

Subsidies, Information, and the Timing of Children’s Health Care in Mali

Anja Sautmann*, Samuel Brown† and Mark Dean‡

June 2016

Abstract

We study how healthcare subsidies and improved information affect over- and under-use of primary healthcare in a randomized control trial of 1544 children in Mali. In a dynamic model of healthcare demand, misuse relative to policymaker preferences (here given by WHO care-seeking standards) arises from seeking care too early or too late during an illness spell. Using nine weeks of daily data, we show that the barrier to optimal care seeking is cost, not information: subsidies increase demand by over 250%, but overuse is rare with or without the subsidy. Information, contrary to intent, appears to increase underuse, as our model predicts.

*Brown University, corresponding author, email: anja_sautmann@brown.edu.

†Brown University, email: samuel_brown@brown.edu.

‡Columbia University, email: mark.dean@columbia.edu.

We have received invaluable advice and support from Andrew Foster, Duncan Thomas, and Ira Wilson. We would also like to thank Anna Aizer, Dan Björkegren, Caitlin Cohen, Jishnu Das, Seydou Doumbia, Chris Gill, Jessica Goldberg, Sanjat Kanjilal, Bentley McLeod, Emily Oster, Simone Schaner, Dan Silverman, and Dean Yang, and participants of the “Natural Experiments and Controlled Field Studies” workshop 2015. Special thanks go to Hamadoun Bocoum, Zorina Curry, Judith Kom, and Fred Wensing, Mali Health and its staff and Board, and Innovations for Poverty Action. We are grateful for funding from Brown University through the Brown Seed Fund Award, the Population Studies and Training Center, and the Rhodes Center, from the Aga Khan Foundation, and from the ESRC/DFID Development Frontiers Award (ES/K01207X/1). All errors are ours.

1 Introduction

Even though child health has improved greatly and child mortality rates have more than halved since 1990, six million children under five years, around 16,000 per day, died in 2015. An important cause is a lack of adequate care for acute illness. For example, less than 26% of malaria cases in children are estimated to receive adequate treatment, yet at the same time, 411,000 children died of malaria in 2013 in Africa alone. Some estimate that improved access to primary care could prevent 29%–40% of post-neonatal deaths in developing countries.¹

The provision of primary healthcare for children of poor families is subject to a policy debate centered on whether healthcare should be funded via subsidies, or via fees paid by patients at the point of use.² The standard economic viewpoint is that, absent other distortions, subsidies are likely to cause overuse and waste precious resources that are needed elsewhere. On the other hand, several large aid organizations advocate subsidized care as a means to offset barriers to healthcare access, which may arise e.g. from poor health knowledge or credit constraints (e.g. UK Secretary of State for International Development [2009], Save the Children [2008]). A number of African countries have introduced (partially) free healthcare for mothers and children in an effort to reduce high mortality and morbidity rates (Ridde and Morestin [2012], Yates [2007]), most recently Burkina Faso in April 2016. Many times, complementary policies focus on improving health education and information of parents in order to encourage more efficient use of healthcare resources.

Even while these efforts are underway, it remains an open question if subsidies, possibly supplemented by information policies, can curb any underuse of care, without creating overuse. Central to this issue is the definition and measurement of over- and underuse. Traditional demand analysis of price changes is uninformative about the value of care in the presence of distortions, for example if credit constraints create barriers to access, parents have poor information about their children’s healthcare needs, or healthcare has important externalities. In this paper we therefore take a different approach: we propose a new model of healthcare timing to frame the analysis, and then use daily health data from a randomized control trial of 1544 children in Mali to estimate the effects of subsidies and information on healthcare use, *relative to the preferences of an outside policymaker*. We derive these preferences from WHO guidelines for children’s healthcare that are in use in over 80 countries (Bryce et al. [2004]).

¹United Nations Inter-Agency and Expert Group on MDG Indicators [2015], WHO Global Malaria Programme [2015], Bhutta et al. [2008].

²See Akin et al. [1987], Litvack and Bodart [1993], Jimenez [1995].

The main challenge in measuring healthcare misuse arises because the value of acute healthcare depends on the patient’s health status: an identical increase in utilization could mean that sick children are getting timely access to desperately needed care, or that healthy children are being dangerously overmedicated. What is more, we argue that efficient care seeking depends on the timing of care *during* an ongoing illness spell: there may initially be a good chance that the child will get better without medical care, and only when symptoms persist it becomes likely that the child is suffering from a serious illness and will not recover on her own. We therefore propose a formal framework that characterizes care-seeking as an optimal stopping problem: the decision is not just if, but *when* the expense of a doctor visit is warranted.³ Within this framework, parents and policymaker may disagree on the optimal timing of care for two reasons. First, they may assign different costs and benefits to treatment, for example due to credit constraints or externalities. Second, parents may judge the probability that the child will recover without care differently from the policymaker, and in particular, they may not have the medical knowledge to interpret the child’s symptoms. These disagreements may lead to *underuse* in the form of *inefficiently late* care-seeking, or to *overuse* in the form of *seeking care too soon*, relative to the policymaker’s preference.

In this model, subsidies reduce the cost threshold for care-seeking, leading to increased demand for care and earlier doctor visits. Depending on the policymaker’s preferences, this may translate into a reduction of underuse, but also an increase in overuse. Providing parents with health information, on the other hand, has the potential to reduce *both* overuse and underuse, by aligning the parents’ beliefs about recovery chances with those of the policymaker. However, since teaching parents to spot serious illness automatically also teaches them to discern when an illness is less serious than they previously believed, better information can paradoxically *increase* underuse if parents apply a higher cost threshold than the policymaker. This lends support to the policy argument that information provision and subsidies are complements: only when both cost thresholds and beliefs are aligned will parents make optimal choices according to the policymaker preferences.

In order to test the predicted effects of subsidies and information on care seeking behavior, we collaborated with the NGO Mali Health to carry out an RCT of their Action for Health program, which removes (certain) user fees (the “free care” treatment) and provides free community healthworker visits (the “healthworker” treatment) for children

³This contrasts with the workhorse model of healthcare demand in development economics, where a health production function translates expenses into a health capital stock (see Akin et al. [1986], Foster [1995], Gertler et al. [1987], Gertler and Van der Gaag [1990], Sahn et al. [2003] for examples).

under five years of age in Bamako, Mali. Our design addresses selection bias and endogeneity problems: first, the random assignment into treatment arms means that we observe an unbiased demand response from families, and second, the quality and cost control checks that Mali Health conducts at their partner clinics prevent offsetting supply effects, for example price increases or a deterioration in the quality of care.

We randomly assigned our sample population to one of three treatment groups – healthworker visits, free care, or both combined – and a control group. The free care arm was provided through two local clinics that are the main providers of formal health care for our population and collaborate with Mali Health. The healthworkers improve information by checking the child’s health and encouraging the use of formal healthcare according to WHO and UNICEF recommendations, codified in the Community portion of the Integrated Management of Child Illness guidelines (C-IMCI, WHO [2014], WHO Department of Child and Adolescent Health and Development [2005], Unicef Health Section [1999], Rosales and Weinbauer [2003]).

Six months after program roll-out, we carried out a nine-week panel survey of the 1544 study children, collected weekly in order to overcome recall bias (Das et al. [2011]). We recorded daily information on 14 symptoms, as well as data on all care sought and treatments received. This enables us not only to group illness days into spells and observe care seeking conditional on the child’s health status, it also allows us to apply the C-IMCI standards to assess *when* a child should see a formal care provider. The IMCI is the global standard for children’s healthcare in resource-poor settings and represents the preferences of both the implementing NGO and the WHO. The C-IMCI handbook describes common and easy to spot symptoms, classifies them into sets of symptoms corresponding to illness groups, and determines after how many days of a given symptom profile the child should be seen by a doctor. This allows us to classify each spell day we observe as ‘care required’ or ‘early’ for care.

Using our data we first estimate the unconditional demand effects of our interventions. We find that subsidies reduce the cost of a doctor visit for the family by about 65%, increasing demand from 0.18 to 0.57 formal consultations per child over the 9 week period. Private expenditure on care does not change, meaning that none of the subsidy is diverted into non-healthcare consumption, and the total value (cost) of care the children receive nearly doubles. In fact, we observe additional significant increases in non-monetary costs in the form of waiting times. The healthworker visits, on the other hand, do not have strong average demand effects.

These unconditional effects are good news for a policymaker interested in raising uti-

lization cost-effectively, but they do not tell us if the additional care goes to the right children at the right times. Thus, we also report the effect of the treatments on important over- and underuse indicators. We find that underuse is rampant in the control, while overuse is rare: only 11% of spells that reach “care required” status eventually receive formal care, while a consultation occurs in 3% of spells that never leave “early” for care status. Subsidies reduce the proportion of care-required spells that do not receive care from 89-90% to 69-73%, while they increase utilization to 6-8% of “early” spells. Thus, they reduce, but do not eliminate underuse, with relatively small effects on overuse: with or without free care, around 85% of all consultations occur on a day when care is required according to the C-IMCI. Healthworkers do not improve this allocation substantially by reducing overuse or underuse; rather it appears that they may increase underuse.

The above measures of misuse arise from the parents’ choices, combined with the characteristics of the illness spells, including the probability that a given spell recovers without formal medical treatment. In order to analyze the underlying changes in behavior that drive the observed outcome changes, and guided by our dynamic model of healthcare timing, we estimate the probability of care-seeking conditional on individual and spell characteristics in a (proportional and non-proportional) hazard model. The estimates show that the probability of care seeking is near zero on any given “early” day, even under the subsidy. Within each treatment group and spell day, the hazard of care is higher if care is required according to the C-IMCI; nevertheless, in the control, it reaches at most 6% probability. The subsidy increases care-seeking probabilities by more than 250%, which leads to a much stronger increase on care-required days in absolute terms. By contrast, we find that the healthworkers have a *negative* effect on care-seeking on care-required days (significant at the 10% level). Our point estimates suggest that this effect occurs mostly in the healthworker-only group rather than the full program group. The healthworkers also do not prevent any overuse, because it is so rare to begin with.

Our results demonstrate that the binding constraint to efficient care seeking in our population is financial, not informational. Parents seek care significantly more often and respond more strongly to the subsidy on care-required days. Subsidies do not lead to overuse, but also do not eliminate underuse, likely due to residual costs of care-seeking. The adverse healthworker effect points to an underappreciated problem when providing financially constrained parents with additional health information, as predicted by our model: if families learn from the healthworkers how to tell when their child is not in immediate danger, and they disagree with the policymaker on the need for care in less severe cases of illness, underuse may actually increase.

Our estimation approach has two important advantages. First, the possibility of spontaneous recovery is integral to the decision to seek acute healthcare, and in our sample, more than 85% of illness spells end without the use of formal care. This is a form of censoring, because the choice of when to seek care is effectively not observed. A hazard model can appropriately account for this, and we show that as a consequence of censoring, average spell length responds much less to the subsidy treatment than the probability to seek care. We also show that this problem would be compounded if we were to analyze only uncensored spells, for example from data collected at the healthcare provider: we would falsely conclude that subsidies *delay* access to care, due to the selection of families into care who tend to wait longer. A second advantage is that the hazard model can flexibly condition on underlying health status, which allows us to make out-of-sample predictions for healthcare demand in other populations and disease environments. We demonstrate this by showing how care-seeking would change as a result of the subsidy and information policies in the event of an Ebola outbreak.⁴ This could be used for example to predict the effect of our interventions on the spread of the illness.

Last, while our study was not designed to estimate long-run health outcomes, we show that better access to care in the free-care arms of our sample had a small but significant effect on average illness spell length, and reduced mothers' concern about their children. This suggests real short-term health effects as well as a potential for long-term health improvements.

Within the literature on healthcare demand in developing countries, our data is to our knowledge the first of its kind. The one other spell data set we are aware of is the 1987 National Medical Expenditure Survey for the US, which only includes spells that lead to medical care use or work absenteeism (Gilleskie [1998] addresses the resulting selection problem by estimating the probabilities of unobserved health events parametrically).⁵ Early research in development economics, instead of estimating demand directly, has focused on the effects of user fee changes on health, labor market participation, and income (Gertler et al. [1987], Gertler and Van der Gaag [1990], Dow et al. [2001], Nabyonga et al.

⁴By combining Ebola spell descriptions from the public health literature with different hazard model specifications.

⁵With administrative data e.g. from insurance claims (see Manning et al. [1987], Finkelstein et al. [2012]), illness spells can only be studied indirectly by grouping claims into treatment episodes, see Stoddart and Barer [1981], Hornbeck et al. [1985], Keeler et al. [1988], Santos Silva and Windmeijer [2001]. Recent work in this context focuses on specific data challenges like censoring and endogeneity, e.g. Kowalski [forthcoming]. So-called two-part models deal with unobserved health shocks by accounting separately for the decision to use care and the amount spent conditional on seeking care (see e.g. Jones [2000]). Note, however, that there is a public health literature which studies the covariates of the delay in seeking care for specific diseases (see Storla et al. [2008] and Nguyen et al. [2010] for reviews).

[2005], McIntyre et al. [2006]).

A more recent literature studies healthcare utilization in field experiments. Most relevant here are Cohen et al. [2015], who examine the effect of providing subsidies for malaria combination therapy (ACTs) and information about health status in the form of subsidized malaria tests at pharmacies. Like us, they find that better access to information (testing) does little to change healthcare usage; but in contrast with our results, there are significant rates of overtreatment that increase with subsidization, at least within their sample of people who come to the pharmacy with intent to buy. Their data was not designed to study underuse, although they can infer some effects indirectly.⁶

We are also aware of three randomized trials on health insurance subsidies, in Nicaragua, Mexico, and Ghana (Thornton et al. [2010], King et al. [2009], Powell-Jackson et al. [2014]). The programs in Latin America had few measurable health or utilization effects, and uptake and retention were low. The trial in Ghana shows some increase in average utilization, but the price elasticity of healthcare demand appears to be small; however, there are significant positive health outcome effects in the subpopulation of children who were anemic at baseline. An explanation for the low demand response in these trial could simply be that the populations studied are relatively healthy and rarely reach “care required” status. Alternatively, there is underuse that the interventions do not remedy, possibly due to unobserved supply-side effects. Our study highlights that we cannot know which is the case without measuring healthcare demand conditional on health status.

Section 2 introduces a dynamic model of care-seeking during an illness spell, defines overuse and underuse, and discusses possible effects of subsidy and information policies. Section 3 describes the policy environment and the RCT intervention and survey design. Section 4 first shows how unconditional demand as well as overuse and underuse outcomes respond to the subsidy and information policies, and then estimates a hazard model of demand behavior, used to assess our model predictions and to project healthcare demand in a hypothetical Ebola outbreak. Section 5 concludes.

⁶Also related are Ashraf et al. [2010], Dupas and Cohen [2010] and Dupas [2014], who test targeting, utilization and behavioral “sunk cost” effects when preventive goods are subsidized. Ashraf et al. [2013] show that better information can increase the purchase response to a subsidy for water disinfection. Preventive goods pose a limited risk of overuse and their value does not depend on demand shocks, so they lend themselves more readily to traditional demand analysis. Fischer et al. [2014] test for learning and reference dependence in the demand for different over-the-counter drugs, but do not focus on over- or underuse. More broadly there has been a surge in randomized studies on healthcare use and quality and the resulting health effects. Examples are Miguel and Kremer [2004], Das et al. [2015], Currie et al. [2011], see also Kremer and Glennerster [2012].

