

Does Patient Demand Contribute to the Overuse of Prescription Drugs?*

Carolina Lopez Anja Sautmann Simone Schaner[†]

May 16, 2019

Abstract

We conducted an experiment in Mali to test whether patients pressure doctors to provide unnecessary medical treatment. The experiment varied patients' information about a discount for antimalarial tablets and measured demand for both tablets and costlier antimalarial injections. We find evidence of patient-driven demand: informing patients about the discount, instead of letting doctors decide to share this information, increased discount use by 35 percent and overall malaria treatment by 10 percent. These marginal patients rarely had malaria, worsening the illness-treatment match. In contrast, we find no evidence that doctors drive demand by exploiting their information advantage to sell more injections.

*We are grateful to Pascaline Dupas, Rema Hanna, seminar participants at the University of Toronto, Loyola Marymount University, MIT, USC, Rice University, UCSD Rady, the Barcelona Summer Forum, the BREAD pre-conference, and the Stanford Rosenkranz Global Health Policy Symposium for many helpful comments. We thank the staff of IPA Mali and especially Alassane Koulibaly for expert field coordination. This project would not have been possible without the advice and support of Dr. Seydou Doumbia and Dr. Issaka Sagara (Malaria Research and Training Center, University of Bamako) as well as Dr. Seydou Fomba (Programme National de Lutte contre le Paludisme (PNLP)). We thank the Ministry of Health and the participating community clinics for allowing us to conduct this research. The study was funded by ESRC/DFID Development Frontiers Award ES/N00583X/1. The research protocol was approved by the IRB of La Faculté de Médecine de Pharmacie et d'Odontostomatologie (FMPOS) de Bamako (protocol No. 2016/47 /CE/FMPOS) and registered in the AEA RCT registry (AEARCTR-0001707). All errors and opinions are our own.

[†]Carolina Lopez: Brown University, carolina_lopez@brown.edu. Anja Sautmann: Massachusetts Institute of Technology, sautmann@mit.edu. Simone Schaner: University of Southern California, BREAD, and NBER, schaner@usc.edu.

1 Introduction

Low income countries must address two opposing challenges when crafting health policy: while many patients do not seek care when they are ill (Whitehead et al., 2001; Sautmann et al., 2017), those who do seek care often consume a great deal of unnecessary treatment (Das and Hammer, 2014; Li et al., 2012). One solution is to subsidize critical treatments in the public sector, while staffing facilities with trained doctors to act as gatekeepers to subsidized care. In practice, however, physicians often fail to live up to their gatekeeping mandate. For example, in an audit study Das et al. (2016) find that 74 percent of public sector providers in India dispensed at least one unnecessary treatment; the Centers for Disease Control and Prevention estimate that 30 percent of antibiotic prescriptions in the United States are unnecessary (Fleming-Dutra et al., 2016); and studies across Sub-Saharan Africa have documented very large shares (36-81 percent) of malaria-negative patients receiving antimalarials at health centers (Reyburn et al., 2004; Hamer et al., 2007; Bisoffi et al., 2009; Ansah et al., 2010).

Why are so many doctors such poor gatekeepers? Most existing literature studies *doctor-driven* reasons for overtreatment, including misaligned financial incentives (e.g. Iizuka, 2012; Currie et al., 2014; Das et al., 2016) and knowledge gaps, compounded by low motivation (e.g. Banerjee et al., 2008; Das et al., 2008).¹ Doctors can overprescribe to patients because health care is a credence good: since patients cannot verify the true cause of their illness, they must rely on the doctor’s recommendation. But by a similar argument, there is scope for *patient-driven* demand for unnecessary treatment: if the doctor cannot persuade the patient that a medication is not needed, she may write a prescription just to keep the patient satisfied. While explicit or perceived patient demand is often cited by doctors as a reason for overtreatment (Kotwani et al., 2010; Linder et al., 2014), there has been much less research on this channel.

Our paper aims to fill this gap. We begin by developing a tractable theoretical model of doctor-patient interactions. We assume that doctors and patients may have different preferences about the best course of medical treatment for an illness when they have access to the same information, and may in particular disagree about whether some drug or procedure is truly needed. To open up the possibility of patient-driven demand, we allow for “gatekeeping costs” – time or hassle costs that the doctor incurs if her prescription deviates from the patient’s preferred option.² Within the experiment we designed, the model generates predictions for when and how a doctor will selectively share information with patients to reduce

¹For a detailed review of this literature, see Dupas and Miguel (2017).

²Our conceptualization of gatekeeping costs is related to Prendergast (2003)’s model of bureaucracy, where the risk of complaints distorts the decisions of a bureaucrat tasked with allocating a good.

gatekeeping costs and steer treatment outcomes towards her own preferences. The theory guides our analysis, allowing us to (1) perform a general test for the presence of gatekeeping costs, (2) assess the extent to which increased use of treatment in response to a price subsidy is driven by patient versus doctor preferences, and (3) estimate how this change in demand impacts the misallocation of treatment.

The experiment was conducted with 2,055 patients visiting 60 public-sector health clinics in Bamako, Mali and focuses on (over)treatment for malaria. Mali’s medical guidelines for appropriate malaria treatment are extremely clear: patients should only be prescribed an antimalarial if they have a positive malaria test result (Ministère de la Santé, 2013). The vast majority of malaria patients present with “simple” or “uncomplicated” malaria, which should be treated with artemisinin combination therapy (ACT) tablets. A small share of patients present with “severe” malaria, which requires more intensive (and expensive) medication, usually delivered via injection or IV drip. As in other Sub-Saharan African settings, we document widespread overuse of antimalarials even though diagnostic tests are readily available: 58 percent of malaria-negative control group patients received an antimalarial prescription, and 41 percent of these prescriptions were for severe malaria treatment.

The experiment introduced a discount for simple malaria treatment, while varying (within clinic, across different days) whether this discount was a doctor’s private information or known to both doctor and patient. On “Patient Voucher” days, all patients visiting the health clinic were given a voucher for free treatment for simple malaria. This subsidy increased the attractiveness of simple malaria treatment relative to both no treatment and severe treatment on Control days. On “Doctor Voucher” days, the same free treatment vouchers were available, but were given to doctors to dispense at their discretion. This means the doctors could *choose* to increase the attractiveness of simple treatment by revealing the voucher to the patient. Vouchers in both conditions had identical appearances, terms, and conditions: a voucher could only be used on the day of the visit, was not valid without a doctor’s signature, and was not valid if the patient received severe malaria treatment.³

Our model shows that treatment outcomes will not vary across the Doctor and Patient Voucher conditions *unless there are gatekeeping costs*. Intuitively, the doctor’s financial incentives and the antimalarial subsidy level are the same across the two conditions; consequently the doctor’s preferred action will also be the same. However, if gatekeeping costs are large enough, the doctor will deviate from her preferred prescription when it conflicts with the patient’s preference. A doctor facing gatekeeping costs can exploit the information

³We used clinics’ regular stocks of malaria medicine and reimbursed clinics for all free medication dispensed through the voucher program at the prevailing retail price. This kept the financial incentives for doctors constant across the three treatment groups.

advantage in the Doctor Voucher condition to get closer to her preferred outcome. Consider, for example, a doctor who does not want to prescribe an antimalarial: she may give in to patient pressure for free treatment in the Patient Voucher condition, but can conceal the discount and avoid prescribing in the Doctor Voucher condition. This provides us with a test of patient-driven demand. By contrast, suppose the doctor would like to induce as much treatment as possible, for example to increase clinic revenue: in this case she will reveal the voucher to patients who would otherwise not buy malaria treatment, but conceal the voucher from patients who are prepared to buy the more expensive “severe” malaria treatment. Thus, we can also test for doctor-driven demand.

