them, and their decisions about where to live and work and how much formal and informal home care to consume likely depend on which if any home care benefits they might be eligible for. The Cash and Counseling experiments likely came as a surprise to many participants, and it is unclear what participants might have expected about the persistence of this policy. Would it continue indefinitely or would it soon revert back to traditional Medicaid home care? Both the surprise aspect and the uncertainty about how long cash benefits might last likely dampened responses relative to what they would have been under an anticipated, permanent policy.

These considerations suggest caution in applying the results of the Cash and Counseling experiments to other contexts. But the robustness of our welfare analysis to even large changes in the price sensitivity of demand for formal care greatly limit this concern in our context. And, despite the uncertainty about its generalizability, the major strength of the Cash and Counseling experiments—the large, exogenous price variation—makes it a valuable piece of evidence about the demand for formal care and the effects of alternative home care-related policies.

A.5 Targeting Effects of In-Kind Provision: Additional Evidence from the Cash and Counseling Experiments

Those who take up Medicaid home care benefits are a highly selected subset of the population eligible for benefits, in terms of both their observable and unobservable determinants of demand for formal care (see Table 4 and Appendix Table A.7). Among those who take up Medicaid home care, recipients whose observable characteristics would normally suggest a low demand for formal care are likely to have unobservable characteristics that are strongly associated with having high demand for formal care; otherwise they would have been unlikely to take up benefits. (Of course, there may be important heterogeneity in participation costs and awareness of the program as well.) Such selection complicates comparisons of benefits received by different groups of recipients based on their observable characteristics.

For example, although being married is associated with having below-average demand for formal care in the population as a whole, among Medicaid home care recipients being

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

Since selection could significantly bias such levels comparisons, we pursue a differences-in-differences approach that likely mitigates, though does not eliminate, this issue. We also separately analyze the subset of participants of the Cash and Counseling experiments who had not been receiving Medicaid home care before the experiments, who are likely to be more representative of the eligible population as a whole. Even so, selection issues are a major caveat of the results that follow, which at best provide suggestive evidence of the effects of in-kind provision on targeting on the intensive margin. This is one reason why our preferred evidence, discussed in Section 5, is on targeting by formal care demand.

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

$$benefits_i = \beta_0 + \beta_1 inkind_i + \beta_2 X_i + \beta_3 (inkind_i * X_i) + \varepsilon_i \tag{3}$$

where $benefits_i$ is the dollar cost of benefits received by participant i, $inkind_i$ is an indicator for whether i was randomized to the in-kind group, and X_i is a particular demographic characteristic. The coefficient of interest, β_3, tells us whether people with greater values of X_i receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do people with lower values of X_i. For example, if X_i is the number of ADL limitations, $\beta_3 > 0$ would imply that those with more ADL limitations receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do those with fewer ADL limitations. This compares the in-kind benefit to the Cash and Counseling tagged near-cash benefit. Because of the tagging, based on an individual medical exam, the near-cash benefit targets resources more than a hypothetical pure (untagged) cash transfer would. As a result, this analysis likely understates the degree to which in-kind provision targets particular groups relative to a pure cash transfer.

Appendix Table A.8 reports the effects of in-kind provision on average benefits, estimated with OLS regressions, and on the right tail of the benefit distribution, estimated with quantile regressions. The right tail of the distribution is of particular importance because that is where there is the greatest scope for targeting to provide insurance value. If in-kind provision concentrates transfers, the OLS estimates will reflect an average of negative effects at the bottom of the benefits distribution and positive effects at the top. The quantile regressions, by contrast, estimate the effects at the top of the distribution, where targeting is likely to

have the greatest impact on utility.

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

Columns 5 through 8 repeat the analysis for the subset of participants who had not been in the Medicaid home care program at baseline. This group is likely more representative of the roughly 90 percent of eligibles who do not take up Medicaid home care. The patterns are qualitatively similar, though with larger standard errors.

This is suggestive evidence that, on the intensive margin among recipients, in-kind provision targets recipients in worse health and with worse informal care options.

A.6 Welfare Analysis: Further Details and Robustness

A.6.1 The utility function and optimal first-best insurance

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

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

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

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

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

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

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

which is independent of θ. The first-best contract does not distort consumption, and it fully insures all risk. By making larger transfers in states of the world with greater demand for formal care, it fully compensates the individual for her expenditures on formal care and any residual utility costs she faces from coping with her health problems.

A.6.2 Estimating the distribution of demand for formal care

As discussed in the text, we use our estimate of the price sensitivity of demand for formal care, β, to convert the observed joint distribution of formal care consumption and formal care prices in the NLTCS into a distribution of the level of demand for formal care in the benefit-eligible population, $G(\theta)$. We express the level of demand for formal care in terms of satiation points, θ. The only part of this calculation that is not entirely straightforward arises because observed formal care consumption does not point-identify θ for people consuming zero formal care, it only bounds it: $\theta_i \leq \beta p_i$. We estimate the full θ distribution, including the θ's of people who consume zero formal care, in three steps.