2 A Dynamic Model of Demand for Healthcare

We analyze the impact of subsidies and information policies, and define what constitutes “overuse” or “underuse” from the perspective of a policymaker, in a model that is based on two core ideas. First, demand for acute care arises in response to a negative health shock that causes discomfort and possible harm to the child.⁷ Second, there is learning about the severity of an illness over time. Initially, it is likely that the child will recover without the aid of a clinic visit, and it is worth tolerating some disutility from sickness while waiting. As time goes by, the probability of recovery without treatment decreases, and further suffering can only be avoided by seeking care. Thus, the demand decision is *when* during an illness spell to visit a clinic. Overall demand can increase as a result of seeking care earlier, if this means that the consultation happens more often before the child recovers on her own. Misallocation in this context occurs if parents seek care too early or too late relative to the policymaker’s preferences.

2.1 Basic Model

We focus our exposition on the main properties and insights from the model, relegating a full formal treatment with proofs to appendix A.

Consider a child in an ongoing illness spell with symptoms $\gamma \in \Gamma$. Let $t \geq 1$ denote the number of days the child has been sick. Parents experience sickness disutility $-S < 0$ every day their child is ill. On any day, the parent can take the child to the doctor which (for simplicity) leads to immediate recovery. The expected utility cost – expenses as well as travel and opportunity cost – of a formal consultation is $-C < 0$.⁸

Symptoms γ can be caused by one of a set of possible underlying conditions $\{I_1, \dots, I_N\}$, each with different probability of spontaneous recovery $(1 - \pi(I_n))$. The expected probability of remaining ill in period t , π_t , depends therefore on the parents’ belief about the possible underlying illnesses causing γ . As the illness spell continues, parents become more concerned about their child, because Bayesian updating implies that illnesses with

⁷The implicit assumption is that there is some (minimal) effort cost and no benefits to an acute care visit besides relieving illness.

⁸For ease of notation we suppress the dependence of C and S on symptoms γ in this section. The analysis does not change substantially if illness disutility and doctor costs change over time, as long as $\frac{C}{\delta(C+S)}$ is weakly decreasing. In reality, treatment may also be unsuccessful, for example due to low quality of care. This matters for the value of a doctor visit, which is contained indirectly in S ; but in addition it introduces the possibility of repeat visits. We did not account for this, because it adds complexity to the model without changing the basic conclusions about parental behavior (see also next section).

high recovery probability are less and less likely:

Lemma 1. *Denote by $\pi = \{\pi_t\}_1^\infty$ the sequence of beliefs over time that the child will not recover. The probability of spontaneous recovery declines over time, so that*

$$\pi_t \leq \pi_{t+1}.$$

The crucial point of lemma 1 is that there is learning about illness severity during the spell, and therefore it can be optimal to delay care. Based on their current beliefs, parents decide between consulting a doctor today and incurring cost C in exchange for certain recovery, and waiting another day. With probability $1 - \pi_t$ the child recovers anyway, thus saving a costly doctor visit, but with probability π_t , the child remains sick and disutility S is incurred, and moreover, future recovery also becomes less likely. It can be shown that a solution to this dynamic decision problem is based on a simple cutoff strategy in beliefs:

Proposition 1. *An optimal strategy is to seek formal healthcare if and only if*

$$\pi_t \geq \frac{C}{\delta(S + C)} = K.$$

In other words, parents wait until the probability of remaining ill rises above some threshold K before seeking care. This threshold is increasing in the utility cost and decreasing in the benefit of treatment. If there was no learning, π_t would be constant, and the choice to seek care would be a simple yes/no decision on the first day of illness. Note that a low quality of care or a probability of treatment success below one would have similar effects as a decrease in S , by reducing the value of receiving treatment (see footnote 8).

2.2 Comparison with the Social Planner

In order to conduct welfare analysis based on the observed price elasticities, we would have to assume that there are no barriers to optimal care seeking, like credit constraints and poor health knowledge, and no differences between the private and social optimum, for example due to externalities or an altruism gap between parents and children. Our approach is instead to compare individual behavior with the recommendation of an external policymaker recommendation, like the WHO (see section 4 below). While all the usual

caveats for welfare implications apply, this provides us with an observable benchmark for efficiency.

The policymaker solves a similar problem as the parent⁹ yet may have a different optimal policy for two reasons. First, she may have a different threshold K^* . For example, the parents' cost of treatment C may differ from the policymaker's C^* if the parents are very poor and credit-constrained, so that the utility cost of taking the child to the doctor is very high (perhaps because it means the family cannot eat that day). Similarly, parents' evaluation of benefits S of curing a disease may not account for the positive externalities through reduced infection rates or the long-run human capital effects for the child that are incorporated in the policymaker's S^* .¹⁰

Second, disagreement may occur if parents cannot interpret the child's symptoms. In order to capture this, we assume that the policymaker can differentiate sub-cases $\{\gamma^{*i}\}_{i=1}^M$ of γ with distinct probability distributions over possible illnesses (with the beliefs of the policymaker and parents connected via Bayes' rule). We denote as π_t^{*i} the beliefs of the policymaker at time t after having observed symptoms γ^{*i} .

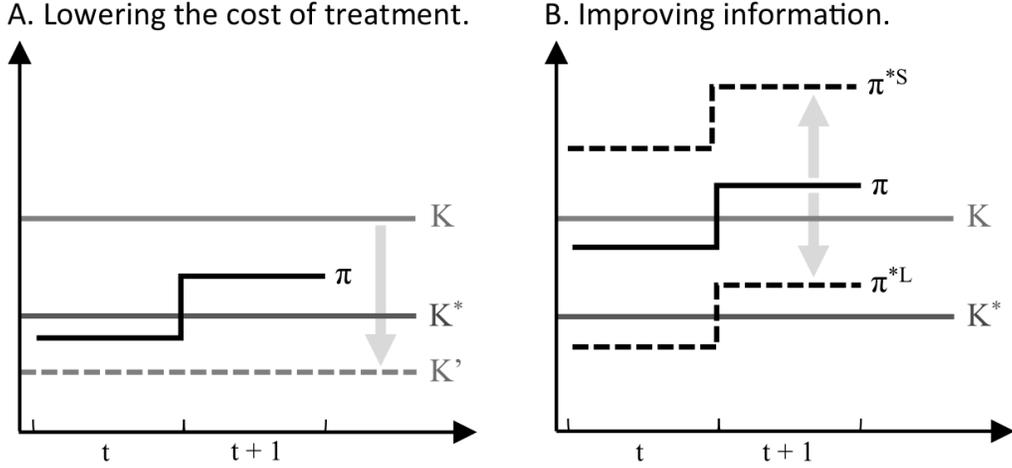
For the remainder of the paper, we define over- and underuse in relation to the policymaker's preferences. "Overuse" occurs if the policymaker considers the chance of spontaneous recovery high enough to warrant the disutility of waiting when the parent does not. "Underuse" occurs when, to the policymaker, a clinic visit is warranted but the parent is not yet prepared to go. Given the optimal cut-off strategy, this implies that overuse leads parents to seek care *too early* within the illness spell, while underuse means they seek care *too late*, relative to the policymaker preferences. Note that this definition is based on the *ex-ante expected benefits and costs of care*. Ex post, it is possible that a child who got care too early would have required care a few days later anyways; similarly, a child who did not receive care when it should have may nevertheless recover shortly after. This does not change the fact that the care-seeking decision was suboptimal from the policymaker perspective at the time of decision-making.

Within our model, we interpret the two interventions of Action for Health as follows. The "free care" component *lowers the care-seeking threshold* of parents to $K' < K$ by reducing the treatment cost C . The healthworkers teach the families to differentiate

⁹Medical guidelines commonly take into account the length of a given set of symptoms as an indicator for treatment needs. Longer duration of a spell typically indicates greater severity. This cannot be explained without learning about persistence – if, for example, all illness lasts the same amount of time, earlier treatment is always better.

¹⁰A similar effect would occur if parents underestimate the value of treatment, for example because they believe that it has little effect on the probability of recovery (see footnote 8). This is particularly a concern given that healthcare is an experience good that is difficult to evaluate *ex ante*.

Figure 2.1: Parents seek care in t if $\pi_t > K$, the policymaker optimum is to seek care if $\pi_t^* > K^*$.



between different subsets of symptoms $\{\gamma^{*i}\}_{i=1}^M$, thereby *aligning the beliefs of the parents with those of the policymaker*.

We demonstrate the effects of the two policies with a simple example (appendix A contains numerical details). Suppose that symptoms γ can be caused by two illnesses, one severe (I_S) and one less severe (I_L), with $\pi(I_S) > \pi(I_L)$. Bayesian updating implies that $\pi_t < \pi_{t+1}$, as illustrated in Figure 2.1 by the solid black line. We assume in what follows that the parents' treatment threshold K is weakly greater than that of the policymaker K^* .

Different cost cutoffs. Panel A shows a parent with a treatment threshold above the policymaker's, $K > K^*$, but where both hold the same beliefs π . The policymaker would like the child to receive care in period $t+1$ (conditional on not having recovered), but the parent will not take the child to the doctor until after $t+1$. This is an incidence of *underuse*.¹¹ Subsidies shift K down, and therefore (weakly) increase utilization in all periods. This means they can remedy some or even all underuse, but if K is reduced by too much, for example to K' , overuse may now occur in period t (since $\pi_t > K'$). In general, since beliefs and treatment thresholds may differ between symptoms and illnesses as well as between families, an untargeted subsidy policy may have heterogeneous effects: it may eliminate some, but not all, underuse, yet simultaneously create some overuse.

¹¹The parent will eventually take their child to the doctor if K is below $\pi(I_L)$, otherwise they will never go.

Different beliefs. Now suppose that the policymaker can additionally interpret the symptoms γ better than the parent. Specifically, she can distinguish symptoms γ^{*L} , which indicate a higher probability of I_L , and γ^{*S} , which indicate that the child is more likely to be seriously ill with I_S . An example is the distinction between simple diarrhea and diarrhea with blood in the stool; the latter is a sign of dysentery and therefore potentially dangerous illness. Bayesian reasoning implies that $\pi^{*L} < \pi < \pi^{*S}$, where the beliefs of the parents π are the average of the beliefs π^{*L} or π^{*S} of the policymaker. Panel B of Figure 2.1 illustrates this.

Misallocation of care is inevitable when parents have bad information, unless the distinction between π^{*L} and π^{*S} is irrelevant for the optimal timing of care. To see this, consider parents with the treatment threshold K in Panel B: they will always seek care in period $t + 1$. However, from the policymaker's perspective, it is *never* optimal to take the same actions under both π^{*S} and π^{*L} . For example, even if the policymaker shared the parents' threshold K , she would want the child to receive care both in t and $t + 1$ under beliefs π^{*S} , but in neither period under π^{*L} . There is therefore underuse when the illness is severe, but overuse when it is harmless. Shifting the treatment threshold via subsidies cannot address both types of misallocation, because it can only monotonically increase or decrease utilization.

Can the healthworker policy correct over- and underuse? The answer depends on the relative cost thresholds. For example, suppose the policymaker shares the parents' K in Panel B. In this case aligning the parents' beliefs with those of the policymaker will lead to perfect agreement, and thus simultaneously eliminate both over- and underuse. By contrast, suppose the policymaker has threshold $K^* < K$. In this case the *uninformed* parent with threshold K takes the policymaker's preferred action in period $t + 1$, but underuses care in period t if the illness is severe. Parents who are able to discern symptoms γ^{*S} from γ^{*L} will seek care in both periods when the illness is severe, and as a result, underuse is reduced under π^{*S} . However, they will seek care in neither period, and therefore underuse is increased, under π^{*L} . In other words, teaching parents to tell apart the most severe cases of illness, e.g. by looking for blood in the stool, can paradoxically *increase* underuse in less severe cases. Better information leads to an unambiguous improvement in the allocation of care only if the cost thresholds of the policymaker and the parent are also aligned.

The above generates different predictions for the effects of our interventions, depending on the importance of cost and information barriers to the efficient allocation of care.

- *Subsidies* increase utilization and lead parents to seek care earlier. They can there-

fore *reduce underuse*, but potentially at the expense of greater *overuse*, especially if parents have *poor information*.

- Providing *information* can *improve the allocation* of care and reduce both over- and underuse if policymaker and parent agree on the treatment threshold. However, it may be counterproductive if it helps parents to discern illness states in which their optimal action differs from that of the policymaker. In particular, *if $K > K^*$, better information may increase underuse* and lead parents to seek care too late.
- If there are both knowledge and cost barriers to efficient care seeking, information and subsidies are *complements* in the optimal allocation of care: *underuse can be unambiguously reduced without increasing overuse* only if π and K are both aligned with π^* and K^* .

This last point is a key motivation to combine healthworker and subsidy policies as the Action for Health program does: free care gives families access health services, while better information ensures that they make efficient use of this benefit.

3 Study Background and Data Collection

The Malian healthcare system builds on a network of community health clinics or *centres de santé communautaires* (CSCOMs). A CSCOM has typically one or two trained doctors on duty, along with a handful of other staff, and sells prescribed medications through an attached pharmacy. CSCOM care is partly subsidized by the government and by private NGOs, but primarily funded by user fees, in accord with the Bamako Consensus from 1981 which advocates self-sustaining, decentralized primary healthcare in West Africa. The public healthcare system is flanked by a private formal sector and a large informal sector.

This study was conducted in Sikoro, a hilly area on the outskirts of the capital, Bamako. The majority of its roads are unpaved, and most dwellings are not connected to the water supply or the sewage system. During the rains from roughly July to October, the incidence of diseases like diarrhea, malaria, and respiratory infection is highest. Mali has very high rates of maternal and child mortality, especially in rural areas. Poor urban areas also often lack basic health services, despite better accessibility. The peri-urban population of Sikoro is representative of one of the fastest-growing demographics in West Africa.

3.1 Study Design

Mali Health started their Action for Health (AfH) program in 2010 in collaboration with the two local CSCOMs in Sikoro. The CSCOMs in this study are financially supported by Mali Health and are subject to regular quality control. Mali Health in particular controlled costs in the free care group (see also section 4.4). This is important because it keeps the supply side constant across treatment groups. We assume that the standard of care provided at these CSCOMs reflects the level of care that the WHO expects when they formulate treatment recommendations for a developing-country context. If they recommend care-seeking, they assume that the child will be evaluated and treated by a medical doctor, as opposed to, say, by an untrained individual or not at all (e.g. if the doctor is actually absent). A separate issue is the sometimes low quality of care in public healthcare systems in developing countries, which may affect how desirable the use of the healthcare system is both from the parents' and the policymaker's point of view (see e.g. Das [2011], Das and Hammer [2005, 2007, 2014], Das et al. [2015], Das and Sohnesen [2006], Leonard and Masatu [2010], Leonard et al. [2002]). The focus of this study is to understand the barriers that prevent families from using healthcare of the typical quality a medical doctor (likely trained and certified in-country) can provide.