We find robust evidence that gatekeeping costs affect doctors’ choices: patients are nine percentage points (35 percent) more likely to redeem a voucher on Patient as compared to Doctor Voucher days. This implies that doctor and patient preferences are frequently misaligned. More voucher use under the Patient Voucher condition could reflect higher total consumption of antimalarials (consistent with patient-driven demand) or lower consumption of severe malaria treatment (consistent with doctor-driven demand). In practice, we find consistent evidence that excess demand for antimalarials on Patient Voucher days is driven by *patient preferences*.

First, doctors are 4 percentage points more likely to write an antimalarial prescription (for simple or severe treatment) on Patient Voucher versus Doctor Voucher days, and patients are 6 percentage points more likely to purchase malaria treatment. The excess demand for antimalarials on Patient Voucher days is entirely driven by patients with the fewest malaria symptoms – antimalarial purchases by this group are a striking 9 percentage points (23 percent) higher on Patient Voucher days. This supports the hypothesis that doctors prefer to withhold treatment from these marginal low-risk cases, but are willing to accede to demand when patients know that treatment is free. Moreover, there is no evidence that these additional purchases are stored for future use, so they likely constitute overuse of malaria treatment.

Second, we find no evidence that doctors strategically withhold information about vouchers in order to “upsell” patients the more expensive severe malaria treatment. Use of severe malaria treatment decreases in both the Doctor Voucher condition and the Patient Voucher condition. We also find no evidence of increased revenue in the Doctor Voucher condition.⁴

Finally, we use data from follow-up home-based malaria tests with patients to assess the match between underlying illness and treatment. In the Control group, we estimate that 57 percent of patients received “appropriate” malaria treatment, in that they either were

⁴Although doctors in our setting are salaried, they have a financial incentive to increase revenues because revenues support the continued functioning of the clinic. See Section 2 for more detail.

malaria positive and purchased malaria treatment, or were negative and did not purchase treatment. The Patient Voucher condition reduced the match by 8 percentage points, while the Doctor Voucher condition reduced the match by 4 percentage points.⁵ The match quality reduction in the Doctor Voucher condition is entirely driven by patient behavior. While doctors' prescription writing does not change, there is an increase in antimalarial purchases conditional on receiving a prescription (some patients do not purchase prescribed medication, which in our context works to improve the match, since most patients do not actually have malaria). In contrast, in the Patient Voucher condition, over half of the worsening of the match is driven by a change in doctors' prescription-writing behavior. These results underscore why understanding doctors' gatekeeping ability is key for informing policy: if doctors were perfect gatekeepers, subsidies would always improve the match, because they would exclusively increase use of antimalarials by malaria-positive patients.

We find no evidence that our results are driven by Hawthorne effects or experimenter demand effects. On the physician side, Hawthorne effects should work to increase physician compliance with official malaria treatment guidelines, which would reduce differences across the voucher conditions. On the patient side, our voucher distribution protocols were specifically designed to avoid signaling a need for treatment. Moreover, evidence from a separate intervention indicates that patient actions did not reflect inferred "cues" from the experiment: a cross-randomized patient information treatment that emphasized the importance of taking a malaria test before purchasing treatment did not increase rates of test taking.

The main contribution of this study is to provide some of the first well-identified evidence of patient-driven demand for excess medical treatment. Most directly related to our work are Kravitz et al. (2005) and Currie et al. (2014), who conduct audit studies where "simulated patients" or SPs (trained actors who are not actually sick) request a specific treatment. Kravitz et al. find that doctors are more likely to prescribe antidepressants when SPs ask for them. Currie et al. find that direct SP requests for an antibiotic increased prescription rates at hospitals in China, although this effect disappeared when SPs indicated they would purchase their medications outside the hospital pharmacy. Audit studies such as these show that doctors at least sometimes comply with increased demand for treatment. This is what we also see in the Patient Voucher condition. However, on its own, doctor compliance with patient demand does not isolate *patient-driven* demand: the physician might be selling the treatment she wanted to sell all along, or act against her own preferences in order to avoid gatekeeping costs. Our design can separately identify the latter because the Doctor Voucher condition allows doctors to manipulate patient preferences by withholding information about the price discount; put another way, doctors who want to sell more can share information

⁵We reject equality of these two treatment effects at the one percent level.

about the voucher, while doctors who want to sell less can conceal the voucher. Our experiment also has the advantage that it uses real patients, drawing from the existing distribution of patient symptoms, demographics, and illnesses .⁶

Note that the prevalence of patient-driven demand in our context does not preclude that some share of overprescription by doctors happens in “collusion” with patients (e.g. both doctor and patient prefer an antimalarial prescription despite a negative malaria test). As the findings by Currie et al. suggest, doctors might fulfill patient demand because it is in their financial interest; other reasons could be knowledge gaps or low diagnostic effort. This may explain some of the high “baseline” level of overtreatment in our setting.

Our work also contributes to a growing literature on the pricing of health goods in developing countries. Most recent work in this area focuses on preventive health products like mosquito bednets and water purification, where beneficiary demand is low, use generates positive externalities, and large subsidies typically improve outcomes.⁷ Use of curative products, by contrast, can generate negative externalities (e.g. by contributing to disease resistance (WHO, 2014)). Recent work on the pricing of curative health products has found mixed results. Cohen et al. (2015) find that subsidizing antimalarials at private-sector “drug shops” dramatically improves access to lifesaving medication, but making the subsidy too high generates large amounts of overtreatment with only marginal gains in access. Importantly, these shops were staffed by retail workers, who were not expected to perform a gatekeeping function. In the public sector, Sautmann et al. (2017) find that subsidizing health care for children (combined with close cost control) can improve access at the extensive margin significantly. Our paper contributes by offering evidence on misallocation – and its drivers – conditional on seeking care (the intensive margin).

Our results highlight that the value of subsidies in the public sector will depend on the extent to which doctors serve as effective gatekeepers. It is therefore critical for policy makers to understand both whether doctors are suboptimal gatekeepers and *why* they fall short. If it is patients who drive overtreatment, then interventions that strengthen doctor authority and reduce gatekeeping costs (like patient communication tools) could help sustain subsidies while limiting overtreatment. In contrast, if doctors drive overuse of medication, policy should aim to limit doctor discretion and empower patients, unless doctors’ preferences can be realigned with proper treatment protocol.

The rest of the paper is structured as follows: in the next section we provide background

⁶Although slightly more patients visit the clinic on voucher as compared to Control days, there is no significant difference in the number of patients attending the clinic on Doctor versus Patient Voucher days. Further, average patient characteristics are well balanced across Doctor and Patient Voucher days.

⁷For bednets see Cohen and Dupas (2010), (Dupas, 2014), and (Tarozzi et al., 2014). For water purifications see (Ashraf et al., 2010) and (Kremer et al., 2011).

on health care and malaria treatment in Mali and present some stylized features of status quo malaria treatment. With this in mind, we develop a model of doctor-patient interactions in section 3. Section 4 describes our experimental design, section 5 presents our empirical results, and section 6 concludes.

2 Background

2.1 Health Care and Malaria in Mali

The public health system in Mali is organized around decentralized community-based primary health care funded by user fees, as laid out in the Bamako Initiative of 1987. At the foundation of this system are *Centres de Santé Communautaire* (CSComs) – community-based primary care clinics managed by a local health association. The local health association retains revenues from sales of medications and other user fees, which are used to fund the operations of the CSCom. Similar community-funded public healthcare is found all over West Africa. This structure – where medical staff are salaried but have an implicit incentive to generate revenue to sustain the clinic – exists in many other settings across the developing world, including private, mission, and NGO-run medical facilities.

The quality of care across CSComs varies and is typically poorest in rural areas. Our study was conducted in the capital city of Bamako and a nearby suburb, where clinics offer a higher standard of care and are usually headed by a trained medical doctor (as compared to a national average of roughly one third (PMI, 2016)). CSComs are one of the most important sources of care: according to the 2012-2013 Demographic and Health Survey, 47 percent of mothers in Bamako who sought care for a child under five with fever or cough took the child to a CSCom.