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

Second, we convert each q^* to its corresponding θ using the estimated price sensitivity of demand for formal care, $\theta_i = q_i^*(p) + \hat{\beta}p$. This adjusts (potentially latent) formal care consumption by our estimate of the impact of the price on consumption. Finally, we estimate the kernel density of the implied θ distribution. Figure 6 shows the resulting θ distribution.

It is mostly just a rightward-shifted version of the observed distribution of formal care consumption, with adjustments for the censoring of people who consume no formal care.

For the tags analysis, we repeat the same procedure for estimating the θ distribution separately for different subsets of the benefit-eligible population, as defined by their tagged characteristics. Appendix Figures A.5 and A.6 show the θ distributions of people who do vs. do not live alone and of people with different numbers of ADL limitations. All of the distributions are shaped like the corresponding (observed) distributions of formal care consumption, 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 ADL limitations.

In the quantitative analysis, we further constrain θ to be non-negative and, as a baseline, no larger than 150 hours per week. While negative satiation points are not implausible in theory, since some people might consume no formal care even at small negative prices, in practice they are awkward with the baseline utility function, since someone with $\theta < 0$ is worse off than someone with $\theta = 0$. Moreover, behavior when $\theta < 0$ is identical to behavior when $\theta = 0$ as long as the net-of-subsidy price of formal care is non-negative. We truncate the baseline θ distribution at 150 hours per week in order to reduce the influence of outliers. Given the importance of right-tail risks for insurance, though, we also report results under different assumptions about the right tail of the θ distribution.

We test the robustness of our results to making different extreme assumptions about how to fill in the unidentified θ values. In one case, we set every unidentified θ 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 θ's equal to their (point-identified) upper bound, $\theta_i = \hat{\beta} p_i$.

A.6.3 State-dependent utility

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

People in worse health likely have different utility functions from people in better health; they likely have a lower level of utility, for example. But what matters for insurance is marginal utility, and a priori it is not clear in which direction a reduction in health might shift marginal utility. On one hand, activities like eating out and traveling likely become

less attractive, which tends to reduce marginal utility. On the other hand, home upgrades and equipment likely become more attractive, which tends to increase marginal utility.

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

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, state-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 states. States 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, state-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, states 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, states 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 exactly offset, 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 than in the log case, such as constant relative risk aversion preferences with a coefficient of risk aversion greater than one, the latter effect dominates and states with greater scaling factors have lower marginal utility for any given level of net consumption.

We consider a wide range of state-dependence in utility. In Table 5, we report the results of different models of state-dependent utility in which $\mu(\theta)$ is linear in θ and varies by a factor of 100 over the range of θ, $\max_\theta\{\mu(\theta)\}/\min_\theta\{\mu(\theta)\} = 100$. We also analyze the effects of state-dependent utility based, as closely as possible, on the estimates of Finkelstein et al.

(2013). Finkelstein et al. (2013) estimate the state-dependence of utility in the number of chronic health problems someone has.[31] It is important to emphasize that their estimates do not map perfectly to our context, whether to the level of demand for formal care, θ, or to the number of ADL limitations someone has. But it is the best evidence on the likely extent of state-dependent utility in a related context.

Finkelstein et al. (2013) estimate, using "outer state-dependence" in our language, that a one-standard deviation increase in the number of chronic health problems is associated with a 10–25 percent decline in marginal utility. We adapt this evidence to our setting by assuming that a one-standard deviation increase in the number of chronic health problems corresponds to a one-standard deviation increase in the level of demand for formal care, θ, while continuing to assume that $\mu(\theta)$ is linear in θ. The midpoint of Finkelstein et al.'s (2013) range of estimates is similar to, though slightly more severe than, the state-dependence implied by the 100-fold variation in $\mu(\theta)$ over the range of θ.

A.6.4 Robustness and intuition

This section provides additional information about the results in Table 5. The results of the alternative specifications provide information about the robustness of the conclusions to different assumptions and shed light on the key factors driving the results. We first summarize the results as a whole. Then we discuss each in more detail and provide intuition.

As Table 5 shows, the results are highly robust to plausible changes in the model. The price sensitivity of demand for formal care must be quite large—over 10 times greater than our estimate based on the Cash and Counseling experiments—in order to overturn the conclusion that the optimal subsidy is large. Even if the distribution of partially-identified θ values is in the "worst-case" configuration (i.e., each θ_i equals the maximum value consistent with i's behavior), the optimal subsidy rate is still 86 percent. The utility function must exhibit strong state dependence of just the right kind—greatly decreasing relative 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 means 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 θ values are scaled down. 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, as discussed in the main text, 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 provision per se.

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

The results are highly robust to changes in β, the price sensitivity of demand for formal care. The net benefit of subsidizing formal care is decreasing in β, but only at a low rate.[32] The great extent to which the results are robust to changes in β might seem surprising given the importance of moral hazard costs in determining the welfare effects of in-kind provision. Part of the explanation is that β affects more than just the moral hazard cost of in-kind provision; it also affects the utility cost of uninsured home care risk, and this latter effect can partly offset the moral hazard effect.[33] The other reason for the robustness is the large targeting benefit from in-kind provision, which comes from the combination of the significant heterogeneity within benefit-eligible states, the positive correlation between marginal utility and the demand for formal care, and, as noted by Kaplow (2011), the rapid rate at which marginal utility diminishes in consumption under standard utility functions.