The full AfH program is designed to improve the health of children under the age of five. It combines healthworker visits and subsidized care. The healthworker component provides biweekly visits from community healthworkers (CHW). CHWs live in the area and do not have any formal medical education, but are trained by Mali Health to recognize basic signs of illness. They monitor children's health by tracking a set of simple indicators, and they advise families both about when it is necessary to seek formal care in the event of illness, and also when to provide care themselves if an ailment does not (yet) require medical attention – for example through the use of oral rehydration salts to treat mild cases of diarrhea. Their training follows the C-IMCI (Rosales and Weinhauer [2003]), a set of guidelines for community healthworkers that incorporates the WHO's and UNICEF's "Integrated Management of Childhood Illness" recommendations for when to refer children to formal healthcare (WHO [2014], WHO Department of Child and Adolescent Health and Development [2005]; see also below).¹²

The free care part of AfH is administered via a personalized card that entitles the child to unlimited free consultations with a doctor at the two local CSCOMs, and free treatment and medication for any illness due to diarrhea and malnutrition, malaria, vaccine-

¹²Outside the focus of this paper, CHWs also advise households on preventive care, and dispense water chlorination tablets for households with no access to clean water.

preventable diseases, and respiratory infection (the five main causes of child mortality). Any costs covered by AfH are billed directly to Mali Health. For cost control purposes, Mali Health’s medical director has developed treatment guidelines for common diagnoses and conducts spot checks on diagnosis and prescriptions using bills submitted to Mali Health and the clinic’s treatment records and accounting. The program reduces the average visit cost to the families substantially (see section 3.3.2), but they have to cover remaining costs that are not part of AfH. These include for example vitamins or injections and other services that are not part of the standard treatment course for a given diagnosis.

The research design took advantage of the second planned roll-out wave of AfH in late 2012. Mali Health conducted a census in their new expansion area in summer 2012 to enumerate all eligible families based on geography, the presence of children under five years of age (or a pregnant mother), and a set of wealth indicators. These indicators serve as a basic proxy-means test, designed to select about the poorest third of families in the area. Data was collected in two survey rounds in 2012 and 2013 in the rainy season. Households identified by the Mali Health census were revisited for the baseline survey in 2012. All households that were found at baseline were included in the random assignment to the different treatment groups.

Data was collected at the level of the household, defined as all persons who identify the same individual as their household head. In this study we use demographic, location, and household asset data collected in the baseline survey, and daily health diary data on children collected in the follow-up round.¹³ The unit of randomization is the compound. A compound may house more than one household, and typically consists of a few rooms around a common courtyard with shared latrines and other facilities. After stratifying compounds by average household assets, number of eligible children, and location, each was assigned to one of four groups: a full treatment group that received healthworker visits and free care, a free care only group, a healthworker only group, and a control. The healthworkers in the two healthworker treatment groups were trained and managed separately by Mali Health, and no healthworker visited families in both the healthworker-only and the full treatment group, to avoid spillovers on the provider side.¹⁴

¹³Detailed information on the survey design is available on request.

¹⁴It was not possible to assign healthworkers across the entire intervention area while keeping their travel distances manageable. Mali Health therefore paired healthworkers of similar experience and quality, and one of each pair was assigned to the healthworker only or the healthworker and free care group. The pairs were then assigned in an overlapping pattern by stratum. For example, pair 1 and pair 2 were each randomly assigned half of the families in the same stratum, and half of the families in a different stratum, each shared with another pair of healthworkers. In this manner, each stratum was assigned four healthworkers (two in each treatment group), the quality of the healthworkers in the treatment groups

Table 1: Sample and Balance.

		Baseline (2012)			Baseline variables in households included at follow-up (2013)				ANOVA for group differences
		Mean	SD	N	Control Mean	Healthworker Mean	Free care Mean	HW & FC Mean	
<i>Household</i>	Number of members	6.24	3.14	961	6.15	6.13	6.32	6.44	0.73
	Assets (USD)	6169	12880	940	5233	6299	7512	5722	0.29
	Weekly income (USD)	63.28	76.45	902	58.40	62.67	66.19	66.62	0.65
	Numer of rooms	2.06	1.56	936	1.98	2.11	2.14	2.04	0.75
	Owner-occupier	0.45	0.50	940	0.42	0.47	0.48	0.44	0.62
<i>Household head</i>	Male	0.89	0.32	968	0.90	0.88	0.88	0.89	0.82
	Literate	0.48	0.50	969	0.46	0.44	0.47	0.55	0.11
	Bambara 1st language	0.69	0.46	969	0.73	0.69	0.68	0.64	0.18
	Salaried	0.12	0.32	969	0.12	0.13	0.13	0.10	0.59
<i>Child</i>	Age (Months)	40.42	0.47	1544	40.61	40.47	39.51	40.26	0.66
	Male	0.51	0.01	1543	0.50	0.48	0.52	0.52	0.67
	Adopted	0.06	0.01	1487	0.06	0.06	0.06	0.03	0.30
	Weight for Height Z-Score	-0.61	0.05	1307	-0.65	-0.62	-0.55	-0.63	0.83

3.2 The Sample Population

Mali Health originally identified 1764 eligible children in the census. At the baseline survey about three month later, a total of 81 children were not found by the surveyors. The roll-out of the Action for Health program started another three months later in early 2013. By the second survey round in Fall 2013, an additional 139 children were not found (including 5 children who passed away, see table 12 in Appendix C.1 for a breakdown). In total, the sample in 2013 consists of 1544 children present in both survey rounds, from a total of 957 households and 592 compounds. This corresponds to an attrition rate of 8.3% between baseline survey (and assignment to treatment group) and follow-up. Attrition rates by treatment group vary between 6.9% (HWFC) and 10.0% (FC); these differences are not significant (p-value .20) using an ANOVA test.¹⁵

Table 1 provides an overview of the demographics and finances of the households in our sample. There are on average more than six people in a household, due in part to polygamous marriages and multi-generational households. Most households are headed by a man. About half of the household heads are literate and two thirds speak the lingua franca of the region, Bambara. Average household assets are USD 6169, or almost CFA 3 million, with an owner occupancy rate of 45%. Average weekly income is about \$63, though only 12% of household heads earn a regular salary.

The average age of the children in our sample was about 3.4 years at the time of the 2013 survey. Our study children are more than half a standard deviation below their age-typical weight for height, and about 6% of children are fostered or adopted, a common

was matched, and no two strata had identical sets of healthworkers.

¹⁵In addition, we have 2013 data on children who were not present in the 2012 baseline, in particular infants born after the 2012 census. This data is not used in our analysis.

practice in West African households. An ANOVA analysis shows no significant differences between the treatment groups for any of the 13 variables listed in table 1.

3.3 Health Calendar

The core element of the data collection is a detailed health calendar for each child, containing daily information on any symptoms the child exhibited, all consultations that occurred during the week with respect to the child’s health, and all medications taken. Health calendar data was collected on a weekly basis over the course of nine weeks from the child’s primary caretaker (usually the mother). To aid mothers’ memory, they were given pictorial diaries with images representing the different symptoms (Appendix B, figure B.1). They were asked to mark off any symptoms and health-related events on the day they occurred, and the surveyors reconstructed the child’s health history together with the mother during the visit following each week. This method provides us with a uniquely detailed record of all the health events involving the study children, without the problems associated with long recall periods (Das et al. [2011]).

3.3.1 Symptom Calendar

The list of symptoms was designed in collaboration with Mali Health staff and based on the C-IMCI. The C-IMCI is designed for use by community healthworkers (CHWs) who have no prior medical experience, and is used in Mali Health’s own healthworker training. It consists of simple rule charts that use important and easy-to-spot symptoms to broadly classify a child’s illness and establish need for care. These properties make it well-suited for use in this study: first, they allow the use of symptom reports from mothers and surveyors who are not medically trained, and second, the treatment rules allow us to map observed symptoms into when a child should seek care according to the policymaker.

The survey explicitly names ten symptoms: convulsions fits, or spasms; lethargy or unconsciousness; inability to drink; vomiting; coughing; difficulty breathing; more than three loose stools; blood in the stool; sunken eyes; and unusually hot (cold) skin. Other symptoms are recorded free-form. From those records we constructed four additional symptom groups: skin conditions, cold symptoms, ear pain, and injuries.

We observe on average 60 days per child, and of those the child exhibited symptoms on 18 days. The most frequent symptoms were cold symptoms, coughing, and unusually hot skin. Section 3.4 below gives more detail on the symptoms and how they translate into illness spells. Appendix B provides a description of how the health calendar information

Table 2: Consultation numbers and average costs and wait times by type of provider (left), and effect of the subsidy on individual cost and value of care received at the CSCOM (right).

Healthcare visits	Formal		Informal	N	Cost per CSCOM visit		<i>Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Private costs for consultation and treatment as reported by the parents. Total costs imputed from private cost and median prices, by consultation type and treatment received.</i>
	CSCOM	Other			Private	Total	
Number of visits	514	67	2253	Control	50	2577 (381)	3766 (322)
Average total cost in CFA (Std. dev.)	1347 (2759)	5578 (5479)	222 (557)	Healthworker visits (HW)	54	3324 (446)	3884 (447)
Average wait time in hours (Std. dev.)	0.73 (0.95)	1.28 (2.09)	0.04 (0.49)	Free Care (FC)	203	948*** (191)	4042 (227)
				HW and FC	207	837*** (179)	3857 (218)
				All	514	1312 (136)	3926 (139)

CSCOM: CSCOM or associated reference hospital (CSREF). Other: private doctor, hospital, or health center provided e.g. by the Red Cross. Informal: peddlers, market, traditional or religious healers, midwives, and pharmacy visits without a formal consultation. Costs and wait times of all connected consultations, e.g. drug purchases.

was collected and gives an overview of the raw numbers of recorded symptom days in table 11.¹⁶

3.3.2 Consultations

At each surveyor visit, mothers are asked if they discussed their child’s health with anyone in the previous week, and the surveyors recorded the role or occupation of the person seen, the type of facility, the waiting time, and the treatments received with all costs incurred.¹⁷

Table 2 (left panel) provides an overview of these healthcare “consultations” by type, their average cost to the household as reported by the parents, in CFA (476 CFA equaled 1 USD in 2013), and the associated wait time in hours. The 581 formal consultations in our data are provided by CSCOMs and CSREFs (associated reference hospitals), as well as by private doctors, hospitals, and clinics. Note that parents often visit several care providers in the same day, for example when they get a prescription from a (formal) doctor and then purchase medication at a pharmacy or store. We group visit records that are linked in this manner and classify them by the “most formal” consultation in the

¹⁶The table shows that the full treatment group experiences or reports on average slightly higher rates of illness compared to the control, and it is possible that the increase in symptom reporting is partly an effect of the treatment. However, the difference is driven by injuries and cold symptoms. Cold symptoms only indicate a need for care if combined with fever and rash (measles suspicion) or after 14 days (bacterial respiratory infection). Injuries are not listed in the C-IMCI. Mali Health does not pay for healthcare related to a simple cold or to injuries.

¹⁷In order to facilitate the recording of drugs prescribed, purchased, and taken, mothers were asked to keep the packaging of any drugs their child received. The surveyors could search and fill the drug from a database of brand and generic names of about 300 medications commonly sold in Mali. For commonly used sources of health care in the area we also have GPS location data, otherwise we collected approximate distances from the respondent’s house.

group. Table 2 shows that the average cost of a CSCOM visit as reported by parents is low compared to other formal sources of care, in part due to the free-care intervention. Private formal providers have the longest wait times, but the wait for a CSCOM visit is still almost 45 minutes on average. By contrast, informal visits and purchases, which comprise the vast majority of consultations, come at low monetary, and almost no time cost.

For our analysis of utilization and demand below, we focus on public and private formal healthcare as the (only) way to address a medical need for care. This is in line with the C-IMCI recommendations, and so the preferences of the policymaker. It assumes that adequate care requires a formal consultation with a medically trained doctor. Typical informal sources are peddlers and market sellers, stores, healers, Islamic marabouts, or “elders”, individuals who, to our knowledge, have no formal medical training that enables them to diagnose illness, and are not bound by the hippocratic oath. We include 203 pharmacy visits that are not associated with a doctor visit in the “informal” category because pharmacists do not carry out a diagnosis, so these purchases amount to self-medication. While informal consultations are not *necessarily* inadequate, there is by definition no way to control quality.

Table 2 (right panel) shows the effect of the free-care component of AfH on the costs of a CSCOM visit to parents and the total (social) costs. Social costs are given by the undiscounted price of care (which consists of all private costs born by the parents or other contributors, e.g. relatives, plus the costs incurred by Mali Health) and provide a proxy for the value of healthcare consumed.¹⁸ The table shows that the subsidy reduces the price of a CSCOM visit by between CFA 1629 (FC) and 1740 (HWFC), compared to a baseline average of CFA 2577. This constitutes an effective price decrease of 63-68%. At the same time, the value of care received per visit is unchanged. This confirms that the supply of healthcare, conditional on visiting a clinic, is not affected by our intervention, allowing us to focus on demand changes.¹⁹ The healthworkers have no significant effect

¹⁸The source for all cost measures are the payments made for consultations and medications. Social costs are calculated using prices recorded for those purchases where the respondent reported having paid the full price themselves. We use these prices combined with provider information and medication brand names and point of purchase to construct the median price of each consultation and medication by source. If the respondent reported another person paying for care or received CSCOM care in the free care group, we compare the sum of median prices for the services received with the private expenses the respondent reported and use the higher value. This is intended to reflect the total cost of care to the family, the NGO, or a third party.

¹⁹In general, one might expect a moral hazard problem on the provider side, namely that doctors respond to the subsidy by raising treatment costs. This is not the case here, due to Mali Health’s quality and cost control, and in Appendix C, Table 13 we show that this holds also when controlling for day of

on the cost of care at the CSCOM, as would be expected. Note that the sizable difference in private and social costs even in the control is partly due to a few large outliers. In such emergencies, families seek outside help to pay for the child’s care.

3.4 Illness Spells and Optimal Care

We use the symptom data from the health calendar to construct illness spells. An illness spell is any contiguous sequence of days on which a child exhibits one or more symptoms. The object of interest for much of this paper is the time until care is sought, or equivalently the “pre-care” spell, which lasts until the first formal health consultation occurs or until recovery, whichever comes first. For example, a child who is coughing for three days and takes cough drops bought from a peddler, then develops a fever on day three and is prescribed and antibiotic and paracetamol by a CSCOM doctor on day five, so that the fever subsides right away and the cough disappears after day eight, would have a eight-day illness spell and a five day pre-care spell.

Our data contains 3160 useable pre-care spells (i.e. spells without left-censoring – see section 4.2) with an average length of 6.5 days. Incidence is remarkably symmetric between the four treatment groups; almost exactly a quarter of all spells occurs in each of the treatment groups (between 781 and 798 spells, see table 7 below).²⁰

We next determine when care should be optimally sought in each spell according to the C-IMCI. The C-IMCI is part of the community and family arm of the IMCI itself, the WHO’s and UNICEF’s primary children’s health program, which was adopted by over 80 countries (Bryce et al. [2004]). These guidelines are explicitly tailored to a developing country context and therefore take into account budget constraints and expected quality of care.

The C-IMCI goes through a set of simple checks and questions in order to spot acute signs of danger and to classify symptoms into those related to gastrointestinal illness, respiratory illness, fever and malaria, and so on. It then makes fairly intuitive recommendations for care-seeking based on the length of time a set of symptoms is observed, as our

the spell and IMCI “early” classification (to rule out potential selection effects – treatment at different points in the spell may differ in price). The estimates in the appendix do reveal that treatment prescribed on an early day costs about CFA 1800 less compared to when care is required. This lends some credibility to the WHO guidelines as an indicator of need for care.

²⁰This is consistent with the fact that the main sources of illness – malaria, respiratory, and gastrointestinal – are infectious, and that better access to primary care is not expected to reduce infection rates. Even if the children in Action for Health receive treatment earlier and more often, they constitute only a small percentage of the population. In other work we document that the healthworkers did not increase prevention significantly.

Table 3: Days of (untreated) illness per child, by need for care, with proportion on which each symptom is observed.