One of the most commonly-treated illnesses in CSComs is malaria. In spite of recent progress, malaria is still the leading cause of mortality in Mali, accounting for roughly 20 percent of all deaths and nearly a third of deaths among children under five (IHME, 2016). Although the parasite is endemic in all parts of the country except the sparsely populated northern desert, rates of transmission are substantially lower in urban areas. For example, in 2015 the estimated prevalence of the parasite in children under five in Bamako was six percent, as compared to 36 percent nationwide (PNLP et al., 2016). Malaria infections are classified as either “simple/uncomplicated” or “severe”. Simple malaria is not life threatening if treated promptly, and is characterized by non-specific, flu-like symptoms including fever, chills, and headache. If left untreated, simple malaria can progress to severe malaria. In this stage of the disease, patients often suffer convulsions and experience life-threatening com-

plications, including loss of consciousness/coma, respiratory failure, renal failure, and severe anemia (Trampuz et al., 2003). Patients with severe malaria require prompt, aggressive treatment to avoid death and should be hospitalized until their symptoms stabilize.

Mali’s national malaria policy requires that suspected malaria cases be confirmed via a positive microscopy or rapid diagnostic test (RDT) before dispensing treatment (Ministère de la Santé, 2013).⁸ RDTs are meant to be free in public health facilities (including CSComs) to ensure that cost is not a barrier to accurate diagnosis.⁹ Artemisinin combination therapies (ACTs) are recommended for simple malaria, while severe malaria cases should be treated with injectable artesunate followed by a dose of ACTs once the patient is stable.¹⁰

In public facilities, ACTs are meant to be free for children under five and subsidized for older individuals, but there are no subsidies for other components of severe malaria treatment (e.g. monotherapies, injection services) (PMI, 2016). In our study area, the ACT subsidy policy is only partially enforced, as evidenced by the fact that 37 percent of children under five receiving ACTs in our control group paid a positive price.

2.2 Status Quo Malaria Treatment in the Study Sample

We now turn to our own data to quantify overuse of malaria treatment at the 60 clinics in Bamako that participated in our study. Here we give a brief overview of the data; we provide additional detail on the sampling frame and data collection protocols in section 4. We stationed surveyors at each clinic for 6 days over a two week period. Enumerators administered a short “intake” interview to all consenting patients seeking care for an acute illness. After the patient’s consultation with the doctor was complete, the enumerator recorded details of all blood tests performed, medications prescribed, medications purchased, and fees paid. A randomly chosen subset of these patients were selected for a more detailed home-based follow up survey on the day after their clinic visit. As part of this visit, a trained enumerator performed an RDT on the patient to confirm his/her malaria status – these tests allow us to precisely identify over- and under-treatment of malaria.¹¹ In what follows, we summarize

⁸This recommendation was first made in the five-year strategic plan for 2007-2011. Previously, presumptive treatment of any fever cases was the main approach to malaria control (Koné et al., 2015).

⁹In our data, RDTs were free of charge 70 percent of the time, while microscopy tests were free less than 3 percent of the time.

¹⁰Artesunate is a derivative of artemisinin. Artemisinin-based antimalarials are the most effective treatments for malaria in Sub-Saharan Africa, where the emergence of drug-resistant parasites has rendered earlier generations of antimalarials ineffective. Quinine can also effectively treat both simple and severe malaria infections in this region, but the drug is less effective than artemisinin and has more side effects (Achan et al., 2011). Malian policy generally reserves quinine for pregnant mothers, though national policy allows for an initial dose of quinine to treat severe malaria if an injectable artemisinin-based therapy is not available.

¹¹These tests provide an accurate measure of a patient’s true malaria status even after taking treatment because RDTs detect antigens for malaria, which remain in the bloodstream after antimalarials have been

(over)treatment outcomes on “Control” days, where no vouchers were available. We also draw on interviews that we conducted with health workers after the end of the in-clinic observation period to characterize providers’ beliefs regarding malaria prevalence and patient preferences.¹²

Overall, our data shows that there is a lot of malaria treatment in general, and that a substantial share of this is overtreatment. *A priori*, we find patterns that could be consistent with either patient- or doctor-driven overtreatment.

Table 1 uses our Control group data to give an overview of how acute illnesses are treated in clinics in Bamako. Column 1 shows results for all acutely ill patients visiting the clinic. Treatment for malaria is very common, with 31 percent of patients prescribed treatment for simple malaria and 30 percent for severe malaria. Twenty-one and 26 percent of patients purchase medication for simple and severe malaria, respectively.¹³ Malaria treatment is expensive: while the average patient not treated for malaria reported spending CFA 2,793 (\$4.58 at the average exchange rate of CFA 610 per USD that prevailed during the study period), simple malaria patients paid CFA 4,475 (\$7.34) and severe malaria patients paid CFA 10,353 (\$16.97). To put this in perspective, the fee for severe malaria patients amounts to 47 percent of average per capita monthly income in our sample.

We quantify overtreatment among the subset of patients who took an RDT at the home follow up. Column 3 shows average outcomes for all patients who took a test, while columns 4 and 5 show means for patients who tested malaria positive and negative respectively. Overtreatment is widespread, with 58 percent of malaria negative patients being prescribed antimalarials, as compared to 76 percent of malaria positive patients. Severe treatment accounts for 78 percent of prescriptions for those who tested positive for malaria and 41 percent for those who tested negative, in spite of the fact that donors estimate that only 10 percent of malaria cases should be severe (PMI, 2015). Use of antibiotics is also widespread, with over 40 percent of malaria positive and negative patients purchasing the drugs. Only 42 percent of patients who received severe treatment were kept for observation or referred to a

taken and parasites have cleared. We used CareStart HRP2(Pf) tests, which detect an antigen that typically takes a couple of weeks to become undetectable in blood samples (Humar et al., 1997; Kyabayinze et al., 2008). This test only detects the *P. falciparum* malaria parasite, which accounts for 92 percent of all malaria infections in Mali (PMI, 2016). In WHO quality assurance testing, these RDTs correctly identified malaria infection in 91 percent of low parasite density blood samples and 100 percent of high parasite density samples. False positives were rare, occurring less than 1 percent of the time, and the invalid test rate was 0 percent (WHO, 2017).

¹²Appendix B gives additional detail on the health worker survey.

¹³We classify treatment as consistent with severe malaria if any of the following are observed: (1) the patient received injectable artemisinin monotherapy or quinine or (2) the patient received artemisinin monotherapy or quinine recorded in the survey as tablets *and* received an ACT. These cases are consistent with administering monotherapy/quinine rectally, which is in line with Malian guidelines for severe malaria treatment.

higher-level hospital, even though all true severe malaria cases would require such referrals. This misallocation provides leading evidence of a gatekeeping failure, especially given the very straightforward treatment guidelines for malaria (prescribe only for patients with a positive test). Yet only 56 percent of Control group patients prescribed an antimalarial took a test, even though all clinics in our study had malaria testing capability. As we discuss below, it is not *ex-ante* clear whether this poor adherence is driven by doctor- or patient-sided factors.

There are some indications that overprescription might be driven by doctors. The patterns of severe malaria treatment we observe match anecdotal reports that doctors in our sample treat positive malaria tests (especially RDTs) as a sign of severe malaria, while interpreting negative tests accompanied by malaria-like symptoms as a sign of simple malaria. This, coupled with the fact that microscopy tests for malaria are not very accurate in field settings (Wongsrichanalai et al., 2007), could explain some of the misallocation we observe. The revenue especially from severe treatment provides an additional motive if doctors are concerned with the clinic’s financial health.¹⁴

We also note that twenty-five percent of patients do not purchase their antimalarial prescription, with the prescription-purchase gap mostly driven by individuals declining to purchase available treatment (77 percent of cases) rather than stock outs. Consistent with this, our home survey found that just 6 percent of individuals who declined to purchase at the clinic obtained an antimalarial from another source after leaving the clinic. This could be a consequence of credit constraints – a reason to subsidize treatment in the first place – but may also be a sign of misalignment between doctor prescriptions and patient preferences.