The results are robust to making either of the two possible extreme assumptions about the values of the partially-identified, smallest θ's. This robustness is partly because the bounds are relatively tight, which can be seen by comparing mean formal care consumption under each of these assumptions. Under a pure-cash benefit policy, mean formal care consumption when the unidentified θ values are all set to zero is 5.5 hours per week, whereas it is 6.3 hours per week when each unidentified θ takes its upper bound. This robustness is also because the left-most part of the θ distribution has relatively little effect on the value of insurance, since, being the thick part of the distribution, there are a large fraction of states of the world that differ little in their marginal utilities or levels of demand, so changes in the exact distribution have little effect on the moral hazard cost or targeting benefit of in-kind provision.

The results are also robust to plausible levels of state dependence in utility. The targeting

[32]When demand for formal care is completely inelastic ($\beta = 0$), a 100 percent subsidy achieves the first best. One caveat about this result is that it is based on a model in which formal care is borderline inferior. In a more general model with income effects of demand for formal care, a full subsidy might not achieve the first best even if substitution effects are zero.

[33]Increasing β increases the moral hazard cost of subsidizing formal care and reduces the value of insurance against home care risk, both of which tend to reduce the value of in-kind provision, while at the same time reducing the role of the consumption floor, which tends to increase the value of in-kind provision. The net effect of increasing β on the welfare effects of in-kind provision are therefore ambiguous in theory. In practice, as Table 5 shows, in the model of Section 6 the net gain from in-kind provision is decreasing in β, but only slowly. Increasing β decreases both the value of insurance against home care risk and the role of the consumption floor. As discussed in Appendix Section A.6.1, as β approaches zero, the model approaches one in which home care risk is equivalent to a wealth shock equal to the cost of reaching satiation with formal care, $p_i\theta_i$. This minimum value of β, 0, maximizes the utility consequences of any given θ, which maximizes the risk people face. At the other extreme, as β approaches infinity, home care risk becomes irrelevant; the utility cost of any shortfall of formal care consumption relative to its satiation level approaches zero, so in the limit as $\beta \to \infty$ people do not purchase any formal care and there is no risk. This risk reduction effect of increasing β has two effects on the welfare gain from in-kind provision that work in opposite directions. On one hand, the smaller risk means a smaller value of insuring it, which reduces the value of in-kind provision. On the one hand, a smaller risk reduces the role of the consumption floor, since a greater β means people optimally spend less on formal care and as a result have more resources available for non-care consumption and don't rely on the floor as much. This reduces the implicit taxation of home care insurance by the consumption floor, which tends to increase the insurance value of in-kind provision.

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). State dependence that increases the marginal utility of high-demand states relative to low-demand states, columns 8 and 11, tends to increase the optimal subsidy and its equivalent variation substantially. State dependence that decreases the marginal utility of high-demand states relative to low-demand states, columns 9 and 10, decreases the optimal subsidy and its equivalent variation. State-dependent utility of the inner increasing form, column 9, is one of only three specifications in which the optimal subsidy is negative. This is because, in this specification, states of the world with a high demand for care are also states in which the individual is most effective at converting spending into composite consumption. As a result, for any given level of actual consumption, the individual's effective consumption is greater in high-demand states. This tends to significantly reduce the marginal utility of additional spending, given the curvature of utility. Perhaps most relevant given the existing evidence about health-related state dependence in utility is the outer decreasing state dependent utility, column 10. This column reduces the marginal utility of high-demand states relative to low-demand states at a rate similar to what would seem to be implied by the midpoint of the range of estimates of Finkelstein et al. (2013) (acknowledging the difficulty of adapting their estimates to our setting; see Appendix Section A.6.3). Even with state dependence based on the top of Finkelstein et al.'s (2013) range of estimates, the optimal subsidy is large and improves welfare substantially relative to a pure-cash benefit.

The welfare gain from 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. 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 in states 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 fraction of retirees. It should be noted that while the first-best contract is dominated by no insurance

from the perspective of someone eligible (or potentially eligible) for home care, the first-best contract is better from the perspective of society as a whole. From the latter perspective, 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. The key role of the consumption floor is also apparent in column 13, which reduces the consumption floor from $5,000 to $2,500 per year. This increases the welfare gain of the optimal policy dramatically, from $5,528 to $14,194—over two times the cost of the optimal policy. The role of the consumption floor can also be seen in the results of additional specifications, not reported but available upon request, that vary the level of income. The higher is income, the smaller are the effects of the consumption floor, since the individual relies on it in fewer states of the world. When income is $30,000 instead of $15,000 per year, the gain from the optimal subsidy over a pure-cash policy increases from $5,528 to $16,035.