Pre-care days per child:	All		Early		Care required	
	Mean	SD	Mean	SD	Mean	SD
Total number	15.93	(14.52)	7.55	(8.27)	8.38	(11.54)
Percentage of total number on which each symptom is present:						
Convulsions, fits, or spasms	0.36%		0%		0.53%	
Lethargic or unconscious	4.06%		0%		8.02%	
Unable to drink or breastfeed	1.08%		0%		2.06%	
Vomiting everything	5.08%		0%		10.93%	
Coughing	32.47%		28.68%		36.48%	
Difficulty breathing	4.54%		3.93%		5.34%	
> 3 loose stools	7.47%		6.66%		10.17%	
Blood in the stool	0.60%		0.05%		1.10%	
Sunken eyes	2.05%		0.16%		3.60%	
Unusually hot skin	31.62%		0.15%		62.49%	
Other: rash, spots, or itch	2.97%		3.49%		1.67%	
Other: cold symptoms	51.11%		59.24%		42.71%	
Other: ear ache	1.00%		0.85%		1.57%	
Other: wound, injury, or burn	4.08%		5.98%		1.74%	
Other symptoms	5.56%		7.22%		2.87%	

Symptoms that indicate need for care always occur on "care required" days, e.g. danger signs. Symptoms that do not themselves indicate need for care according to C-IMCI can co-occur on care-required days, if care is required based on other symptoms, e.g. injuries.

model suggest. For example, a fever without any other symptoms could indicate malaria and requires urgent care; a child with simple diarrhea should be seen by a doctor only if the symptoms have been present for five days, but if the diarrhea is accompanied by blood in the stool, care should be sought immediately. Applying these rules to mothers' symptom reports, every day in the pre-care spell can be classified as either an 'early' day (i.e. before care should be sought) or a 'care required' day (on or after the day at which care should be sought). In the example spell described above, the first three days would count as 'early days', while all subsequent days would be 'care required' days (24 hours after the fever started). Appendix B describes the classifications we applied in detail.

Mali Health bases their training on the C-IMCI, and we assume that these guidelines represent the policymaker preferences (but note that different preference specifications could be easily adopted). We take these preferences as given and focus on demand behavior relative to them, under the assumption that the C-IMCI was formulated according to some optimizing process to balance expected health benefits with resource constraints. Our goal is not to assess if adherence to the WHO recommendations for care-seeking in fact *do* have positive long-run health effects. On a larger scale, demand estimates like ours could be a first step in studying this question.

The top of table 3 shows the raw number of untreated days per child (which may

represent several different spells), both in total and split into early and care-required days. Below is the percentage of those days on which each symptom is present. For example, the first four symptoms, convulsions, lethargy, inability to drink, and vomiting, are acute danger signs and should trigger immediate care (so they always occur on care-required days). On average, nearly 16 (27%) of 60 observed days per child were untreated (pre-care) illness days, and on more than half of those there was an (unfilled) need for care. The biggest contributors to need for care are gastrointestinal symptoms (diarrhea and vomiting) and fever (unusually hot skin). While cold symptoms and coughs in most cases do not in themselves indicate a need for care, they co-occur on many “care required” days.

Note that the symptoms we selected for collection do not cover all possible illness, but rather focus on conditions that both mothers and surveyors could easily recognize and report. As an example, some symptoms mentioned in the C-IMCI as indicators of respiratory illness, specifically the number of breaths per minute and the presence of “chest indrawing”, require experience and training (and a stopwatch) to be collected reliably. Note also that wounds and injuries are not mentioned in the C-IMCI and were therefore not used to determine need for care.

This means that our classification as “care required” contains some noise and generally represents a *lower bound for the true need for care*. On some days that we designate “early” the child may actually require care according to the (full) C-IMCI guidelines. Furthermore, parents may have private information that indicates a need for care according to the “true” policymaker preferences, even if these preferences are not specified in the C-IMCI (e.g. the child may have a life-threatening infected wound, but the symptoms of blood poisoning are too rare to make it practicable to train community health workers to spot them).

Table 4 shows some summary statistics on the recorded pre-care spells. The model in section 2 highlighted the role of spontaneous recovery for the optimal timing of care and the determination of over- and underuse. The possibility of spontaneous recovery introduces a form of right-censoring, because we do not observe when parents would have sought care if the child had *not* recovered. This type of censoring is an integral part of the data generating process, and as we argue below, it is important for understanding healthcare demand, especially when a large number of spells do not end in formal care. Indeed, only 494 of the 3564 spells in our data set are uncensored, and most of the censoring is due to the child’s recovery (not missing data). The average uncensored or right-censored spell is nearly one week long, with a high standard deviation. Within these

Table 4: Pre-care spell lengths and numbers by type of censoring.

Pre-care spell length in days	Uncensored and right-censored			Left-censored
	Total	Early	Care required	Total
Mean days within spell	6.543	3.216	3.327	10.292
SD	(7.142)	(4.053)	(6.311)	(10.582)
N	3160	3160	3160	404

A pre-care spell is uncensored when a formal consultation occurs. It is right-censored if no consultation occurs on the last day of the (observed) illness spell. It is left-censored if the last day before the first observed spell day is missing from the data.

spells, there are on average 3.2 “early” days and 3.3 “care-required” days. Note that, due to survey interruptions and spells recorded in the first week of the survey, a number of illness spells are left-censored, that is, they may have started before the first recorded day. These spells are on average longer. Left-censoring means that need for care cannot be determined, because the day of the spell is unknown. We discuss the implications for estimating healthcare demand below.

4 Results

We use the above data to perform three sets of analyses. First, we study the impact of subsidies and healthworkers on *unconditional* utilization of acute care. Due to the fact that all CSCOM care in the area is provided by Mali Health’s cooperation partners, we can observe these effects without the interference of offsetting supply responses, like lower quality or higher prices. Second, we study the impact of our three treatments on healthcare demand *conditional* on healthcare status. We apply WHO standards of care to judge the proportion of demand generated (or avoided) which is wasteful, and how the program improves access to care for children who need it; in other words, to learn about over- and underuse. Third, we want to understand the *underlying choices* which generate these outcomes. Our theoretical model shows that care-seeking arises from the dynamic decision process of the parents, combined with the stochastic processes that drive illness incidence and spontaneous recovery. In order to separate the former from the latter, we estimate a hazard model of care seeking during an illness spell, which characterizes the underlying individual demand for formal care. This allows us to evaluate our model predictions, and it provides us with symptom-day specific estimates for demand behavior that can be used to make out-of-sample predictions for different disease environments, seasons, or populations. We demonstrate this by predicting care-seeking behavior with and without subsidies for typical spells of hemorrhagic fever (Ebola and Marburg virus),

Table 5: Utilization, private expenditure, and social cost of treatment per child, by treatment group.

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
		Number of visits per child			Private spending per child (CFA)				Wait time	Social cost per child	
		Total formal	CSCOM	Other formal	Total	CSCOM	Other formal	Informal	per child (hrs)	Total	CSCOM
Control	<i>Mean</i>	0.178	0.125	0.053	1274	323	476	362	0.2	1480	472
N = 399	<i>S.E.</i>	(0.022)	(0.018)	(0.015)	(170)	(62)	(135)	(50)	(0.05)	(185)	(75)
Healthworkers (HW)	<i>Mean</i>	0.203	0.142	0.061	1740*	493	690	363	0.28	1850	576
N = 379	<i>S.E.</i>	(0.031)	(0.025)	(0.017)	(211)	(118)	(120)	(46)	(0.06)	(223)	(130)
Free Care (FC)	<i>Mean</i>	0.571***	0.532***	0.039	1279	532	350	285	0.54***	2991.2***	2190***
N = 389	<i>S.E.</i>	(0.053)	(0.051)	(0.013)	(188)	(121)	(104)	(49)	(0.08)	(324)	(265)
HW and FC	<i>Mean</i>	0.56***	0.538***	0.021*	1101	451	265	287	0.48***	2812***	2077***
N = 377	<i>S.E.</i>	(0.051)	(0.049)	(0.01)	(153)	(103)	(82)	(57)	(0.07)	(255)	(215)
All	<i>Mean</i>	0.376	0.333	0.043	1347	449	445	325	0.37	2277	1322
N = 1544	<i>S.E.</i>	(0.023)	(0.022)	(0.007)	(91)	(52)	(57)	(25)	(0.04)	(132)	(103)

*Standard errors clustered at the compound level; *** p<0.01, ** p<0.05, * p<0.1 (differences from the control). Private costs for consultation and treatment as reported by the parents. Social costs imputed from private cost and median prices, by consultation type and treatment received. Note that (4) is larger than the sum of (5) to (7) because some consultations could not be classified (e.g. due to contradictions between location and person seen).*

an important application in the context of West African public health.

4.1 Utilization

Overall Healthcare Utilization We begin by describing the impact of subsidies and healthworker visits on unconditional measures of healthcare utilization. Columns (1) to (3) in table 5 show the number of formal consultations per child, in total and split into CSCOM and other formal care. (4) to (7) report the effect of the program on private expenditure, first in total and then for CSCOM, other formal, and informal care separately. Column (8) reports wait times per child, and (9) and (10) report total and CSCOM social costs (value of care, see above) per child. All standard errors are clustered at the compound level.

The demand elasticity for care is clearly high: with free care, the numbers of CSCOM visits and of total formal visits per child respectively increased by 430% and 314% (with healthworkers) and by 426% and 320% (without), relative to the decrease in the CSCOM visit cost of 68% (HWFC) to 63% (FC) (see table 2 above). There is some substitution out of other formal care, significant at the 10% level in the HWFC group.

Remarkably, column (4) shows that the subsidy has almost no effect on total private health expenditure. If anything, columns (5) to (7) suggest that it leads to somewhat *higher* expenditure per child at the CSCOM, only partially offset by lower spending for other care. This implies that there is no crowding out of private spending; all subsidies

Table 6: Acute formal consultations by treatment group, if occurring on “early” or “care-required” day.

Acute formal visits by "care required" status	Control		HW			FC			HWFC		
	Avg. visits per (S.E.)	Percent of visits	Avg. visits per child (S.E.)	Percent of visits	Demand in % of control	Avg. visits per child (S.E.)	Percent of visits	Demand in % of control	Avg. visits per child (S.E.)	Percent of visits	Demand in % of control
Acute visits	0.163 (0.020)	100%	0.169 (0.022)	100%	106%	0.488 (0.034)	100%	302%	0.443 (0.034)	100%	288%
early	0.030 (0.009)	18%	0.029 (0.009)	17%	97%	0.067 (0.013)	14%	222%	0.077 (0.015)	17%	269%
care required	0.133 (0.018)	82%	0.140 (0.019)	83%	108%	0.422 (0.032)	86%	320%	0.366 (0.031)	83%	292%

Acute care refers to the first consultation during a spell. All standard errors are clustered at the compound level.

translate into additional healthcare consumption, and as a result the value of care received by each child almost doubles under free care (columns (9) and (10)). Furthermore, the free-care groups pay higher non-pecuniary costs for their child’s healthcare, as evidenced by the increase in time spent waiting for treatment (column (8)).

Overuse and Underuse According to WHO Criteria Next, we use WHO standards of care to classify the observed increases in utilization. Table 6 shows all acute consultations (i.e. those that took place during an illness spell), split by whether the spell had entered “care-required” status at that time.²¹ The first column shows very little evidence for overuse in the control group: 82% of all formal consultations take place when the child is in the “care required” status. Moreover, even though there is a large increase in acute visits under free care, almost all of these are appropriate according to the WHO classification. Relative to the control, the free care group has 0.29 more visits per child during “care required” periods, relative to 0.04 during “early” periods (0.23 vs. 0.05 for the HWFC group). This means that subsidies leave the proportion of visits that reflect overuse relatively unchanged: on average 15% in the group with free care and 18% in the group without. The proportion of overuse visits is in fact lowest in the free care-only

²¹Table 14 in the appendix shows a complete breakdown of all consultations inside and outside a spell. Only 5% of visits occur outside a spell, reported by mothers as follow-ups, visits for prevention (e.g. vaccinations), or “other”. Of the remaining consultations, 88% (or 494) constitute the first formal visit during a spell of illness, or what we call an “acute care” visit. These acute visits are the target of the WHO care-seeking recommendations. Visits after a first acute visit cannot directly be judged using the C-IMCI standards. However, almost all of them occur after a spell entered “care required” status and when the child is still exhibiting symptoms. This might indicate a different treatment approach by doctors when care is free: anecdotally, CSCOM doctors lament that it is very hard to convince patients to come back after a first visit, so that they often prescribe treatments for a range of conditions at once. When the cost of care is reduced, they may instead be able to ask the parents to return to try a different treatment approach when the child’s condition has not improved.

Table 7: Total spells vs. spells with a consultation by group; conditional on spell entering “care required” status or remaining in “early” status.

	Control		Healthworker		Free care		HW & FC		Total	
Total number of spells	781		790		798		791		3160	
Spell did/did not enter "care required" (care req./early)	early	care required	early	care required	early	care required	early	care required	early	care required
(A) All spells	374	407	327	463	368	430	353	438	1422	1738
(B) Spells with a consultation	11	44	7	45	23	134	28	120	69	343
Percent spells w. cons. (B/A)	3%	11%	2%	10%	6%**	31%***	8%***	27%***	5%	20%

Significance levels: *** 1%, ** 5%, * 10%, t-test on mean difference from control.

group. Healthworker visits have essentially no effect on overuse.

Since our data is collected at the household level, we can observe utilization per *spell*, which allows us to study rates of underuse, as well as overuse. As a starting point, table 7 reports the pre-care spells in each treatment group, the number of those spells that do (or do not) enter “care required” status at any point according to the C-IMCI, and finally the number and proportion of those that receive acute formal care. Note that a spell can only transition from “early” to “care required”, not the other way around. Moreover, since we are considering pre-care spells (not illness spells), the consultation always occurs on the last day of the spell.

The table shows, first, that overuse occurs only in a very small proportion of spells; on average 2.5% of spells that do not enter “care required” end in a consultation without free care, rising to 7.1% with free care. Second, underuse is rampant: in the control, an illness spell that requires care ends without a formal consultation 89% of the time. The subsidy treatment reduces this proportion to 73% (with healthworkers) and 69% (without). The effect of information provided by the healthworker is small and does not indicate an improved allocation of care. In fact, the point estimates of proportion of “care required” spells that end in a formal consultation is *lower* with healthworker visits (although the effect is not significant). We will return to this point below.

An implication of tables 6 and 7 is that expanding free care to the other two treatment groups would have led to 191 additional doctor visits, of which at least 82% (157) would have occurred when care was required. Table 7 reveals that there is very little room to reduce overuse further – that 16% of consultations occur on “early” spell days is partly a result of substantial *underuse*. If the proportion of care-required spells that receive care had gone up from 29.3% to 92.9%, for equal rates of underuse and overuse, the proportion of consultations that constitute overuse under free care would have been only 6%. In order to achieve the same by reducing overuse further, the proportion of early spells with consultations would have had to be at 2.4% lower than in the control.

4.2 Probability of Care-Seeking by Spell Day

The results above give us a first idea of the effects of subsidies and information on observed care-seeking *outcomes*. These outcomes are a combination of individual behavior and exogenous stochastic processes: the probability of falling ill determines if acute care will be needed at all; spell length and utilization per spell – if parents seek care before the spell ends – depend on spontaneous recovery rates. In order to fully understand the effect of the policy interventions on *behavior*, we want to separate it from exogenous illness incidence. Based on our model, parents make the choice to visit a doctor anew on every illness day. The object of interest for demand estimation is therefore the probability that the parents seek care conditional on continued illness; or in other words, the “hazard” of care. This approach also resolves the censoring issue that arises from the possibility of spontaneous recovery (see below).