On the other hand, there are also signs of patient-driven overprescription. The share of patients who are prescribed treatment exceeds doctors’ average estimate of true malaria prevalence of 48 percent by 13 percentage points (see Appendix Table B1). Fifty-seven percent of health workers report that they feel pressure from patients to prescribe unnecessary medication, with over half of these health workers mentioning antimalarials specifically. Patients clearly over-estimate their malaria risk: 55 percent of respondents name malaria when asked what they think their illness is before they see the doctor, and among those with a valid RDT test at home, this proportion is 61 percent, even though only 25 percent have a positive test.

¹⁴While we do not observe profits, higher revenue from an individual patient will indicate higher profits provided profit is positively correlated with revenue. This is likely to hold in our case, especially since the higher cost of severe treatment in part reflects service fees e.g. for injections: we estimate service fees by taking the difference between the total cost of care and the cost of medications, tests, and equipment. Our estimates suggest that service fees are 35-47 percent higher for severe (as compared to simple) malaria. Also note that there was very limited patient queuing in our setting, so in the short run when staff cannot be adjusted, the marginal cost of providing such services is close to zero.

3 Theory

In order to investigate the role of patients pressuring doctors to prescribe, and more generally to understand the drivers of overprescription, we construct a model that matches the features of the market we observe in the Control group, and then use it to predict differences between two experimental pricing conditions. In both conditions, we provide vouchers that make simple malaria treatment free of charge to the patient, while keeping the clinic’s financial incentives constant. In the first condition, which we refer to as “Patient Voucher” or “*PV*”, both the doctor (she) and the patient (he) know that simple malaria treatment is free to the patient. In the second condition (“Doctor Voucher” or “*DV*”), only the doctor knows about the discount going into the consultation, and she can decide at her discretion to reveal this information to the patient. In section 4, we explain in detail how we operationalized these pricing treatments by giving treatment vouchers either directly to patients or to doctors.

The remainder of this section gives detail on the model setup, and then outlines the key testable predictions that allow us to distinguish between patient- and doctor-driven demand. In the main text we focus on the intuition behind these predictions, leaving the formal analysis for Appendix A.

3.1 Model Setup

Malaria Risk. Patients and doctors hold beliefs about the patient’s malaria risk that are determined by three parameters. First, the patient exhibits observable symptoms, described by a vector γ . Second, both the patient and the doctor receive an unobservable signal, ϵ and η , respectively. As a result, patients believe they have malaria with probability $\pi(\gamma, \epsilon)$, while doctors believe the malaria probability is $\hat{\pi}(\gamma, \eta)$. Since doctors have medical expertise and access to diagnostic tests, we assume that the signal on the patient side is strictly less informative than the doctor signal; that is, ϵ is correlated with the true malaria status of the patient and therefore with η , but does not contain additional information for the doctor that can improve her diagnosis.¹⁵

In practice, doctors and patients may form separate beliefs over whether the patient has a malaria infection, and whether it is severe. We collapse these beliefs into a single index because it simplifies the model, and matches both patterns in the data and our anecdotal

¹⁵In our model, doctors cannot influence patient beliefs, π , and patients cannot influence doctor beliefs, $\hat{\pi}$. This setup differs from a standard informed expert problem, where the doctor’s actions (e.g. prescription decisions) can influence patient beliefs. As we discuss in Appendix A, we abstract away from this channel because it does not shed additional light on doctors’ decisions to conceal or reveal vouchers in *DV* (which is what we use to identify patient- versus doctor-driven demand), while adding significant complexity to the analysis.

understanding of prescription decisions (specifically, rates of severe malaria prescription increase with the number of simple malaria symptoms in our data, despite the fact that just 2 percent of patients exhibited severe malaria symptoms like reduced consciousness, confusion, coma, or convulsions).

Patient Preferences. Patients receive different expected utilities from taking either simple malaria treatment ($t = l$), severe malaria treatment ($t = v$), or no malaria treatment ($t = n$). We normalize the expected utility of no treatment to zero. The expected utility of $t \in \{l, v\}$ is the expected benefit minus the cost:

$$U_t(\pi, P_t) = \pi B_t - C_t(P_t),$$

where $B_v > B_l$ are the benefits a malaria-positive patient receives from treatment for severe malaria versus simple malaria, respectively. Denote by $P_v > P_l$ the price of severe versus simple treatment to the patient, with $P_l \in \{0, P\}$, depending on whether the patient has a voucher. The cost of treatment, $C_t(P_t)$, includes both monetary and non-monetary costs like side effects. We assume $C_l(0) < C_l(P) < C_v(P_v)$.

Figure 1, Panel A illustrates patient preferences with and without the simple malaria discount. First consider preferences when $P_l = P$. Assuming that $\frac{C_l(P)}{B_l} < \frac{C_v}{B_v}$, this creates four ranges for the malaria probability π : individuals with the lowest levels of π are in N_P and prefer no treatment; as π increases, we progress to region $SL1_P$, where patients prefer simple treatment and would not purchase severe; next, in $SL2_P$, patients prefer simple but would purchase severe; and at the highest levels of π , patients are in SV_P , where they prefer severe treatment over simple or no treatment. This captures the idea that the benefits of both simple and severe treatment – and thus the patient’s willingness to buy them – are increasing in the number and severity of symptoms, but that the benefits of severe treatment only outweigh the higher costs when symptoms are most pronounced.

Figure 1, Panel A also shows that when the voucher lowers the price P_l to 0, the regions $SL1$ and $SL2$ expand, leading fewer patients to prefer both no and severe treatment. Thus, the voucher *increases* patient demand for treatment on the extensive margin, while *decreasing* demand for (more powerful) treatment on the intensive margin.

Doctor Utility. Doctor utility is a weighted average of what the doctor considers medical best practice, the patient’s utility, and clinic profits. Formally, we write this as follows:

$$V_t(\hat{\pi}, P_t) = (1 - \alpha) \left[\hat{\pi} B_t - \hat{C}_t \right] + \alpha \left[\hat{\pi} B_t - C_t(P_t) \right] + \beta \Pi_t$$

where $\hat{\pi}B_t - \hat{C}_t$ is the expected net benefit of treatment from the doctor’s perspective, $\hat{\pi}B_t - C_t(P_t)$ is expected patient utility given the doctor’s belief, $\hat{\pi}$, and Π_t is clinic profit from prescribing treatment $t \in \{l, v\}$. Note that Π_l does not depend on P_l – this is because we specifically designed the voucher treatments to leave financial incentives to sell unchanged. The altruism parameter α governs how much weight the doctor puts on patient utility relative to her own assessment of treatment benefit. Simple and severe treatment incur costs $\hat{C}_l < \hat{C}_v$ respectively from the doctor’s perspective.¹⁶ We assume $\frac{\hat{C}_l}{B_l} < \frac{\hat{C}_v}{B_v}$, so that doctors prefer no treatment for lowest malaria beliefs $\hat{\pi}$ (region \hat{N}), simple treatment for intermediate malaria beliefs (region $\hat{S}L$), and severe treatment for highest $\hat{\pi}$ (region $\hat{S}V$). As with patient preferences, we normalize the utility of prescribing nothing (V_n) to be zero.