Columns 14 and 15 report results based on different assumptions about the θ distribution. If states in which a person consumes more than 50 hours per week of care are dropped (column 14), the optimal subsidy is 59 percent. If all of the θ values are cut in half (column 15), the optimal subsidy is 75 percent. These results partially address 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 lead 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.

Appendix Figures and Tables

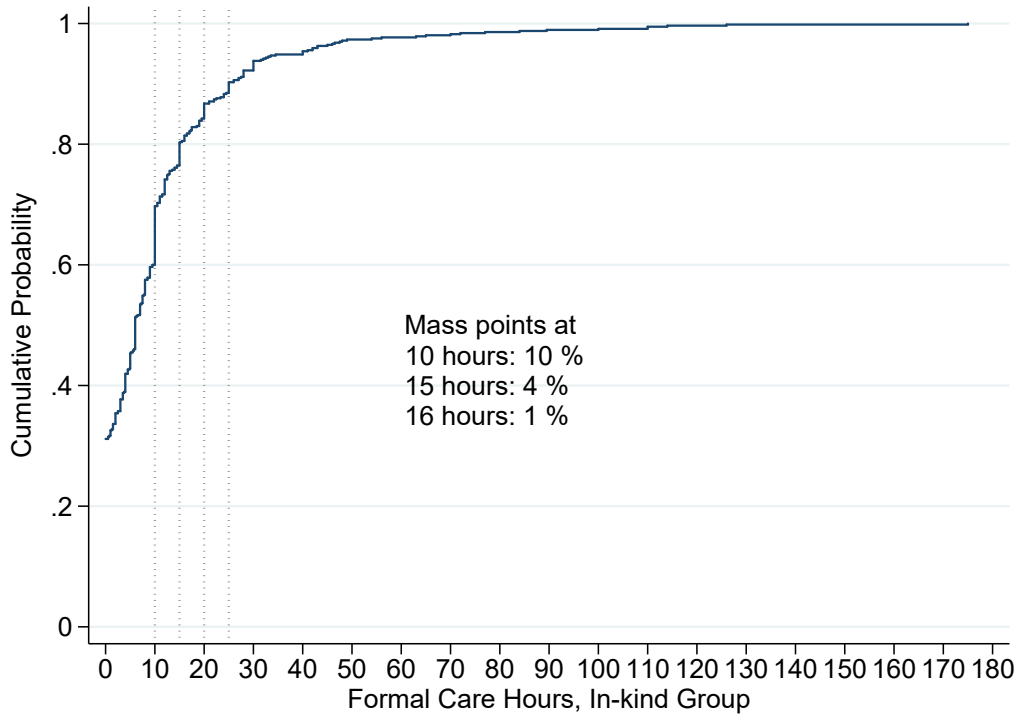


Figure A.1: CDF of Formal Care Consumption in Cash and Counseling, Arkansas

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

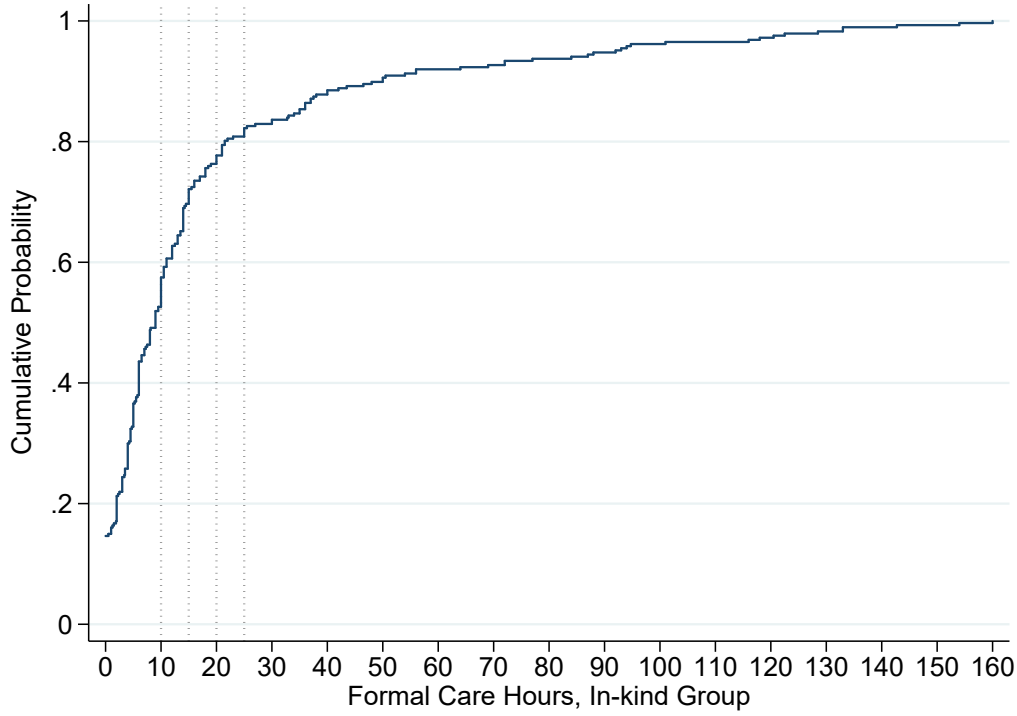