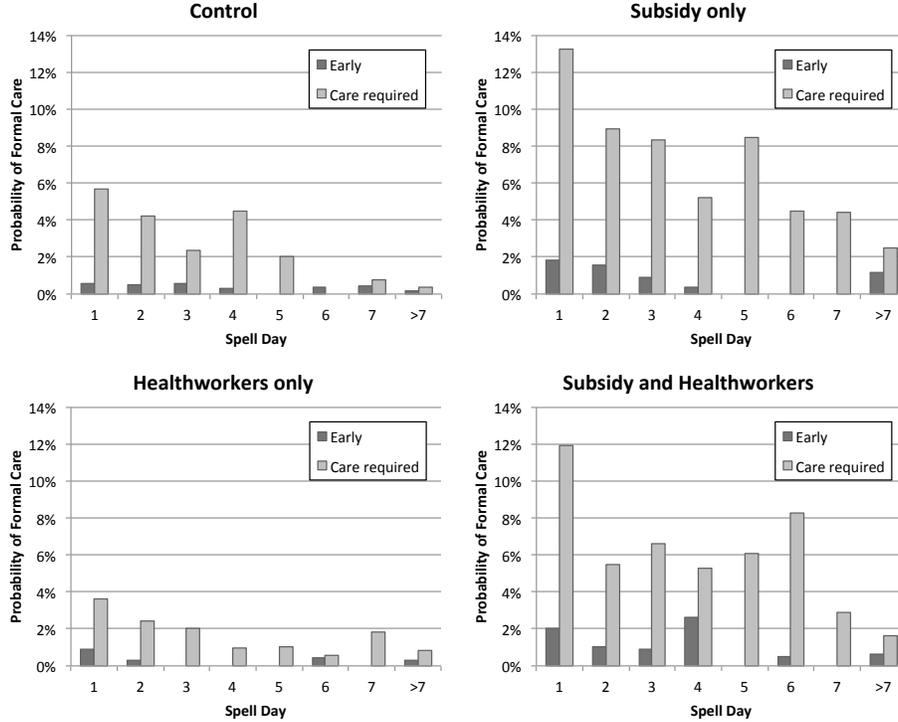
Non-Parametric Hazard of Care We analyze the probability of receiving formal care on each illness spell day t conditional on whether it is an early or care-required day according to the C-IMCI. This is just one possible way of controlling for the characteristics of the spell; other specifications are equally possible, e.g. estimating care-seeking probabilities for different symptom combinations (see section 4.3).

Figure 4.1 shows the simple average probability of seeking care on each spell day by treatment group and “care required” classification. Again, we see little evidence of overuse, coupled with significant underuse. The hazard of care is generally higher early in the spell. The probability of care-seeking on “early” days is below 1% in the control and the healthworker-only group, and almost entirely below 2% in the subsidy groups. Free care more than doubles care seeking; in the first five spell days, the probability is above 8% on average in the free care-only group. The hazard of care appears *lower* in the groups with healthworker visits than those without.

Cox Proportional Hazard Model We use a Cox proportional hazard model in order to concisely estimate the average effects of subsidies and information. The Cox model imposes no restrictions on the daily baseline hazard, but assumes that the effect of each covariate shifts the probability of seeking care proportionally on every day t of the illness spell. Formally the hazard at t conditional on covariates x_{it} is

$$h(t|x_{it}) = \lambda_0(t)e^{x'_{it}\beta}.$$

Figure 4.1: Each panel shows the average probability of care seeking per spell day within the control and the treatment groups, separately for spell days classified as early and as care-required according to the C-IMCI. Spell days past day 7 are grouped.



As we show in Appendix C, figure C.1, the proportionality assumption delivers a good description of our data.

Besides child and household controls, the explanatory variables are an “early” dummy that captures C-IMCI need for care, and indicators for the free care and healthworker interventions and their interaction (treatment dummies). The coefficients on these variables tell us to what degree individual behavior agrees with the policymaker’s preferences and how it responds to the policy interventions. We interact treatment and “early” dummies to allow for the demand response to differ depending on the policymaker’s preferences. In particular, if better information reduces both overuse and underuse, the effect of the healthworker treatment on the hazard of care should be *positive* on days when care is required, but *negative* on early days.

Table 8 reports the estimated hazard ratios from this model, that is, the proportional increase in the probability of seeking care when the independent variable increases by one unit. Models (3) and (4) control for the household’s assets at baseline and the distance to the closest formal care provider; the gender, age, and literacy level of the household

Table 8: Estimates from a Cox proportional hazard model (uncensored and right-censored spells).

COX HAZARD MODEL	(1)	(2)	(3)	(4)
Early: care not yet required (IMCI)	0.147*** (0.0200)	0.150*** (0.0524)	0.140*** (0.0207)	0.159*** (0.0565)
Healthworker visits (HW)	0.769 (0.182)	0.774 (0.196)	0.628* (0.157)	0.636* (0.167)
Early x HW (EHW)		0.960 (0.491)		0.958 (0.537)
Free care (FC)	2.792*** (0.498)	2.892*** (0.569)	2.522*** (0.464)	2.644*** (0.542)
Early x FC (EFC)		0.805 (0.333)		0.758 (0.323)
HW and FC (HWFC)	1.147 (0.320)	1.068 (0.320)	1.525 (0.452)	1.453 (0.457)
Early x HWFC (EHWFC)		1.528 (0.920)		1.303 (0.859)
Total assets in US\$ (log)			1.048* (0.0229)	1.048* (0.0229)
Dist. to closest formal (log)			0.851 (0.0934)	0.852 (0.0935)
Child controls	-	-	YES	YES
Household head controls	-	-	YES	YES
Observations	20,599	20,599	18,080	18,080
Partial LL	-2880	-2879	-2479	-2479
Wald chi-square	259	267.4	257.5	270.5

*Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

Standard errors clustered at the compound level (in parentheses).

head; and indicators for the child’s age, relation to the household head, and gender.²² Specifications (2) and (4) include interactions of the treatment dummies with the “early” indicator. Note that the model is estimated using uncensored and right-censored spells only; in left-censored spells, the day of the spell is unknown, so we cannot estimate the baseline hazard or determine “care required” status. However, we re-estimate the model under inclusion of these spells (treating the first observed day as the first spell day) and find only small changes in the estimated coefficient sizes (see Appendix C, table 15).

Table 8 shows that the probability of care-seeking increases significantly and by at least 250% under free care, both on early and on care required days. This is consistent

²²Of these controls, only assets and child age affect healthcare demand positively and significantly, while the distance to the closest provider as a measure of non-monetary costs has the expected (negative) effect, but is not significant.

with figure 4.1: whereas the *proportional* impact of free care is roughly the same on early and care-required days, its *absolute* impact is much larger on care-required days. The reason is that the hazard of care on early days is on average 14-16% of that on care-required days (in specifications (2) and (4), the hazard ratio of “free care x early” is in addition below one, although not significantly so).

Given the extremely low probability of early care seeking, it is perhaps not surprising that the healthworkers have no significant effect on overuse. Table 8 shows in fact that better information may *reduce* demand for formal care: in all specifications, the coefficient on the healthworker dummy is negative, and the effect is significant at the 10% level in (3) and (4). In absolute terms, care-required days are again more affected. This can also be seen in figure 4.1.

While the point estimates on the interaction effect of healthworkers and free care from specifications (3) and (4) in table 8 are not significant, their size implies an increase in utilization that reverses the negative effect of healthworkers alone. The combined hazard ratio on the healthworker effect in the free care group is 0.96 in (3) and 0.92 vs. 1.15 on care-required vs. early days, respectively, in (4) (none significantly different from 1). This suggests a form of complementarity – while healthworker visits on their own reduce care-seeking, this affect disappears when they are coupled with free care.

4.3 Discussion and Application

Our results have implications for understanding the barriers to care faced by the study households, and for the estimation of healthcare demand. We also apply the results of the hazard model estimation to provide counterfactual estimates of the impact of free care on care seeking for a child with Ebola.

Implications: Barriers to Care Our results are consistent with the view that the primary barrier to effective care seeking in our study population is the (monetary) cost of care, not information. The low level of utilization suggests that parents apply a significantly higher cost threshold than the policymaker, and even when care is subsidized, alignment is not complete, so that a significant amount of underuse remains. Moreover, parents recognize differences in illness severity, as care seeking is much less likely on ‘early’ days compared to ‘care required’ days.

On the other hand, our model showed that better information may *increase* underuse, provided the parents apply a much higher cost threshold than the policymaker. Indeed,

this is the pattern we find in our data: the healthworkers *increase* underuse when there is no subsidy. This can occur if the parents use the information from the healthworkers to learn when an illness is *not* extremely serious, leading them to use less healthcare.²³ This result both confirms the important role of cost barriers, and suggests that information is not perfect, either. However, since the parents in our sample are rarely seeking care too early, there is very little overuse for the healthworker to correct, and therefore information is not the binding constraint in our population.

Implications: Censoring and Selection Bias Without spontaneous recovery, all pre-care spell observations would be uncensored, and estimating hazard of care, time-to-care, and the length of the pre-care spell would all be equivalent: the average probability of an event occurring (per period) would translate directly into the average time until this event occurs. However, our model points out that spontaneous recovery is an integral part of the decision to delay care in the first place, and with this censoring – the child recovers before care is sought – the observed pre-care spell length systematically underestimates time-to-care.

This is important because the *degree* of censoring is affected by any treatment that changes time to care. For example, if the healthcare subsidy lowers K , some spells are censored at high K , but uncensored at low K (this is what leads to increases in utilization). As a result, the average pre-care spell length underestimates time-to-care at high K by more, and will therefore show a smaller reduction than true time to care. This is exacerbated in data on uncensored spells only, for example consultation data collected at the healthcare provider.

This point is clear from table 9, which shows the average pre-care spell length by treatment group. The top of the table shows that free care significantly reduces the time to care on average, but only by at most about 1.2 days, or between 9% and 17%. The bottom of the table shows that, while the total number of uncensored spells (spells that end in a consultation) is much higher in the groups that receive some or all of Action for Health, the average uncensored spell length is weakly longer. Our hazard estimates show that the apparent increase in time-to-care is entirely a result of selection bias, and not,

²³This interpretation is also consistent with a scenario in which the parents do not learn from the healthworker how to interpret symptoms themselves, but use the signal from the healthworker’s assessment of the child to infer illness severity. The healthworker visit is essentially used as a substitute for a formal consultation. Note that part of the purpose of a formal consultation *is* the evaluation of the child, including possibly a conclusion that treatment is not needed (yet). However, the WHO wants this evaluation to be carried out by a trained doctor, even *after* a community healthworker has assessed the child’s health status.

Table 9: Length of an untreated (pre-care) spell by treatment group, all spells and uncensored spells only.

Pre-care spell length in days	Control	Health worker	Free care	HW & FC	Total
Uncensored and right-censored spells					
Mean length	7.014	6.959	5.810***	6.402*	6.543
SD	(7.341)	(7.924)	(6.027)	(7.103)	(7.142)
N	781	790	798	791	3160
Uncensored spells only (ending with a formal consultation)					
Mean length	3.764	6.346*	4.401	4.959	4.762
SD	(4.303)	(8.733)	(4.861)	(7.100)	(6.278)
N	55	52	157	148	412

Significance levels: *** 1%, ** 5%, * 10%, t-test on mean difference from control.

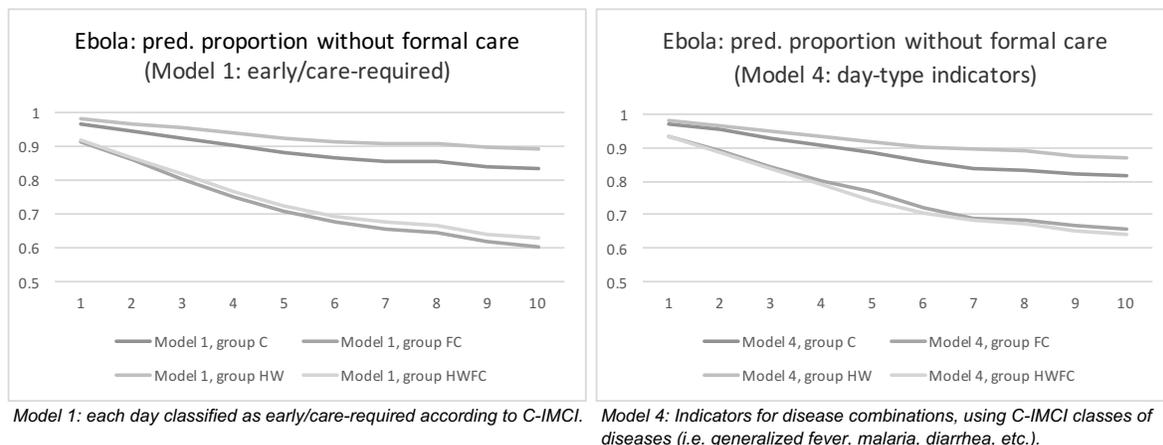
for example, due to a form of moral hazard, where families go to the doctor *later* during a spell when the costs of a visit fall, and only get last minute “emergency” care.

We can assess the magnitude of the bias by reestimating the hazard model on the uncensored spell data only. The results, shown in table 16 in Appendix C, indicate that we would have wrongly concluded that free care has no or even a negative effect on time-to-care, and that healthworker visits alone lead to substantial delay. Using only uncensored data also almost doubles the estimated hazard ratio on early days, so that parents appear less able to discern the health care needs of their children. This bias arises because changes to care-seeking on early days are more frequently observed, due to the fact that they occur earlier in the spell than care-required days.

This is an important methodological point: in order to estimate the unbiased effect of a policy intervention on health care demand conditional on need for care, it is imperative to collect data at the household level and estimate demand conditional on health status with a hazard model. In a population like ours where underuse is rampant, data collected at the point of use (e.g. the healthcare provider or insurer) is subject to selection effects that are so strong that they reverse the estimated impact of free care on some aspects of care-seeking behavior.

Application: Predicting Care-Seeking for an Ebola Spell An important reason to estimate demand conditional on health status is the ability to conduct counterfactual policy analysis and make accurate out-of-sample predictions. As an example, we predict care-seeking behavior under our different treatments for a hypothetical outbreak of Ebola. There is no cure for Ebola, although early supportive care can sometimes prevent fatal outcomes. However, care-seeking is essential for case-tracking and containment of the virus, and early detection can alter the course of the epidemic. This is a situation with

Figure 4.2: Effect of free care and healthworker treatments on predicted average care-probabilities for ten spells of hemorrhagic fever (Ebola) for the median child. Sources of spell descriptions and additional models in Appendix.



potentially high welfare losses from underuse, given the large externality of infecting others.

We estimate four different specifications of a Cox hazard model that include child and household controls and indicator variables that describe the observed spells, interacted with treatment dummies. The models differ by how the symptoms in a spell are coded (see appendix for details). Figure 4.2 shows the survival rates – i.e. the cumulative probabilities of not seeking care – for the simplest and most complex model by treatment group. Model 1 simply uses “early” vs. “care required” indicators according to the C-IMCI, identical to specification (4) of table 8, whereas Model 4 uses a flexible set of indicators for different symptom combinations.²⁴

The figure shows that free care substantially increases care-seeking for the average Ebola patient. By day 6, when many of the most severe symptoms appear (e.g. a characteristic rash, vomiting and diarrhea, and hemorrhaging), a child in one of the free-care groups would be up to 20% more likely to have received care. Since early fluid replacement and respiratory support can decide over life and death, and bodily fluids are highly infectious, this suggests a high potential for avoided mortality as well as reduced contamination risk for caretakers and family members. Nonetheless, access to care is incomplete – successful containment would likely require aggressive supplemental policies. This example

²⁴It should be noted that most spell descriptions come from adult patients who received (some) treatment. Ebola is arguably so aggressive that most cases eventually receive care; treatment is largely supportive. Nonetheless it is possible that this causes some bias in our predictions – see appendix for more detail.

demonstrates how our method of data collection could be used to simulate policy effects for a range of possible public health scenarios and the associated welfare consequences.