Note that the doctor’s sympathies for her patients’ beliefs may create a shared preference for overtreatment. Alternatively, doctors’ revenue interest and incomplete altruism towards the patient’s expenses may lead to an average incentive to overtreat relative to the patient’s preferences.¹⁷ But it is also possible for doctors to prefer less treatment than patients, given that patients do not take the negative externalities of overtreatment into account and may be positively selected for high treatment utility. The latter two cases of disagreement are our object of interest. In what follows, we say patients drive demand for treatment when the doctor prescribes (and patients purchase) more aggressive treatment than the doctor herself thinks is optimal. We say the doctor drives demand when the doctor leads patients to purchase more (powerful) treatment than patients find optimal.

Gatekeeping Costs. If the doctor only cares about V , patients have no way to “induce” overtreatment; the doctor simply will not prescribe what she does not want the patient to purchase.¹⁸ To allow for a wedge between what the doctor prescribes and what she truly prefers, we assume that the doctor incurs a gatekeeping cost if she fails to prescribe the patient’s most preferred treatment. This cost encompasses factors like the risk of losing the patient to another clinic, any direct cost of dealing with an upset or angry person in the office, and the effort cost of explaining the prescription and cajoling the patient to comply.

¹⁶These costs may be lower than those of the patients, since treatment decisions based on medical considerations do not take into account what the patient has to pay. However, they could also be higher, because they encompass the cost of provision, side effects, and public costs like disease resistance concerns.

¹⁷In the interest of simplicity, we abstract away from the possibility that doctors overtreat to avoid additional diagnostic effort. This could be incorporated into the model by giving the doctor the option to obtain a more precise signal that allows her to distinguish which patients not to treat, versus simply treating all patients. Doing so does not change the main insights of our analysis.

¹⁸Note that our definition of patient-driven overtreatment does not include overtreatment that arises due to doctor altruism, since this is incorporated in V through weight α on patient preferences. If the two sides are aligned in their preference for overtreatment, including through altruism, neither side can be said to be “driving” overtreatment rates.

We assume that the doctor experiences a share g of the patient’s lost utility (relative to his first-best option) as a gatekeeping cost. Note that when a patient prefers no antimalarial, he can always implement this first best even when the doctor writes a prescription, as we assume he can decline to purchase a prescription at no cost. This form of overprescription therefore incurs no gatekeeping cost.¹⁹

The range of π that incurs a gatekeeping cost from prescribing no or severe treatment increases when patients know simple treatment is free. This reflects the fact that patient demand for simple treatment is higher when the price is lower. As a result, the doctor has a strategic advantage in the Doctor Voucher condition (relative to the Patient Voucher condition), because she can selectively conceal the voucher to manipulate patient demand in a way that reduces her gatekeeping costs.

Doctor-Patient Communication. In the interest of realism, we assume that there are limits to how much doctors and patients can communicate. Specifically, doctors cannot learn ϵ , and patients cannot learn η , though the joint distribution of the signals is known. This means the doctor cannot be completely sure what gatekeeping cost she will incur when prescribing treatment t – rather, she must base her decision on expected gatekeeping costs given γ and η .²⁰ This allows for the possibility that some patients do not purchase all the medications they are prescribed, which matches the patterns in our data. In the face of this uncertainty, we permit doctors to adopt a strategy that avoids gatekeeping costs completely, by letting the patient *choose* between simple and severe treatment. In this case purchasing patterns are completely determined by patient preferences.²¹

3.2 Comparing Doctor Voucher and Patient Voucher Treatments

Our first testable prediction focuses on detecting gatekeeping costs when doctors and patients have different prescription preferences:

Prediction (1) If there are gatekeeping costs and doctors and patients do not always agree on the optimal prescription, the use of vouchers will be higher in PV than DV.

¹⁹There may be a similar non-internalized cost from prescribing the patient more than he actually buys, akin to a negative reputation effect from prescribing a “clearly unnecessary” treatment; e.g. if the patient judges the doctor to be greedy or incompetent and decides not to return. However, this cost would only be incurred when the patient chooses not to buy the prescribed treatment, and its negative effect on the doctor’s utility cannot be distinguished from the direct effect of the patient not getting treated.

²⁰This structure is more general than assuming doctors can perfectly observe ϵ . This is because our setup nests the perfect information case, which is covered by a fixed ϵ with no variance.

²¹In an earlier version of this paper, we discuss the model without this option (available on request) – the key insights are unchanged.

Obtaining this result is straightforward: absent gatekeeping costs, the doctor will always choose the action that maximizes her utility, V . This action is independent of voucher condition, which in turn implies that voucher use will be the same in PV and DV . When gatekeeping costs are present, note two things: First, if it is optimal to use the voucher under DV , it is also optimal to use the voucher under PV . Second, under DV the doctor can strategically lower gatekeeping costs by concealing the voucher from patients to whom she does not want to prescribe simple treatment. In PV , by contrast, the gatekeeping cost she incurs when refusing to prescribe simple treatment may lead her to switch her prescription to simple, despite her own preferences.

We motivate the remaining predictions with two stylized examples, illustrated in Panels B and C of Figure 1. For simplicity, the examples consider a case where (a) gatekeeping costs are so high that the doctor is never willing to override the patient’s preferences, and (b) the doctor is able to perfectly observe patient beliefs π . Our predictions hold more generally; we formally establish these results in Appendix A, where we analyze the model in detail.

Prediction (2) If the share of patients who are prescribed and purchase any kind of malaria treatment is lower in DV than PV, this indicates patient-driven demand.

This prediction maps to Figure 1, Panel B, where we consider the optimal strategy of a doctor who would like to sell as little medication as possible (i.e. nothing \succ simple \succ severe) if it were not for gatekeeping costs. This is an extreme case of *patient-driven demand*. The uppermost bar at the bottom of the figure shows patient demand when $P_l = P$ (no voucher), while the middle bar shows patient demand when $P_l = 0$. Since gatekeeping costs are exceedingly high, this bar also reflects prescription and purchase outcomes under PV . The lowermost bar shows outcomes per the doctor’s optimal strategy under DV . Here, to minimize demand, the doctor conceals the voucher from *lower risk patients* with $\pi \in N_P$ (at price $P_l = P$ these patients prefer to buy nothing, so no gatekeeping cost is incurred), but shares the voucher and prescribes simple malaria treatment to higher risk patients with $\pi \in SL1_P \cup SL2_0$. This strategy lets the doctor sell less malaria treatment under DV than under PV , in line with prediction (2). In the appendix, we formally show that any substitution from no malaria treatment in DV to simple malaria treatment in PV will be driven by doctors who prefer not to prescribe, but give in to patient pressure when the discount is known.

Prediction (3) If the share of patients who are prescribed and purchase severe treatment is lower in PV than DV, this indicates doctor-driven demand for more (powerful) treatment.

Figure 1, Panel C, illustrates the optimal strategy of a doctor who would like to sell as much medication as possible (i.e. severe \succ simple \succ nothing). This is our example of *doctor-driven demand*. In this case, the doctor will share the voucher with lower risk patients with $\pi \in SL1_0$, since this maximizes their demand for antimalarials. Thus, in contrast to Panel B, the share of patients receiving any antimalarial does not differ between *PV* and *DV*. However, in *DV* the doctor will conceal the voucher from *higher risk* patients on the margin between preferring simple and severe treatment ($\pi \in SL2_0 \cap SV_P$), thereby increasing the share of patients purchasing severe treatment relative to *PV*. The appendix establishes the general case: that any substitution away from simple to severe in *DV* will be driven by doctors who prefer severe, selling to patients who prefer simple.

Our final prediction summarizes observations made above about the “marginal” patients receiving vouchers under *PV* but not *DV*:

Prediction (4) Higher voucher use and substitution into prescription and purchase of simple treatment in PV are driven by lower malaria probabilities if patients drive demand for treatment, but by higher malaria probabilities if doctors drive demand for treatment.