Figure A.2: CDF of Formal Care Consumption in Cash and Counseling, Florida

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

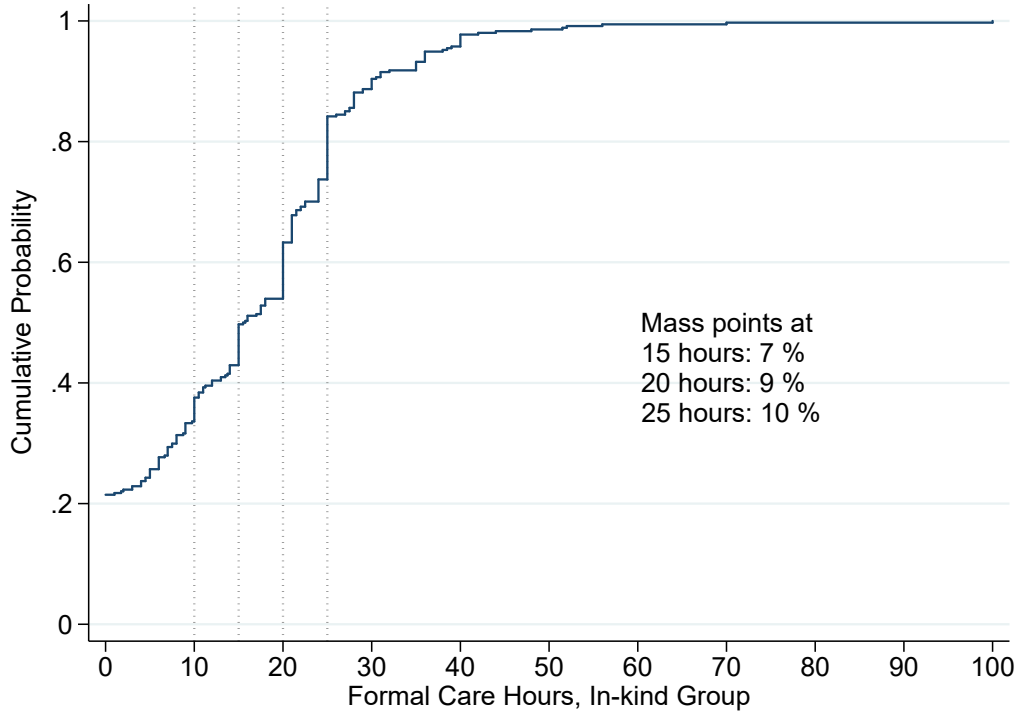


Figure A.3: CDF of Formal Care Consumption in Cash and Counseling, New Jersey

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

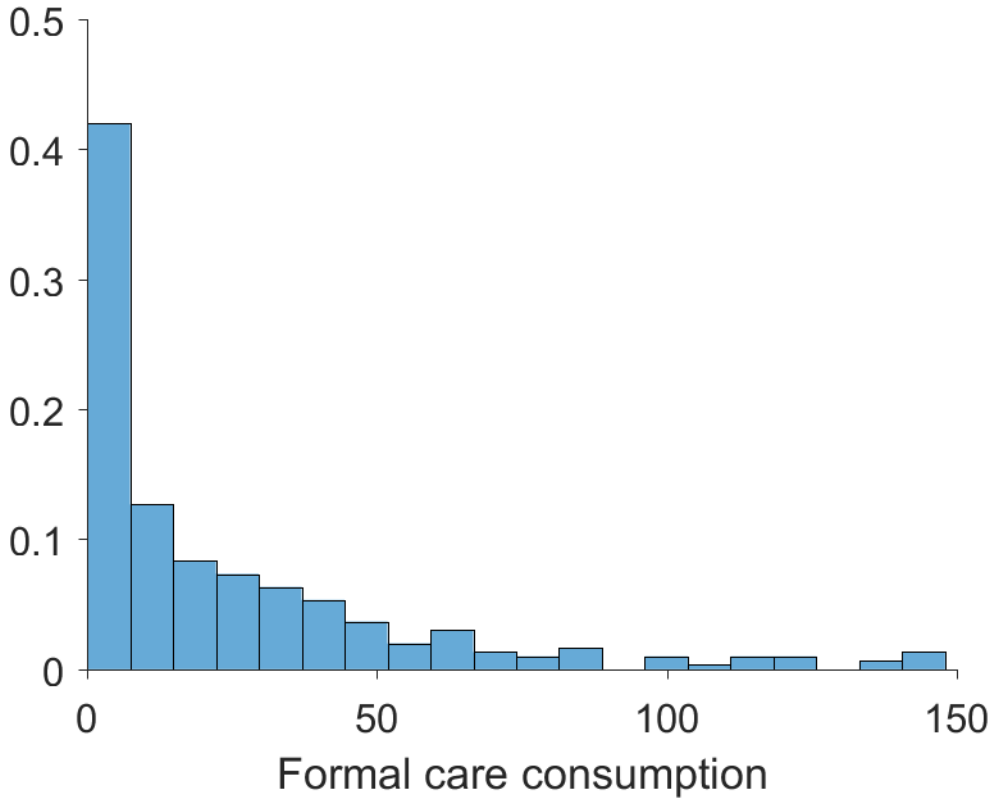


Figure A.4: Distribution of Formal Care Consumption in Benefit-Eligible Population

[Empirical density of formal care consumption among the non-institutionalized population aged 65 and older with two or more ADL limitations. Data from the NLTCS. For readability the figure omits the 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. The mean of the full distribution is 12 hours per week.]