4.4 Demand Effects on Sickness

So far we have shown that free care dramatically increases the value of care consumed, but not how this affects health.

Note that access to health care in our setting encompasses both evaluation by a doctor, and, conditional on evaluation, treatment. An important benefit is to detect (low-frequency) cases where treatment can prevent serious long-term harm or death, as with malaria. Our study was not designed to spot the unconditional impact of the program on these longterm health outcomes.²⁵ Instead, we focus on the unique advantage of our data, namely that we observe the actual day-to-day incidence of illness, a measure of the immediate discomfort experienced. Conditional on a spell occurring, we expect that treatment leads to faster recovery and a reduction in symptoms. Moreover, in our symptom calendar we also recorded the mother’s concern about the child, which serves as a proxy for the (expected) utility experienced by the child’s caregiver. Table 10 shows estimates of the effect of free care on the average duration of an illness spell (including days after care was sought) as well as on the daily probability of the mother being ‘concerned’ or ‘very concerned’ about her child’s health. Regressions (3) and (4) include all days, (5) and (6) only days with symptoms. Each pair of regression equations first shows a simple OLS estimate, and then an OLS with household and child control variables. In (5) and (6), we can include an “early” indicator, and an interaction of free care with “early”. All standard errors are clustered at the compound level.

The OLS estimates suggest that being in one of the free-care groups is associated with slightly shorter average spell length, although the effect is only significant at the 10% level. Since only a small percentage of spells receive formal care, the effect of care on the individual spell is likely larger (but note that there is selection into care). The effect size is consistent with findings from Tanzania (Adhvaryu and Nyshadham [2015]). Free care also significantly reduces the average proportion of illness days when the mother is concerned or very concerned about the child, from 30% to about 22%. The effect is smaller

²⁵Unconditional illness incidence is driven in large part by infection rates and other exogenous events. As a consequence, the effect of the program on average days of sickness is expected to be small. Indeed, the number of spells in the four groups is nearly identical (table 9), and the full program group actually has the *highest* average number of sickness days of the four treatment groups (though this is driven by a higher incidence of wounds, burns and injuries, see table 11 in the appendix).

Table 10: Effect of free care on average illness spell length and mother concern.

	Spells		All days		Illness days only	
	Illness spell length		Mother concerned		Mother concerned	
	(1)	(2)	(3)	(4)	(5)	(6)
Free care (FC)	-0.169 (0.407)	-0.866* (0.513)	-0.0186** (0.00808)	-0.0189** (0.00889)	-0.0827*** (0.0313)	-0.0772* (0.0453)
Early x FC (EFC)					0.0366 (0.0367)	0.0322 (0.0390)
Early					-0.119*** (0.0273)	-0.114*** (0.0285)
Constant	7.357*** (0.328)	6.313*** (1.363)	0.0734*** (0.00663)	0.0461* (0.0247)	0.291*** (0.0252)	0.300*** (0.104)
Controls	-	Yes	-	Yes	-	Yes
Observations	3,144	2,809	98,537	85,299	27,859	25,069
R-squared	0.000	0.007	0.001	0.006	0.022	0.029

*** p<0.01, ** p<0.05, * p<0.1; clustered standard errors in parentheses. Controls: child and household characteristics.

but still significant when including all observed days. Note also that mother concern is significantly lower on early days – this is consistent with the idea that mothers have some information about the child’s health based on observed symptoms.

5 Conclusion

The policy debate around the subsidization of healthcare has been subject to large swings. The Bamako consensus in the 1980s, which led to a broad move towards user-fee financed healthcare systems, is in the process of reversing: several aid organizations now advocate free care (UK Secretary of State for International Development [2009], Save the Children [2008]), and a number of African countries have introduced (partially) free health care especially for mothers and small children (Ridde and Morestin [2012], Yates [2007]) in an effort to reduce high mortality and morbidity rates in these populations. However, as others have pointed out (Powell-Jackson et al. [2014]), there is relatively little convincing evidence of the effect of abolishing user fees on utilization or ultimately health outcomes.

This study aims to open the “black box” of healthcare demand. We model the dynamic problem that parents face when choosing to visit a doctor, estimate healthcare use conditional on health status to remove confounding effects of illness incidence, and benchmark care-seeking against the need for care according to WHO medical standards. We show that the demand analysis needs to account for spell censoring due to spontaneous recovery to avoid biased effect estimates.

Our results have encouraging implications for the user-fee debate. The welfare cost of subsidies could be prohibitive if they lead to substantial mis-targeting and overuse of

healthcare, risking to overburden the severely limited capacity of healthcare systems in poor countries and thus diverting resources away from those truly in need. However, we find that healthcare usage nearly triples in response to subsidies, but does not result in substantial overuse, because demand comes largely from children for whom medical care is required according to WHO guidelines. At the same time, the families in our study are quite adept at recognizing their child's need for healthcare, so that additional information policies do not markedly improve the allocation of care. Overall, it appears that overuse and moral hazard or insufficient health knowledge are not a primary concern in this population, whereas underuse due to diverging spending priorities between parents and policymaker remains high. Thus, policy efforts should be focused on increasing utilization when care is required, rather than further curbing the existing low levels of overuse.

References

- Achyuta Adhvaryu and Anant Nyshadham. Return to treatment in the formal health care sector: Evidence from Tanzania. *American Economic Journal: Economic Policy*, 7(3): 29–57, 2015.
- John Akin, Nancy Birdsall, and David de Ferranti. *Financing Health Services in Developing Countries: An Agenda for Reform*, volume 34. World Bank, 1987.
- John S. Akin, Charles C. Griffin, David K. Guilkey, and Barry M. Popkin. The demand for primary health care services in the Bicol region of the Philippines. *Economic Development and Cultural Change*, 34(4):755–782, 1986.
- Nava Ashraf, James Berry, and Jesse M. Shapiro. Can higher prices stimulate product use? Evidence from a field experiment in Zambia. *American Economic Review*, 100(5): 2383–2413, 2010.
- Nava Ashraf, B. Kelsey Jack, and Emir Kamenica. Information and subsidies: Complements or substitutes? *Journal of Economic Behavior and Organization*, 88:133–139, 2013.
- Zulfiqar A Bhutta, Samana Ali, Simon Cousens, Talaha M Ali, Batool Azra Haider, Arjumand Rizvi, Pius Okong, Shereen Z Bhutta, and Robert E Black. Interventions to address maternal, newborn, and child survival: what difference can integrated primary health care strategies make? *The Lancet*, 372(9642):972–989, 2008.
- Marie-Jo Bonnet, Philippe Akamituna, and Anicet Mazaya. Unrecognized Ebola hemorrhagic fever at Mosongo hospital during the 1995 epidemic in Kikwit, Democratic Republic of the Congo. *Emerging Infectious Diseases*, 4(3):508–510, 1998.
- Jennifer Bryce, Cesar G Victora, Jean-Pierre Habicht, J Patrick Vaughan, and Robert E Black. The multi-country evaluation of the integrated management of childhood illness strategy: Lessons for the evaluation of public health interventions. *American Journal of Public Health*, 94(3):406–415, 2004.
- Jessica Cohen, Pascaline Dupas, and Simone Schaner. Price subsidies, diagnostic tests, and targeting of malaria treatment: evidence from a randomized controlled trial. *American Economic Review*, 105(2):609–645, 2015.
- Janet Currie, Wanchuan Lin, and Wei Zhang. Patient knowledge and antibiotic abuse: evidence from an audit study in China. *Journal of Health Economics*, 30:933–949, 2011.
- Jishnu Das. The quality of medical care in low-income countries: From providers to markets. *PLoS Medicine*, 8(4):1–2, 2011.
- Jishnu Das and Jeffrey Hammer. Which doctor? combining vignettes and item response to

- measure clinical competence. *Journal of Development Economics*, 78(2):348–383, 2005.
- Jishnu Das and Jeffrey Hammer. Location, location, location: residence, wealth, and the quality of medical care in Delhi, India. *Health Affairs*, 26(3):w338–w351, 2007.
- Jishnu Das and Jeffrey Hammer. Quality of primary care in low-income countries: Facts and economics. *Annual Review of Economics*, 6:525–553, 2014.
- Jishnu Das and Thomas Pave Sohnesen. *Patient Satisfaction, Doctor Effort and Interview Location: Evidence from Paraguay*. World Bank, 2006.
- Jishnu Das, Jeffrey Hammer, and Carolina Sánchez-Peramo. The impact of recall periods on reported morbidity and health seeking behavior. *World Bank Policy Research Working Paper*, 5778, 2011.
- Jishnu Das, Alaka Holla, Aakash Mohpal, and Karthik Muralidharan. Quality and accountability in healthcare delivery: Audit-study evidence from primary care in India. *NBER Working Paper*, 21405, 2015.
- W. Dow, P. Gertler, R.F. Schoeni, J. Strauss, and D. Thomas. Health care prices, health and labor outcomes: experimental evidence. *RAND Working Paper*, 97-01, 2001.
- Pascaline Dupas. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228, 2014.
- Pascaline Dupas and Jessica Cohen. Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics*, 125(1): 1–45, February 2010.
- R. T. D. Edmond, Brandon Evans, E. T. W. Bowen, and G. Lloyd. A case of Ebola virus infection. *British Medical Journal*, 2:541–544, 1977.
- Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. The Oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics*, 127(3):1057–1106, 2012.
- Greg Fischer, Dean Karlan, Margaret McConnell, and Pia Raffler. To charge or not to charge: evidence from a health products experiment in Uganda. *NBER Working Paper*, 20170, 2014.
- Pierre Formenty, Christophe Hatz, Bernard Le Guenno, Agnès Stoll, Philipp Rogenmoser, and Andreas Widmer. Human infection due to ebola virus, subtype côte d’ivoire: Clinical and biologic presentation. *Journal of Infectious Diseases*, 179(Supplement 1):S48–S53, 1999.
- Andrew D. Foster. Prices, credit markets and child growth in low-income rural areas. *Economic Journal*, 105(430):551–570, 1995.

- J. S. Gear, G. A. Cassel, A. J. Gear, B. Trappler, L. Clausen, A. M. Meyers, M. C. Kew, R. Bothwell, T. H. and Sher, G. B. Miller, J. Schneider, H. J. Koornhof, E. D. Gomperts, M. Isaäcson, and J. H. Gear. Outbreak of Marburg virus disease in Johannesburg. *British Medical Journal*, 4(5995):489–493, 1975.
- P. Gertler and J. Van der Gaag. *The willingness to pay for medical care: evidence from two developing countries*. Baltimore and London, 1990.
- Paul Gertler, Luis Locay, and Warren Sanderson. Are user fees regressive? The welfare implications of health care financing proposals in Peru. *Journal of Econometrics*, 36(1-2):67–88, 1987.
- Donna B. Gilleskie. A dynamic stochastic model of medical care use and work absence. *Econometrica*, 66(1):1–45, 1998.
- Mark C. Hornbeck, Arnold V. Hurtado, and Richard E. Johnson. Health care episodes: definition, measurement and use. *Medical Care Research and Review*, 42(2):163–218, 1985.
- Emmanuel Jimenez. *Human and physical infrastructure: public investment and pricing policies in developing countries*, volume 3, chapter 43, pages 2773–2843. Elsevier, 1995.
- Andrew M. Jones. Health econometrics. In A. J. Culyer and Joseph P. Newhouse, editors, *Handbook of Health Economics*, volume 1, chapter 6, pages 267–344. Elsevier, 2000.
- Emmett B. Keeler, Joan L. Buchanan, John E. Rolph, Janet M. Hanley, and David M. Reboussin. The demand for episodes of medical treatment in the health insurance experiment. *RAND Working Paper*, 1988.
- Gary King, Emmanuela Gakidou, Kosuke Imai, Jason Lakin, Ryan T. Moore, Clayton Nall, Nirmala Ravishankar, Manett Vargas, Martha M. Téllez-Rojo, Juan E. Hernández Ávila, Mauricio Hernández Ávila, and Hector Hernández Llamas. Public policy for the poor? a randomised assessment of the Mexican universal health insurance programme. *The Lancet*, 373(9673):1447–1454, 2009.
- Amanda Kowalski. Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care. *Journal of Business and Economic Statistics*, forthcoming.
- Michael Kremer and Rachel Glennerster. Improving health in developing countries: evidence from randomized evaluations. In *Handbook of Health Economics*, volume 2, chapter 4, pages 201–315. Elsevier, 2012.
- Kenneth L. Leonard and Melkiory C. Masatu. Using the Hawthorne effect to examine the gap between a doctor’s best possible practice and actual performance. *Journal of Development Economics*, 93:226–234, 2010.

- K.L. Leonard, G.R. Mliga, and D. Haile Mariam. Bypassing health centres in tanzania: revealed preferences for quality. *Journal of African Economies*, 11(4):441, 2002.
- Jenny I. Litvack and Claude Bodart. User fees plus quality equals improved access to health care: Results of a field experiment in cameroon. *Social Science and Medicine*, 37(3):369 – 383, 1993.
- Willard G. Manning, Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz. Health insurance and the demand for medical care: Evidence from a randomized experiment. *American Economic Review*, 77(3):251–277, 1987.
- Gustav Adolf Martini. *Marburg Virus Disease*, chapter Marburg Virus Disease. Clinical Syndrome, pages 1–9. Springer Berlin Heidelberg, Berlin, Heidelberg, 1971.
- Diane McIntyre, Michael Thiede, Göran Dahlgren, and Margaret Whitehead. What are the economic consequences for households of illness and of paying for health care in low- and middle-income country contexts? *Social Science and Medicine*, 62(4):858–865, 2006.
- E. Miguel and M. Kremer. Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217, 2004.
- J. Nabyonga, M. Desmet, H. Karamagi, P. Y. Kadama, F. G. Omaswa, and O. Walker. Abolition of cost-sharing is pro-poor: evidence from Uganda. *Health Policy and Planning*, 20(2):100–108, 2005.
- Hoa L. Nguyen, Jane S. Saczynski, Joel M. Gore, and Robert J. Goldberg. Age and sex differences in duration of prehospital delay in patients with acute myocardial infarction: A systematic review (pmc3072277). *Circulation: Cardiovascular Quality and Outcomes*, 3:82–92, 2010.
- J. Manuel Parra, Octavio J. Salmerón, and María Velasco. The first case of Ebola virus disease acquired outside Africa. *New England Journal of Medicine*, 371:2439–2440, December 2014.
- Timothy Powell-Jackson, Kara Hanson, Christopher J.M. Whitty, and Evelyn K. Ansah. Who benefits from free healthcare? evidence from a randomized experiment in Ghana. *Journal of Development Economics*, 107:305–319, 2014.
- Valéry Ridde and Florence Morestin. A scoping review of the literature on the abolition of user fees in health care services in Africa. *Health Policy and Planning*, 27(3):1–11, 2012.
- Alfonso Rosales and Kristin Weinbauer. *C-IMCI: Community Integrated Management of Childhood Illness*. Catholic Relief Services, 2003.
- David E. Sahn, Stephen D. Younger, and Garance Genicot. The demand for health care services in rural Tanzania. *Oxford Bulletin of Economics and Statistics*, 65(2):241–260,

2003.