All proofs are in Appendix A, where we also discuss additional results on prescribing simple treatment without using the voucher, and on doctors who are purely revenue motivated. We now turn to the experiment we designed to identify the role of patient and doctor preferences in driving demand for malaria treatment.

4 Experimental Design and Sample

4.1 Experimental Design

Sampling Frame. At the outset of our study, we obtained an administrative list of all CSCComs in the city of Bamako and in nearby Kati and Kalaban Coro in Koulikoro. After conducting a census of these clinics, we dropped clinics that had closed or were more than 15 kilometers away from Bamako, and removed 21 clinics that were working with a local NGO to offer subsidized, improved care to mothers and children. This yielded a final sample of 60 clinics. Four health workers at each clinic enrolled in the study were invited to attend a refresher training that covered Mali’s official diagnostic and treatment guidelines for malaria and a hands-on training on how to administer an RDT. The purpose of the training was to ensure that all providers were equally aware that only patients with a positive malaria test should receive treatment per the national malaria control policy, and to enable caregivers to

follow this policy by teaching them how to conduct the test.²² The training materials were prepared by the research team in cooperation with researchers at the University of Bamako, as well as the *Programme National de Lutte contre le Paludisme* (PNLP), Mali’s department of malaria control, and the Malaria Research and Training Center Bamako. All trainings were conducted by five trainers from the PNL P and one trainer from the regional health directorate (Direction Régionale de la Santé, or DRS). All clinics sent at least two and on average 3.9 providers to the training.²³

Doctor and Patient Voucher Treatments. The simple malaria treatment discounts were delivered via vouchers for a free course of ACTs, and randomized within each clinic across six different days during a two-week observation period. The objective of the vouchers was to reduce the cost of treatment for simple malaria, while leaving both the revenue to the clinic and the cost of other types of treatment to the patient fixed. The vouchers were delivered by trained “intervention officers” who arrived at the clinic in the morning. Intervention officers did not perform any survey activities, were stationed at a separate part of the clinic from the survey team, and intervention eligibility was not tied to survey participation.²⁴

In the Patient Voucher condition, vouchers were distributed directly to patients when they arrived at the clinic, before they went to consult the doctor. Patients and/or caregivers were informed that the voucher would pay for simple malaria treatment (ACT tablets), provided the doctor determined that this was the appropriate course of action. Study officers were carefully trained to limit potential endorsement effects: patients were informed that the decision of whether to use the voucher was up to the doctor, and that the patient could not and should not use the voucher if they were not diagnosed with malaria. In the Doctor Voucher condition, patients were not informed about voucher availability before the consultation. The vouchers were directly given to the doctors, who could assign the vouchers to patients at their discretion.

Intervention officers always brought more than enough vouchers to cover a day’s demand,

²²In practice, many doctors failed to adhere to these guidelines in spite of the training: on Control days, just 56 percent of patients prescribed an antimalarial received a malaria test first. This ameliorates concerns that the training made it difficult to detect doctor-driven demand, e.g. due to Hawthorne effects.

²³Half the clinics were randomly selected to receive a more intensive training designed to increase doctors’ trust in the diagnostic accuracy of RDTs. Appendix B has more detail on this treatment, which is analyzed in detail in a separate paper. This treatment has no impact on our analysis of simple malaria discounts. The training invitation letters did not specify which training the clinic was selected for – as a result, average attendance rates were identical across the intensive and basic training groups. Clinics were trained in six groups in November 2016.

²⁴Each intervention officer delivered the same treatment to all clinics and thus rotated between clinics, while the separate survey team stayed at the same clinic across the entire sampling period.

so voucher rationing was not an issue. Both conditions used the same voucher design, which required the doctor’s signature to verify the simple malaria diagnosis. The voucher was not valid unless used on the same day. Signed vouchers were processed at the clinic pharmacy after the consultation was complete, and a copy had to be returned to the intervention officer by the patient, who verified that the correct medication and full discount was received. At the end of the day, unused vouchers were collected (in *DV*), and the clinic was reimbursed for all vouchers redeemed that day at their standard sales price.

The within-clinic randomization and associated data collection were conducted after the doctor trainings in November and December 2016, covering the end of the rainy season and therefore the period of highest malaria risk.²⁵ We divided the 60 clinics into three cohorts of twenty based on geography. Each of the three cohorts rotated through two weeks of data collection and experimental intervention. Within each cohort, we randomly assigned each clinic to one of 20 intervention schedules depicted in Figure 2. Each clinic received two Control days, two Doctor Voucher days, and two Patient Voucher days. Although all clinics were informed of the upcoming study activities and interventions in advance, clinic staff did not receive prior notice of the actual intervention schedule – rather, our field officers informed clinic staff of the day’s intervention on the morning of an observation day.

4.2 Data Collection

In order to identify the source of over-demand for malaria treatment, we require detailed data on prescription and treatment outcomes from patients at the clinic, as well as patients’ true malaria status. Our main analysis uses data from three different sources.

In-Clinic Survey. As illustrated by Figure 2, we collected data at the clinic on six observation days spread over a two week period. On these observation days, enumerators approached all acutely ill patients at clinic intake and requested consent to participate in the study. We classified a patient as “acutely ill” if they were feeling sick and exhibited any of the following symptoms: fever, chills, excessive sweating, nausea, vomiting, diarrhea, poor appetite, headache, cough, weakness, fatigue, or reduced consciousness. Before the doctors’ consultation, enumerators recorded patients’ basic demographic details, symptoms, and information on any prior treatment and/or diagnosis.²⁶ After the consultation was complete,

²⁵The Doctor and Patient Voucher treatments were cross-cut with a third within-clinic treatment designed to give patients better information about the accuracy and availability of RDTs (“Patient Information” or “*PI*”). This treatment is analyzed in a separate paper and is not the focus of the present analysis. Appendix B describes this treatment in more detail.

²⁶We interviewed the patient whenever s/he was aged 15 or older and well enough to take the interview; otherwise a guardian answered the survey.

the enumerator recorded details of all prescribed and purchased medications, blood tests taken, and costs paid to the clinic. All told, we completed 2,055 interviews.

Home Follow-Up Survey. We randomly selected 1,735 patients for a more detailed follow-up survey conducted in the patient’s home the day after the clinic visit, of whom 1,495 (86 percent) were successfully interviewed.²⁷ The home survey collected information on changes in the illness, any treatments or tests obtained after the clinic visit, and recorded whether patients were taking the medications that were purchased at the clinic. During this survey 1,126 patients (75 percent of interviewed patients/65 percent of patients targeted for the home survey) gave separate consent to take an RDT and obtained a valid test result.²⁸ Appendix Table B2 shows that there are no significant differences in the probabilities of being selected for home survey, taking the home survey, or having a valid home-based RDT by treatment arm.

Administrative Voucher Data. Finally, we use our own administrative data on voucher redemption to quantify how informing patients about vouchers impacts redemption rates.

4.3 Background Characteristics and Randomization Verification

Table 2 presents summary statistics on the study population and the results of balance tests. The first column of the table displays means and standard deviations of baseline characteristics in the Control group (patients who visited the clinic on days where no vouchers were available). Columns 2 and 3 present coefficients from the following regression specification, which we use throughout the paper:

$$y_{ict} = \beta_0 + \beta_P PV_{ct} + \beta_D DV_{ct} + \gamma_t + \varepsilon_{ict} \quad (1)$$

where y_{ict} is the outcome of interest for patient i visiting clinic c on day t , PV_{ct} and DV_{ct} are dummy variables indicating the patient voucher and doctor voucher interventions for clinic c on day t , and γ_t are date fixed effects.²⁹ The treatment rotation schedules were assigned at the clinic level, so we cluster standard errors at this level as well. Since the relative difference between giving vouchers to doctors versus patients will be of special focus in the analysis,

²⁷From November 14-17 2016 we randomly selected up to six patients per clinic per study day to interview at home. Between November 18 and December 30 2016 (the end of the study) we randomly selected up to six patients per clinic per study day, and then oversampled at clinics who saw more than six patients, until 60 patients per study day were sampled.