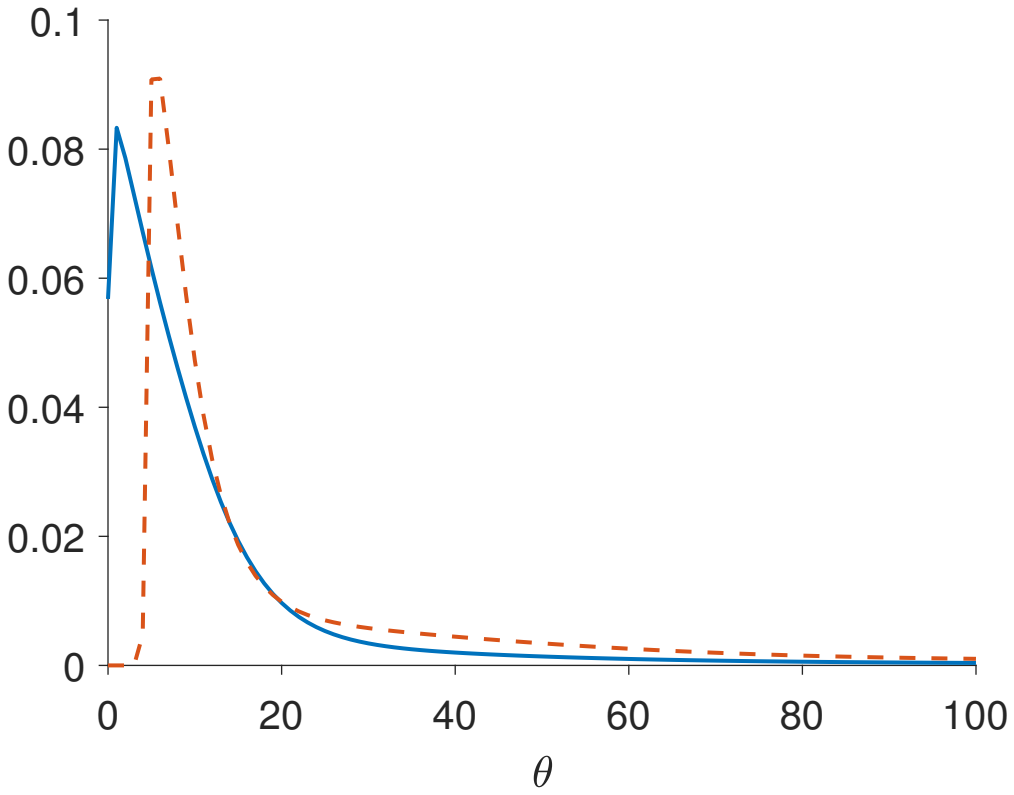


Figure A.5: Distribution of Demand for Formal Care by Whether Someone Lives Alone