João Santos Silva and Frank Windmeijer. Two-part multiple spell models for health care demand. *Journal of Econometrics*, 104:67–89, 2001.

Save the Children. *Freeing up Healthcare: A guide to removing user fees*. McPake, Barbara and Schmidt, Alice and Araujo, Edson and Kirunga-Tashobya, Christine, London, 2008.

W. Stille and E. Boehle. *Marburg Virus Disease*, chapter Clinical Course and Prognosis of Marburg Virus ("Green-Monkey") Disease, pages 10–18. Springer Berlin Heidelberg, Berlin, Heidelberg, 1971.

G. L. Stoddart and M. L. Barer. Analysis of demand and utilization through episodes of medical services. *McMaster University working paper*, 81-06, 1981.

Dag Gundersen Storla, Solomon Yimer, and Gunnar Aksel Bjune. A systematic review of delay in the diagnosis and treatment of tuberculosis. *BMC Public Health*, 8(15), 2008.

Rebecca L. Thornton, Laurel E. Hatt, Erica M. Field, Mursaleena Islam, Freddy Solís Diaz, and Martha Azucena González. Social security health insurance for the informal sector in Nicaragua: a randomized evaluation. *Health Economics*, 19(S1):181–206, 2010.

K. Todorovitch, M. Mocitch, and R. Klačnja. *Marburg Virus Disease*, chapter Clinical Picture of Two Patients Infected by the Marburg Vervet Virus, pages 19–23. Springer Berlin Heidelberg, Berlin, Heidelberg, 1971.

UK Secretary of State for International Development. *Eliminating World Poverty: Building our Common Future*. Department for International Development White Paper, London, 2009.

Unicef Health Section. *Integrated Management of Childhood Illness: Household and Community Component*. Kenya-Mugisha, Nathan and Pangu, Kasa, 1999.

United Nations Inter-Agency and Expert Group on MDG Indicators. *The Millennium Development Goals Report 2015*. United Nations, 2015.

WHO. *Integrated Management of Childhood Illness: Chart Booklet*. World Health Organization, 2014.

WHO Department of Child and Adolescent Health and Development. *Handbook IMCI: Integrated Management of Childhood Illness*. World Health Organization and Unicef, 2005.

WHO Global Malaria Programme. *World Malaria Report 2014*. World Health Organization, 2015.

Rob Yates. The impact of abolishing user fees in Africa – recent developments in six African countries. *SSRN Working Paper*, June 2007.

ONLINE APPENDIX: NOT FOR PUBLICATION

A Examples and Proofs

Proof of Lemma 1

Note that Bayesian updating implies

$$P_{t+1}(I_n) = \frac{\pi(I_n)P_t(I_n)}{\sum_{m=1}^N \pi(I_m)P_t(I_m)} = \frac{\pi(I_n)P_t(I_n)}{\pi_t(P_t)}.$$

Thus, the probability of illness I_n (strictly) decreases if its recovery rate is above average, and increases if it is below average. This means that the time- t distribution of recovery rates first-order stochastically dominates the time $t + 1$ distribution, and $\pi_t(P_t)$ must increase over time.

Proof of Proposition 1

We show that the parent has no desire to deviate from this strategy given that it is used in all future periods. First consider the choice of the parent when beliefs first cross the threshold, so that $\pi_t(P_t) < K$ but $\pi_{t+1}(P_{t+1}) \geq K$. In this case, they can either choose to go to the doctor in the current period, and pay cost $-C$, or wait and follow the equilibrium strategy to go in the next period if the child is still sick. The latter has an expected cost of $\pi_t(P_t)\delta(-S - C)$, so $\pi_t(P_t) < K$ ensures that it is optimal not to consult a doctor today. Since $\pi_{t+1}(P_{t+1}) \geq K$ and π_t weakly increases over time, the same logic ensures that it is also optimal to go immediately in $t + 1$ and any period after. Now consider $t - 1$. Here, the choice is between going immediately or waiting for two periods before receiving treatment. The utility from waiting is

$$\begin{aligned} \pi(P_{t-1})\delta(-S + \pi(P_t)\delta(-S - C)) &< \\ \pi(P_{t-1})\delta(-S - C) &\leq -C \end{aligned}$$

where the first inequality uses that $\pi(P_t)\pi_t$ is below the threshold K . Iterating this argument shows that it is optimal to wait in all earlier periods $t - k < t - 1$ where $\pi_{t-k}(P_{t-k})$ is below K . \square

To generate the examples in figure 2.1, we initially assume I_S and I_L are equally likely, so that $\pi_t = \frac{1}{2}(\pi(I_S) + \pi(I_L))$. Letting $\epsilon > 0$ such that $\pi(I_S) = \pi_t + \epsilon$ and $\pi(I_L) = \pi_t - \epsilon$,

Bayesian updating implies that $\pi_t < \pi_{t+1} = \frac{\frac{1}{2}\pi(I_S)^2 + \frac{1}{2}\pi(I_L)^2}{\pi_t} = \pi_t + \frac{\epsilon^2}{\pi_t}$. This provides the path of beliefs in panel A. For panel B we, assume that under γ^{*L} there is an initial probability of illness I_S of 0.25, and a probability of illness I_L of 0.75. It can be shown that $\pi_t^{*L} = \pi(P_t(\gamma^{*L})) < \pi_t$, and $\pi_t^{*L} < \pi_{t+1}^{*L} = \pi(P_{t+1}(\gamma^{*L})) < \pi_t, \pi_{t+1}$. To the policymaker observing the low severity symptom, the probability of not recovering is lower than to the parent in either period. Under γ^{*S} , the probabilities of illnesses I_S and I_L are reversed, so that $\pi_t^{*S} > \pi_t$ and $\pi_{t+1}^{*S} > \pi_{t+1}$. This implies that $\pi_t^{*L} = \pi(P_t(\gamma^L)) = \pi_t - \frac{1}{2}\epsilon < \pi_t$, and $\pi_{t+1}^{*L} = \pi(P_{t+1}(\gamma^L)) = \pi_{t+1} - \frac{\epsilon}{\pi_t} < \pi_t$: the probability of not recovering is increasing, but actually lower than the initial π_t in either period. Under γ^{*S} , the probabilities of illnesses I_S and I_L are reversed, so that $\pi_t^{*S} = \pi(P_t(\gamma^{*S})) = \pi_t + \frac{1}{2}\epsilon$ and $\pi_{t+1}^{*S} = \pi(P_{t+1}(\gamma^{*S})) = \pi_{t+1} + \frac{\epsilon}{\pi_t}$.

B Health Diary, Symptoms, and IMCI Classifications

Figure B.1 shows the health diary. The diary has entries for nine major symptoms, the mother's concern about the child, and doctor and pharmacy visits.

The full list of 12 symptoms were chosen based on the Integrated Management of Childhood Illness (IMCI) guidelines (see WHO Department of Child and Adolescent Health and Development [2005], WHO [2014]). IMCI is a joint program of the World Health Organization and UNICEF and aims at establishing simple procedures for the management of the main causes of morbidity and mortality of small children. The guidelines differentiate between low and high malaria regions. Rosales and Weinbauer [2003] have adapted the IMCI for use by community healthworkers, who typically have no formal medical training, but can be instructed to follow simple protocols for the basic diagnosis of illness. These protocols are mainly designed to detect a need for formal medical evaluation. Symptoms are classified into acute danger signs and symptoms that point to a particular illness or class of illnesses, e.g. malaria or generalized fever. The choice of symptoms collected was based on ease of observation and description for both surveyors and mothers, and explicit rules found in the IMCI guidelines for referral to formal care. The symptoms are:

- Convulsions, fits, or spasms (danger sign)
- Lethargic or less conscious (danger sign)
- Unable to drink or breastfeed (danger sign)
- Vomiting everything (danger sign)
- Coughing (respiratory disease)
- Difficulty breathing (respiratory disease)

Nom de l'enfant : _____ ID Ménage : _____ ID/Nom Enqueteur : _____ Semaine : _____

		ENFANT				
		JOURS				
SYMPTÔMES						
Jour						
Mère été inquiète pour la santé de l'enfant						
Convulsions, crises, ou spasmes						
Léthargique/ moins conscient						
Refuse l'allaitement ou de boire						
Vomir tout						
Toux						
Difficulté à respirer						
Diarrhée						
Peau chaude à toucher						
Médecin visité						
Pharmacien visité						

Figure B.1: Sample of the health diary. Along the top, the surveyors enter the days of the week. The last day of each calendar, marked by a “writing” symbol, corresponds with the visit of the surveyor, the first day corresponds with the previous visit and used to record events occurring after the surveyor leaves. Mothers were asked to use the diaries as a memory aid, but records were also taken if the mother did not fill in the health diary.

- Diarrhea – If diarrhea: more than three loose stools per day? (diarrheal disease)
- If diarrhea: blood in the stool? (indicator for dysentery)
- If diarrhea: sunken eyes? (indicator for dehydration)
- Unusually hot skin (under 2 months age: unusually cold skin) (fever)

Additionally, we manually classified symptoms recorded under “other health changes” into “cold symptoms”, “ear pain”, “skin rash”, “wound or injury”, “unusually hot” (typically recorded as “fever” or “malaria”), “head/neck/eye pain”, “stomach pain”, and “other”.

Table 11 shows the frequency of reported symptoms by randomized treatment group and in total. The full treatment group experiences or reports more illness compared to the control, significant at the 10% level, although some symptoms are also reported less often, for example “sunken eyes”. This difference may be a statistical accident, but may also be an effect of the treatment if both free care and healthworkers increase parents’ attention to their child’s symptoms. It would be hard to say if symptoms in the full treatment group are “over-reported” or in the control “underreported”. Note, however, that the difference is driven by injuries and cold symptoms; Mali Health does not pay for healthcare related to either. Moreover, the incidence of illness *spells* is almost identical in all treatment groups (see table 6).

Based on the C-IMCI, we used guidelines on urgent or non-urgent referral to a clinic to determine when a child should seek formal care. The following rules were applied:

- Any of the danger signs require immediate (same day) care.
- Diarrhea with blood in the stool or sunken eyes require immediate referral on suspicion of dysentery or severe dehydration, respectively.
- Diarrhea without signs of dysentery or dehydration requires non-urgent referral after at most 5 days of continual illness.
- Fever with a rash and cough or cold symptoms require immediate referral on suspicion of measles.
- Fever without cough, cold symptoms, difficulty breathing, rash, or ear infection requires immediate referral on suspicion of malaria.
- Any other fever requires a non-urgent referral for generalized fever.
- A simple cough requires non-urgent care after 14 days on suspicion of tuberculosis.
- Cold symptoms and difficulty breathing require non-urgent care after 14 days on suspicion of a bacterial rather than viral infection.
- Ear pain should lead to non-urgent referral for acute or chronic ear infection.

In this context, a non-urgent referral is interpreted as “within 24 hours”, that is, at least on the next day. An immediate referral is interpreted to mean on the same day.

Table 11: Symptom incidence by treatment group.

	Control	Health worker	Free care	HW & FC	Total
Symptom days as proportion of observed days per child.					
Any symptom	0.285 (0.258)	0.297 (0.260)	0.292 (0.259)	0.320* (0.257)	0.298 (0.258)
Convulsions, fits, or spasms	0.000 (0.006)	0.002 (0.032)	0.000 (0.003)	0.001* (0.010)	0.001 (0.017)
Lethargic or unconscious	0.013 (0.070)	0.010 (0.034)	0.014 (0.056)	0.019 (0.049)	0.014 (0.054)
Unable to drink or breastfeed	0.005 (0.036)	0.003 (0.016)	0.004 (0.018)	0.003 (0.016)	0.003 (0.023)
Vomiting everything	0.011 (0.030)	0.012 (0.034)	0.010 (0.022)	0.013 (0.029)	0.011 (0.029)
Coughing	0.108 (0.168)	0.119 (0.185)	0.099 (0.164)	0.122 (0.169)	0.112 (0.172)
Difficulty breathing	0.015 (0.051)	0.019 (0.070)	0.013 (0.042)	0.018 (0.046)	0.016 (0.053)
> 3 loose stools	0.024 (0.065)	0.020 (0.055)	0.025 (0.071)	0.020 (0.051)	0.022 (0.061)
Blood in the stool	0.003 (0.015)	0.002 (0.017)	0.001 (0.011)	0.001 (0.011)	0.002 (0.014)
Sunken eyes	0.008 (0.038)	0.005 (0.022)	0.007 (0.030)	0.004* (0.024)	0.006 (0.029)
Unusually hot skin	0.078 (0.115)	0.084 (0.121)	0.088 (0.127)	0.089 (0.120)	0.084 (0.121)
Other: rash, spots, or itch	0.006 (0.037)	0.011* (0.048)	0.011* (0.041)	0.008 (0.037)	0.009 (0.041)
Other: cold symptoms	0.173 (0.223)	0.176 (0.220)	0.181 (0.227)	0.206** (0.235)	0.184 (0.226)
Other: ear ache	0.002 (0.018)	0.003 (0.026)	0.004 (0.024)	0.003 (0.020)	0.003 (0.022)
Other: wound, injury, or burn	0.010 (0.036)	0.011 (0.044)	0.011 (0.049)	0.020*** (0.061)	0.013 (0.048)
Other symptoms	0.016 (0.047)	0.014 (0.047)	0.013 (0.051)	0.014 (0.045)	0.014 (0.048)
Observations					
Total days observed per child	60.115 (9.497)	59.896 (9.172)	59.403 (10.086)	60.31 (8.751)	59.929 (9.393)
Number of children observed N	401	384	390	378	1553

Standard deviations in parentheses. Significance levels: *** 1%, ** 5%, * 10%, mean difference from control.

Note that cough or cold symptoms alone, while very frequent, are typically signs of a simple cold and do not require formal care. Note also that we do *not* classify children with wounds or injuries or children with pain symptoms other than ear pain to require formal care. These are symptoms that were reported frequently but are not mentioned in the C-IMCI guidelines. Lastly, we classified Sikoro as a low-malaria region, on the basis that malaria only occurs seasonally and the Malian government mandates malaria testing for all potential malaria cases. In high malaria regions, any fever is treated as likely malaria, regardless of accompanying symptoms and often without additional testing. In low-malaria regions a fever requires medical care at most a day later.

C Additional Tables

C.1 Attrition

Table 12: Children Per Group and Attrition.

	Control	Healthworker	Free care	HW & FC	All
<i>Original Sample</i>	463	433	451	417	1764
Not Found at Baseline	26	24	19	12	81
Moved Post-Baseline	34	23	35	23	115
Died Post-Baseline	0	1	3	1	5
Refused Post-Baseline	0	0	0	1	1
Unexplained Absence	4	6	5	3	18
<i>2013 Sample</i>	399	379	389	377	1544
Total Attrition	13.8%	12.5%	13.7%	9.6%	12.5%
<i>Attrition Post Baseline</i>	8.7%	7.3%	10.0%	6.9%	8.3%

C.2 Utilization and Pre-Care Spell Length

Table 13: The cost of a formal consultation by spell day, spell day X Early, and treatment group. The average “Early” effect on formal care costs is CFA 1803 (significant at 1% level in a regression on spell day, treatment, and one “Early” indicator).