²⁸1,158 patients gave consent; 32 tests were inconclusive.

²⁹We include date fixed effects because the proportion of PV , DV , and C days in the sample varies depending on the date.

column 4 presents the p-value from an F-test of the null hypothesis that $\beta_P = \beta_D$. Column 5 presents the p-value from an F-test of the null hypothesis that the vouchers treatment effects are jointly equal to zero ($\beta_P = \beta_D = 0$). Each row shows results from a separate regression.

The first row of Table 2 shows that just over five acutely ill patients with eligible symptoms visited the clinic on the average observation day. Note that the voucher treatments could affect outcomes not just through a direct effect on prescription and purchase behavior, but also through a selection effect, whereby the pool or number of patients visiting the clinic changes in response to the experimental conditions. We randomized the voucher treatments within-clinic and administered them only for a day at a time and at unannounced dates, in order to minimize this selection effect. Nonetheless there is a 0.63-0.76 patient increase in traffic on Doctor and Patient Voucher days (significant at the 5 percent level in both cases). Critically, there is no evidence of a differential increase on PV versus DV days, which is most important for our testable predictions.

Panels B and C show that the average characteristics of patients (and their households) are similar across the different voucher treatment groups. Forty percent of patients are male. The average patient is 17 years old, reports 3.6 symptoms and has been sick for 4.3 days. The most commonly reported symptoms are fever (80 percent of patients) and headache (62 percent of patients). Thirty percent of patients report chills or excessive sweating. These symptoms are all commonly associated with malaria (Taylor et al., 2010), but only 25 percent of patients taking a home RDT tested positive. Note that home-tested patients in the Patient Voucher group are less likely and those in the Doctor Voucher treatment to have a positive RDT. While each of these is not significantly different from the Control, the PV-DV difference of six percentage points is significant at the 5 percent level. Malaria-positive patients are more likely to get treated for malaria and especially severe malaria, so this could lead to a difference in treatment rates in *DV* and *PV*. Note, however, that greater rates of simple and severe treatment in *DV* would bias results *against* patient-driven and *towards* doctor-driven demand. This is the opposite of what we find.

Panel C of Table 2 shows that patients come from relatively poor families: just under one third of primary “respondents” (either the patient or the caregiver accompanying the patient to the clinic, in the case of minor children) are literate, households are quite large, with 10.3 members on average, and monthly per capita income is just under CFA 22,000 (\$36). Forty-one percent of household members are children under the age of 15, and households own 0.49 mosquito nets per capita.

Appendix Table B3 further shows that the characteristics of individuals selected for the home survey versus not, and interviewed versus not (conditional on selection), are balanced. However, there is strong evidence of selection into taking the home RDT: patients who

consented to take the home test are significantly sicker and more likely to have purchased an antimalarial than patients who did not give consent.

4.4 Predicting Malaria Risk

We would like to study how prescription and treatment outcomes vary with patients’ underlying malaria risk (π in the model). We therefore follow Cohen et al. (2015) and estimate the relationship between RDT result and observed patient characteristics and symptoms in the clinic survey with the following probit specification:

$$E [pos_{ict} | x_{ict}] = \Phi \left(x'_{ict} \lambda \right) \quad (2)$$

where pos_{ict} is a dummy variable equal to one if a patient tests RDT positive and x_{ict} is a vector including dummy variables for each of the seven screening symptoms (fever, chills/excessive sweating, nausea/vomiting/diarrhea, poor appetite, headache, cough, weakness/fatigue), days since onset of illness, patient age, a dummy equal to one if the patient is under age 5, the interaction between age and the under 5 dummy, patient gender, and patient pregnancy status. We also control for basic demographic characteristics indicative of socio-economic status (which correlates with malaria risk), including the survey respondent’s ethnicity, ability to speak French, literacy in French, and education.³⁰ The results of this regression, shown in Appendix Table B4, are used to impute a malaria risk for each patient. Appendix Figure B1 shows that the distribution of predicted malaria risk by treatment group is very similar, and Table 2 also finds no significant difference in average predicted risk across treatment groups.

The probit is estimated off the selected sample of individuals who consented to a home-based test. If this selection is partly based on unobservables (to the econometrician) that are informative about malaria status, the malaria risk of tested patients may be different, conditional on recorded symptoms, than that of the average patient. However, our most important tests do not require estimated risk to be unbiased; rather we require a proxy that is strongly correlated with true risk. The regression assigns positive weight to symptoms known in the medical literature to indicate malaria, like fever, chills, nausea, and headache (CDC, 2015), which suggests that imputed malaria risk is correlated with true malaria risk. Consistent with this, predicted risk is highly correlated with malaria treatment outcomes. Figure 3 uses data from the Control group to plot local linear regressions, where predicted

³⁰We expect malaria risk to change discontinuously at age 5 because the Malian government had a policy of seasonal malaria chemoprevention for children under 5 at the time of the study (PMI, 2017). We control for the patient and respondent being different people as well. Results are similar if we omit the basic demographic characteristics from the probit.

malaria risk is the running variable. Panel A shows that the shares of patients who are not prescribed an antimalarial (black line) and who did not purchase one (grey line) steadily decline with predicted risk. Panels B and C repeat this exercise for simple and severe malaria prescriptions and purchases. They show that purchases and prescriptions of simple treatment decrease and those of severe treatment increase with predicted risk. As a proportion of prescriptions, the share of patients who decline the prescribed treatment is decreasing in predicted malaria risk for both simple and severe treatment. Appendix Figure B2 shows that these patterns hold for those with valid home-based RDT tests and those without.

5 Main Results

5.1 Empirical Approach

We are now prepared to assess the extent to which overtreatment is driven by patients versus doctors. We first recap our four main testable predictions and map them into null hypotheses that we can impose on the data. We use equation 1 as our core specification for estimating treatment effects. To improve precision and address any concerns of patient selection on PV and DV days, we augment equation 1 to control for additional clinic and patient-level characteristics. To avoid specification searching, we select these additional covariates using double lasso (Belloni et al., 2014).³¹ Appendix Tables B5 and B6 show that results with no additional controls are very similar.

We also study heterogeneity in treatment effects with respect to above/below median predicted malaria risk using the following specification:

$$y_{ict} = \delta_0 + \delta_{PH}PV \times high_{ict} + \delta_{DH}DV \times high_{ict} + \delta_{PL}PV \times low_{ict} + \delta_{DL}DV \times low_{ict} + \theta high_{ict} + x'_{ict}\alpha + \lambda_t + \nu_{ict}$$

where $high_{ict}$ and low_{ict} are dummy variables identifying patients with above/below median predicted malaria risk, x_{ict} is the vector of clinic and individual-level controls (again selected using double lasso), and λ_t are date fixed effects. Since $high_{ict}$ and low_{ict} are generated

³¹Candidate controls include a set of clinic fixed effects, as well as all individual characteristics included in Table 2, except those only measured at the home survey: symptom dummies, illness duration (top-coded at the 99th percentiles), patient age, an under 5 dummy, patient gender, a dummy to identify pregnant women, a dummy to identify patients who were also survey respondents, survey respondent gender, respondent ethnicity, education, literacy and knowledge of French, and a dummy variable identifying patients interviewed in the home follow-up survey. Missing values are dummied out and recoded to the sample mean. We also include squared terms for patient age and illness duration, as well as pairwise interactions between the aforementioned individual-level controls. We partial out date fixed effects, so they are included in all regressions.

regressors, we present bootstrapped standard errors clustered at the clinic level for this specification.³²

Note that our theoretical predictions specify one-sided alternatives. For instance, prediction (1) is that if gatekeeping costs are non-zero, then use of vouchers will be higher in PV as compared to DV. The highest-powered way to test this prediction is to test the null of $H_0 : \beta_P = \beta_D$ against the alternative that $H_A : \beta_P > \beta_D$.