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

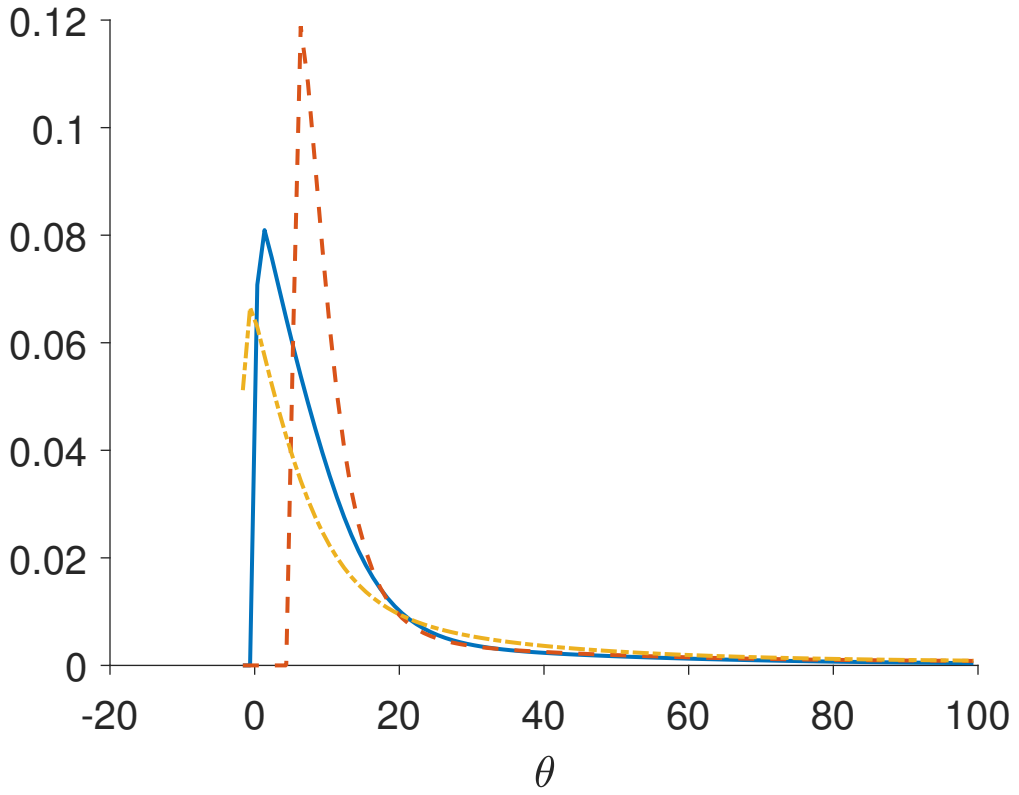


Figure A.6: Distribution of Demand for Formal Care by Number of ADL Limitations

[Estimated probability density functions of formal care satiation points, θ, for each of three groups within the home care benefit-eligible population: 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 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 A.1: Summary Statistics and Balance Tests for the Cash and Counseling Experiments

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

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

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

	NLTCS					Cash and Counseling
	(1) All 65+	(2) 2+ ADLs	(3) Eligible for Medicaid home care	(4) 2+ ADLs, not on Medicaid home care	(5) 2+ ADLs, on Medicaid home care	(6) All 65+
Formal care hours/week	2.87	12.44	15.89	9.19	37.93	12.04
Age	77.43	79.14	79.97	79.18	79.62	78.76
Number of ADLs*	0.70	3.66	3.66	3.56	4.15	2.11
Health fair or poor	0.41	0.70	0.70	0.68	0.77	0.79
Female	0.63	0.66	0.70	0.65	0.74	0.82
Lives alone	0.36	0.25	0.30	0.22	0.41	0.32
Unmarried	0.53	0.55	0.58	0.52	0.70	0.81
Has children	0.81	0.78	0.74	0.80	0.78	.
Income ($/month)	984.13	897.89	659.45	931.03	651.88	.
Observations	5,147	887	493	745	104	2,470

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

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

Table A.3: NLTCS Variables Used to Predict Formal Care Consumption in Table 1

Category	Number	Examples
Demographics and finances	392	Age, sex, race, and other demographics; family structure; insurance coverage; individual's and spouse's income and assets by source.
Formal care prices	7	Whether on Medicaid, whether on Medicaid home care; Formal home care prices in the individual's state: minimum, maximum, and 25th, 50th, and 75th percentiles
ADL limitations	159	Whether individual is physically able to perform tasks on her own such as bathing, walking, dressing, toileting, brushing her teeth, or eating; whether individual uses a cane, orthopedic shoes, a wheelchair, or other special equipment for these tasks; how long the individual has needed help with these tasks; whether individual receives help with these tasks.
IADL limitations	127	Whether individual is physically able to perform tasks on her own such as cooking, driving, keeping track of finances, managing medications, shopping, or using the telephone or computer; whether the individual receives help with any of these tasks.
General health	161	Whether individual has or had within the past twelve months conditions including, but not limited to, heart attack, bronchitis, and Alzheimer's; tobacco and alcohol use; battery of mental health questions (evaluated in part via memory tasks); how often the individual talks to or interacts with family and friends.
Informal care	100	Number of informal caregivers; relationship of informal caregiver to individual; how far away (miles) children are from individual; which ADL- or IADL-related tasks informal caregiver helps with.

Variables from the 1999 NLTCS dataset. ADL is activities of daily living. IADL is instrumental activities of daily living. Some questions are asked of both the respondent and any helpers, so the raw variable counts overstate the (substantial) breadth of coverage.

Table A.4: Price Sensitivity of Demand for Formal Care, First Stage Estimates

	(1)	(2)
Assigned to near-cash	7.66***	7.62***
	(0.23)	(0.23)
Controls	No	Yes
F-Statistic	1,127	1,130
Mean market price	13.68	13.68
Adjusted R-squared	0.32	0.33
Observations	2,440	2,440

Dependent variable is the marginal price of formal care. Data are from the Cash and Counseling experiments. Controls included in column 2 are indicators for sex, education level, race, self-rated health, five-year age bins, and state. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

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

	Censored models (observe $q_i = \max\{0, q_i^*\}$)			Uncensored models (assume $q_i^* = q_i \geq 0$)		
	(1) Normal	(2) Extreme Value	(3) T-location scale	(4) Normal	(5) Negative binomial	(6) Poisson
Price	-1.80***	-2.19***	-1.22***	-0.96***	-0.58***	-1.08***
	(0.13)	(0.22)	(0.07)	(0.08)	(0.11)	(0.20)
Mean hours	10.48	10.48	10.48	10.48	10.48	10.48
Observations	2,440	2,440	2,440	2,440	2,440	2,440