	Total cost of formal visit					
	All formal		Private care only		CSCOM care only	
Healthworker visits (HW)	41.56		-1,677		0.718	
	(773.3)		(2,171)		(651.1)	
Free care (FC)	-581.3		-441.2		-154.6	
	(628.5)		(2,631)		(413.4)	
HW and FC (HWFC)	-381.3		2,653		-217.9	
	(864.0)		(3,490)		(741.8)	
Spell day	Main	x Early Care:	Main	x Early Care:	Main	x Early Care:
Day 1	4,948***	-1,348***	8,424***	-4,118	4,208***	-1,215***
	(696.9)	(507.8)	(2,641)	(2,502)	(418.0)	(452.0)
Day 2	5,445***	-1,169	9,216***	-4,886	4,762***	-915.5
	(763.7)	(760.8)	(3,170)	(5,220)	(548.3)	(657.2)
Day 3	5,499***	-2,156***	10,600***	-5,865*	4,537***	-1,802***
	(753.0)	(704.8)	(2,937)	(2,937)	(616.9)	(661.1)
Day 4	4,057***	-582.4	4,853	-	3,542***	-532.6
	(670.7)	(646.8)	(3,499)	-	(481.8)	(551.7)
Day 5	6,681***	-	14,706***	-	5,407***	-
	(945.6)	-	(2,971)	-	(757.1)	-
Day 6	6,381***	-3,655***	-	-	5,908***	-3,348***
	(1,043)	(1,273)	-	-	(937.6)	(1,222)
Day 7	5,564***	-4,564***	1,000	-	5,248***	-
	(1,111)	(1,111)	(0)	-	(1,015)	-
Day 8	5,138***	-2,557***	6,130***	-	4,891***	-2,726**
	(740.3)	(645.4)	(1,695)	-	(1,268)	(1,229)
Day 9	6,490***	-3,812***	-	-	6,163***	-3,943***
	(1,209)	(1,325)	-	-	(1,144)	(1,330)
Day 10	4,889***	-1,784	5,677**	-	4,748***	-1,892
	(757.3)	(2,005)	(2,171)	-	(742.3)	(1,879)
>10 days	4,795***	-887.5	3,303	-272.4	4,625***	357.2
	(761.6)	(1,334)	(2,645)	(3,247)	(638.7)	(563.5)
Observations	414		46		368	
R-squared	0.664		0.748		0.692	

Table 14: Consultations by “early” and “care required” status, for each treatment group.

Formal visits by cause of visit		ALL	Acute			Follow-up during spell			Follow-up after spell	Prevention/ other
			all	early	care required	all	early	care required		
Control	Avg. visits per child	0.178	0.163	0.030	0.133	0.008	0	0.008	0	0.008
	(Std. error)	(0.015)	(0.020)	(0.009)	(0.018)	(0.006)	(0)	(0.006)	(0)	(0.004)
	Percent of visits		100%	18%	82%	100%	0%	100%		
HW	Avg. visits per child	0.203	0.169	0.029	0.140	0.021	0.005	0.016	0	0.013
	(Std. error)	(0.028)	(0.022)	(0.009)	(0.019)	(0.008)	(0.005)	(0.006)	(0)	(0.006)
	Percent of visits		100%	17%	83%	100%	25%	75%		
Demand in % of control		114%	104%	97%	105%	281%	-	211%	-	175%
FC	Avg. visits per child	0.571	0.488	0.067	0.422	0.059	0	0.059	0.005	0.018
	(Std. error)		(0.034)	(0.013)	(0.032)	(0.014)	(0)	(0.014)	(0.004)	(0.007)
	Percent of visits		100%	14%	86%	100%	0%	100%		
Demand in % of control		321%	300%	222%	317%	786%	-	786%	-	239%
HWFC	Avg. visits per child	0.560	0.443	0.077	0.366	0.090	0.003	0.088	0.005	0.021
	(Std. error)		(0.034)	(0.015)	(0.031)	(0.023)	(0.003)	(0.023)	(0.004)	(0.007)
	Percent of visits		100%	17%	83%	100%	3%	97%		
Demand in % of control		315%	272%	256%	276%	1199%	-	1164%	-	282%

We classify as acute care the first event during a spell. Follow-ups are additional visits during the same spell or visits outside a spell reported as follow-up visits by the parents. Prevention/other are visits not associated with spells and reported as preventative care visits or "other" by the parents.

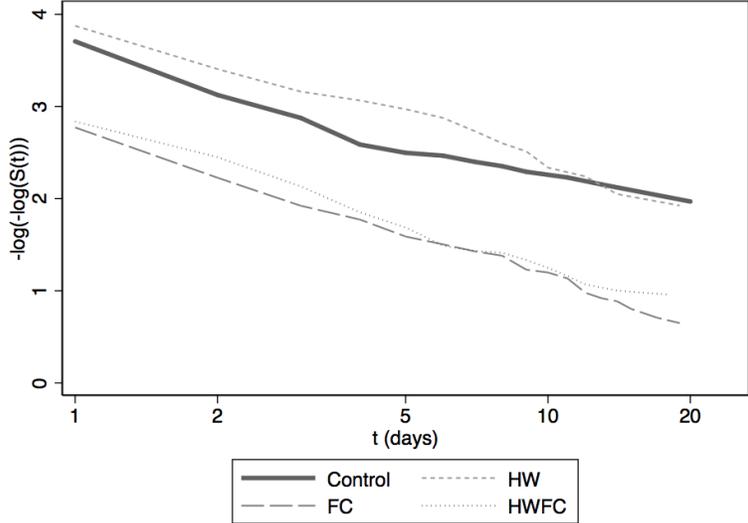
C.3 Proportional Hazard Model

Specification Check Proportional Hazard Model

If the hazard function is indeed proportional, the graph of $-\log(-\log(S(t, x)))$ against the log of time should be parallel for different covariate values. We can check this by graphing the transformed empirical survival curves in the set of all spells for each treatment group. Since treatment assignment was random, initially the unconditional covariate distribution in the four groups should be similar. The test is imperfect if there is heterogeneity in the hazard of care for different population subgroups, since this means that the population composition will change differentially over time. Nonetheless, it can give us some idea how restrictive the proportionality assumption is in our data.

Figure C.1 shows the plot of the transformed survival function against time in log scale. The FC and HWFC curves are nearly identical. The HW and control curves are roughly shifted up in parallel, although the control group seems to show a slow-down of the survival rate at around the fourth spell day, relative to the other groups. Overall, however, the proportionality assumption seems to provide a fairly good approximation of the data, since the difference in slope across control and treatment is small in comparison to the magnitude of the downward shift in $S(t)$ in the groups that receive free care. Before day 10, the HW group also shows higher survival levels than the control group, consistent

Figure C.1: Predicted survivor functions by treatment group



with the idea that the healthworkers delay care seeking.

Cox Hazard Estimates Using All Spells and Only Uncensored Spells

Table 15: Cox model estimates using all spells, including left-censored.

COX HAZARD MODEL	(1)	(2)	(3)	(4)
Early: care not yet required (IMCI)	0.147*** (0.0185)	0.154*** (0.0517)	0.141*** (0.0191)	0.162*** (0.0557)
Healthworker visits (HW)	0.895 (0.188)	0.868 (0.199)	0.779 (0.173)	0.749 (0.179)
Early x HW (EHW)		1.224 (0.552)		1.303 (0.617)
Free care (FC)	2.953*** (0.472)	3.064*** (0.538)	2.639*** (0.428)	2.761*** (0.500)
Early x FC (EFC)		0.787 (0.305)		0.752 (0.299)
HW and FC (HWFC)	0.952 (0.240)	0.939 (0.257)	1.209 (0.324)	1.238 (0.357)
Early x HWFC (EHWFC)		1.083 (0.583)		0.858 (0.488)
Total assets in US\$ (log)			1.052** (0.0229)	1.052** (0.0229)
Dist. to closest formal (log)			0.872 (0.0911)	0.873 (0.0916)
Child controls	-	-	YES	YES
Household head controls	-	-	YES	YES
Observations	24,741	24,741	21,725	21,725
Partial LL	-3455	-3454	-3017	-3016
Wald chi-square	302.7	330.6	303.4	334.6

*Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

Standard errors clustered at the compound level (in parentheses).

Table 16: Cox estimates for uncensored spells only.

COX HAZARD MODEL	(1)	(2)	(3)	(4)
Early: care not yet required (IMCI)	0.397*** (0.0537)	0.472*** (0.133)	0.398*** (0.0580)	0.424*** (0.120)
Healthworker visits (HW)	0.680** (0.123)	0.688* (0.131)	0.550*** (0.117)	0.533*** (0.118)
Early x HW (EHW)		0.957 (0.460)		1.324 (0.739)
Free care (FC)	0.891 (0.113)	0.955 (0.122)	0.795* (0.106)	0.828 (0.117)
Early x FC (EFC)		0.659 (0.244)		0.785 (0.307)
HW and FC (HWFC)	1.364 (0.295)	1.262 (0.292)	1.709** (0.421)	1.692** (0.443)
Early x HWFC (EHWFC)		1.529 (0.890)		0.974 (0.651)
Total assets in US\$ (log)			1.035* (0.0229)	1.035* (0.0229)
Dist. to closest formal (log)			0.927 (0.0700)	0.932 (0.0702)
Child controls	-	-	YES	YES
Household head controls	-	-	YES	YES
Observations	1,969	1,969	1,750	1,750
Partial LL	-2106	-2105	-1811	-1810
Wald chi-square	55.39	60.73	69.98	74.69

Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Standard errors clustered at the compound level (in parentheses).

C.4 Ebola and Marburg Disease Spells; Models, Predictions, and Sources

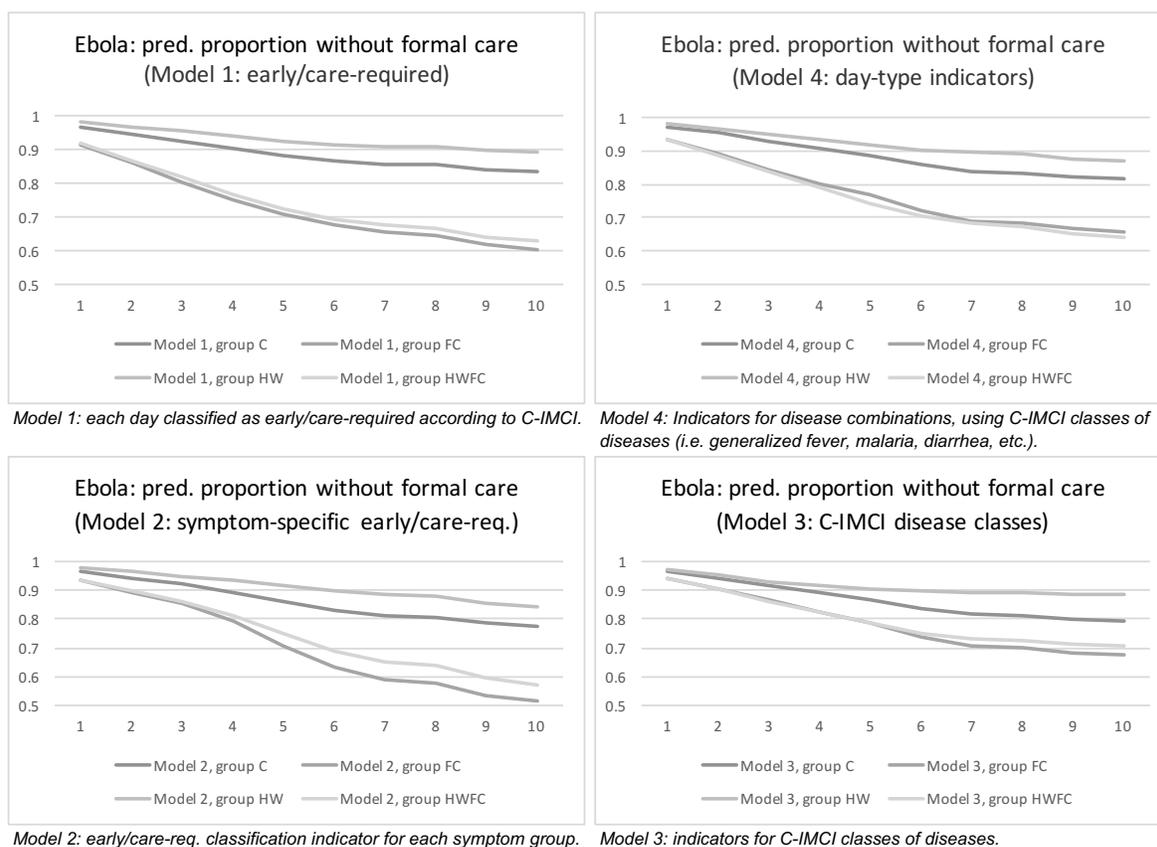
We coded ten spells of hemorrhagic fever, derived from detailed day-to-day descriptions of outbreaks of Marburg disease and Ebola from 1971 to 2015.²⁶ A typical Ebola spell starts suddenly with high fever and headaches, sometimes accompanied by pressure behind the eyes or conjunctivitis. After a few days, severe diarrhea and vomiting set in, and abdominal pain and “tarry” or bloody stool (melena) can occur. A characteristic rash appears around days five to seven, but it is not always recognized. Many patients become listless, apathetic, confused, or unresponsive (asthenia), and (in fatal cases) comatose. Some patients experience swelling and sore throat that causes them difficulty

²⁶Descriptions found in: Bonnet et al. [1998], Edmond et al. [1977], Formenty et al. [1999], Gear et al. [1975], Martini [1971], Parra et al. [2014], Stille and Boehle [1971], Todorovitch et al. [1971].

to swallow (pharyngitis) and triggers a dry cough. We used the symptom descriptions to code these spells as they might have appeared in our data collection. Symptoms are recorded as “lethargic or less conscious”, “unable to drink or breastfeed”, “vomiting everything”, “coughing”, “more than three loose stools per day?”, “blood in the stool”, “sunken eyes”, “unusually hot skin (fever)”, “cold symptoms”, “skin rash”, “head/neck/eye pain” and “stomach pain”. We did not classify other reported symptoms, like other hemorrhaging (nose bleeds, bleeding from puncture wounds, etc.), muscle and joint pains, and swelling, as well as clinically recorded symptoms (e.g. platelet counts or urine color and output).

We estimated four different specifications of a Cox hazard model in our original spell data and then predicted care-seeking hazards for each Ebola/Marburg spell. Each model includes child and household controls and indicator variables that describe the observed spells, interacted with treatment dummies. The models differ by how the symptoms in a spell are coded. Model 1 simply uses “early” vs. “care required” indicators according to the C-IMCI, identical to specification (4) of table 8. Model 2 uses the classification for “care required” within each symptom group described in the C-IMCI. Model 3 uses dummies for symptom groups as used in the C-IMCI (but regardless if this set of symptoms has entered “care required” status) and finally Model 4 uses dummies for individual symptom groups as well as various combinations of these groups, essentially creating a set of dummies for every daily symptom profile that occurs more than a few times in the data.

Figure C.2: Ebola spell survival, all models.



It should be noted that most spell descriptions come from patients who were hospitalized and sometimes received extensive treatment. However, for the purposes of this analysis, Ebola has the “advantage” that the course of the disease is so severe that a large majority of people who fall ill will eventually see a doctor, limiting selection effects. Moreover, Ebola is a virus and does not respond to antibiotics or antimalarials, including quinine (the treatments usually given when the virus is not recognized. Treatment is otherwise largely supportive, ranging from intravenous rehydration to respiratory support). Nonetheless it is possible that our simulations either overstate treatment probabilities, because our Ebola spells are selected to be more serious than the average spell, or understate them, in cases where treatment reduces the severity of symptoms (the latter is unlikely given the high share of fatal outcomes of the cases described).