Dependent variable is formal care consumption in hours per week. Data are from the Cash and Counseling experiments. Each column presents the estimated price sensitivity of demand under a different distributional assumption on the underlying error term. Column 1 is the baseline specification. Columns 1-3 assume, as in the baseline, that observed consumption is censored to be non-negative. Columns 4-6 assume that everyone with $q_i = 0$ has a marginal value of care of exactly p_i, the maximum consistent with their behavior (i.e., no censoring). All models instrument for price with the participant's randomized treatment status and are estimated via two-stage residual inclusion. Columns 5 and 6 report average marginal effects. * p<0.10, ** p<0.05, *** p<0.01

Table A.6: Price Sensitivity of Demand for Formal Care and Statutory Limits

	(1) Arkansas	(2) Florida	(3) New Jersey
Price, IV Tobit	-1.04***	-2.78***	-1.61***
	(0.22)	(0.42)	(0.15)
Price, IV Tobit Limits	-0.53***		-1.78***
	(0.12)		(0.16)
Controls	Yes	Yes	Yes
Market price, formal care	12.36	15.09	14.59
Mean hours, in-kind group	10.76	18.60	16.10
Observations	1,129	589	722

Dependent variable is formal care consumption in hours per week. Data are from the Cash and Counseling experiments. Seperate regressions are run for each state. First row is IV Tobit. Second row is IV Tobit with statutory limit as upper bound. There is no statutory limit in Florida. All regressions control for sex, education level, race, self-rated health, five-year age bins, and state. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table A.7: Level of Demand for Formal Care Among Those Who Do Vs. Do Not Take Up Medicaid Home Care

	(1) OLS	(2) 90th	(3) 95th	(4) 99th
Medicaid home care	15.70**	14.20	34.90	75.34*
	(7.23)	(18.48)	(23.52)	(40.56)
Age	0.58***	-0.00	-0.00	2.81*
	(0.19)	(0.50)	(0.78)	(1.46)
Four or more ADLs	11.93***	32.00	70.00***	18.26
	(3.86)	(20.58)	(24.13)	(29.03)
If health fair or poor	-1.81	-1.00	-0.00	34.32
	(3.94)	(8.91)	(14.01)	(22.27)
Female	1.22	-0.00	0.00	-43.73*
	(4.15)	(5.82)	(18.55)	(23.60)
Lives alone	11.17*	55.00**	42.00*	5.40
	(6.06)	(26.79)	(24.27)	(34.52)
Unmarried	8.22*	20.00*	28.00	81.06***
	(4.68)	(11.00)	(22.98)	(30.58)
Has children	5.41	7.00	-0.00	23.89
	(5.40)	(12.98)	(14.11)	(15.51)
Income	0.00	0.00	0.00	-0.01
	(0.01)	(0.01)	(0.02)	(0.01)

Dependent variable is price-adjusted formal care consumption, in hours per week. Price-adjusted formal care consumption uses our estimate of price sensitivity to simulate each individual's consumption if she faced a price of $18.50 per hour, the maximum in the data. The sample is those eligible for Medicaid home care, based on the "Income eligible, < 2 cars" measure. The sample has 481 observations. Column (1) reports results from an OLS regression with robust standard errors. Columns 2-4 present results from quantile regressions, with the quantile specified in the column heading, with bootstrapped standard errors. * p<0.10, ** p<0.05, *** p<0.01

Table A.8: Targeting in the Cash and Counseling Experiments, Arkansas

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

Each entry is the estimate of β_3, the coefficient on $inkind_i * X_i$, in equation 3, described in Appendix Section A.5, of separate regressions. The dependent variable is the weekly cost of benefits received in dollars. The estimates reveal whether people with more of the characteristic receive differentially greater transfers in the in-kind group (relative to the near-cash group) than do people with less of the characteristic. Data from the Arkansas Cash and Counseling experiment. Columns 1 through 4 include all participants. Columns 5 through 8 only include the subset who had not been enrolled in Medicaid home care before the experiment. Columns 1 and 5 are OLS regressions with robust standard errors. Remaining columns are quantile regressions with bootstrapped standard errors. The omitted health category is health good or excellent. * p<0.10, ** p<0.05, *** p<0.01

Table A.9: Regressions of Formal Care Use on Individual Predictors, NLTCS

	(1) Coefficient	(2) R-squared
Lives alone	15.57***	0.037
	(2.65)	
Number of ADLs	4.39***	0.034
	(0.78)	
Age	0.76***	0.032
	(0.14)	
Unmarried	12.17***	0.030
	(2.33)	
Female	6.16**	0.007
	(2.47)	
Has children	2.07	0.001
	(2.36)	
Self-rated health fair or poor	-0.08	0.000
	(2.53)	
Income	0.00	0.000
	(0.00)	

Each entry is from a separate regression of formal care on the variable specified in the table row. Coefficients and standard errors from this regression are presented in the first column; the R-squared is presented in the second column. Data are from the NLTCS. Sample is non-institutionalized individuals aged 65 and older with two or more ADL limitations. There are 887 observations. Omitted category for self-rated health is health excellent or good.