We showed that gatekeeping costs can increase overtreatment if patients demand more malaria treatment than doctors would like to supply, or act as a moderating force on overtreatment if doctors are the source of overprescription. The rest of the testable predictions are designed to parse between these two alternatives. Specifically, if overtreatment is driven by patients:

- The share of patients prescribed and purchasing any malaria treatment will be higher in *PV* than *DV* (prediction (2)), mapping to an alternative hypothesis of $H_A : \beta_P > \beta_D$.
- *DV-PV* differences in voucher use and substitution into malaria treatment will be driven by individuals with *low* malaria risk (prediction (4)). The relevant alternative hypothesis is $H_A : \delta_{PL} > \delta_{DL}$.

If overtreatment is driven by doctors:

- The share of patients who are prescribed and purchase severe malaria treatment will be lower in *PV* than *DV* (prediction (3)). This maps into an alternative hypothesis of $H_A : \beta_P < \beta_D$.
- *DV-PV* differences in voucher use and substitution to simple malaria prescription and treatment will be driven by individuals with *high* malaria risk (prediction (4)). The relevant alternative hypothesis is $H_A : \delta_{PH} > \delta_{DH}$ for voucher use and $H_A : \delta_{PH} < \delta_{DH}$ for severe malaria prescriptions and purchases.

We now turn to the data to assess these predictions empirically.

5.2 Impacts on Antimalarial Prescriptions and Purchases

Overall Impacts. Table 3 begins by studying treatment effects on voucher use, as well as effects on broader malaria prescriptions and purchases. We report the coefficients on the

³²Specifically, first we use double lasso to select covariates. Then we create 1,000 bootstrapped samples by sampling 60 clinics with replacement. Then we re-estimate predicted malaria risk in each bootstrap sample, reconstruct *high_{ict}* and *low_{ict}* based on the new risk estimate, and run the regression. Standard errors are based on the distribution of coefficients across bootstrapped samples.

treatment dummies β_P and β_D and the p-value of both a two-sided test and the appropriate one-sided test (see above). We also indicate the theory prediction that corresponds with the alternative hypothesis (existence of gatekeeping costs, doctor-driven, or patient-driven overprescription – GC, DD, or PD).

The first column looks at *voucher redemptions*. We find strong evidence in favor of gatekeeping costs: voucher redemption is 9 percentage points (35 percent) higher in *PV* as compared to *DV*. This difference is statistically significant at the one percent level with the appropriate one-sided test.

The next two columns study the share of patients who were either prescribed (column 2) or purchased (column 3) *any malaria treatment*. The results are consistent with patient-driven demand: more patients are prescribed and purchase malaria treatment in *PV*. This means the subsidy led to stronger substitution out of no treatment when patients were aware of the ACT discount. Individuals who visited the clinic on *PV* days (relative to Control days) were 5.2 percentage points more likely to leave with a malaria prescription and 14 percentage points more likely to purchase an antimalarial treatment. By contrast, *DV* increased purchases – by a highly significant 8.5 percentage points – but had no impact on prescriptions. Thus, doctors used vouchers for patients to whom they were giving prescriptions already, but not to make new prescriptions. The *PV* versus *DV* differences are significant using both the one-sided and two-sided tests.

The last two columns turn to severe malaria treatment to test for evidence of doctor-driven demand. In spite of the strikingly high rates of severe malaria treatment in our sample, we find no evidence that this is driven by doctor preferences. In fact, the substitution from severe to simple treatment is, if anything, stronger in *DV* than in *PV* (see column 5). This is the opposite of our theoretical prediction under doctor-driven demand.

Impacts by Predicted Malaria Risk. If patients are pressuring doctors to prescribe, the differential treatment effects on voucher use should be driven by a stronger substitution from no treatment to simple treatment in *PV* at the lower margin of the predicted risk distribution. By contrast, doctor-driven demand should lead to a stronger substitution from severe to simple treatment in *PV* when malaria risk is higher, because doctors in this condition cannot conceal the vouchers to maintain demand for severe treatment. Figure 4 plots the relationship between treatment outcomes and predicted malaria risk by voucher condition using local linear regressions.³³ For reference, vertical lines demarcate the 25th, 50th, and 75th percentiles of the predicted risk distribution.

³³To account for clustering and the fact that predicted risk is a generated regressor, confidence intervals are bootstrapped using the procedure described in sub-section 5.1.

Again, our predictions for patient-driven demand are borne out in practice. Panel A shows that voucher use is persistently higher in *PV* up to the 75th percentile of predicted malaria risk. Panel B shows the share of patients purchasing any malaria treatment is higher in *PV* up until the 75th percentile. Panel C shows that the Doctor Voucher group has *lower* purchase rates for severe malaria at higher levels of predicted malaria risk. Note again that this cuts *against* the prediction for doctor-driven demand.

To better quantify these patterns, Table 4 looks at heterogeneity in treatment effects by above versus below median predicted malaria risk. While rates of voucher use in the *DV* and *PV* treatments are very similar for above-median (“high”) risk patients, voucher use by low risk patients is 15 percentage points higher when patients know about the vouchers before consulting the doctor. This represents a striking 71 percent increase in voucher use relative to low risk patients in *DV*. We see similar patterns for any antimalarial prescription and purchase, where we are able to reject the null in favor of the one-sided and two-sided alternative at the five percent level or better for low risk patients. *PV* increases low-risk prescriptions by 8.1 percentage points (16 percent) and purchases by 9 percentage points (23 percent) relative to *DV*. Finally, columns 4 and 5 explore effects for severe malaria treatment. While we formally reject $H_0 : \delta_{PH} = \delta_{DH}$ against the two sided alternative at the 1 percent level in column 5, we fail to reject the null in favor of the “doctor-driven demand” one-sided alternative, since prescriptions and purchases of severe malaria treatment at high predicted malaria risk are actually *lower* under *DV* as compared to *PV*.

Adherence to Treatment Guidelines. Why would the substitution from simple to severe malaria treatment be *lower* in *PV*? While this finding is seemingly inconsistent with the theory, Figure 5 offers an explanation: doctors were more likely to bend the vouchers’ rules when patients knew about discounts. Specifically, Panel A graphs the share of patients who purchased severe malaria treatment by treatment arm. In the control group, none of the 26 percent of patients purchasing severe malaria treatment used a free ACT voucher, which is to be expected. However, none of the severe malaria patients in the Doctor and Patient Voucher treatments should have used a voucher either, since the vouchers were only valid for simple malaria prescriptions. This condition, however, relied on doctor certification as well as patients reporting their prescription correctly to the intervention officer, so could in practice be violated.

The third bar of Panel A shows that doctors and patients almost never broke the rules in *DV*: just one percent of patients in this condition purchased severe malaria treatment and used a voucher. In contrast, over four percent of patients in *PV* purchased severe malaria treatment and used a voucher. The graph makes clear that this differential rate of rule

breaking entirely accounts for the differential rate of severe malaria purchases between the two voucher conditions. Columns 1 and 2 in Appendix Table B7 verify that the rate of rule breaking is significantly higher in *PV* versus *DV* (we reject the two-sided test at the one percent level). This behavior is also consistent with patient-driven demand: doctors are sometimes willing to violate the terms of the voucher when patients push for the discounted treatment.

Not Using the Voucher for Simple Treatment. Columns 3 and 4 explore the inverse of this phenomenon: patients purchasing simple malaria treatment without a voucher in *DV* versus *PV*. It is possible that vouchers were sometimes not used because antimalarials other than ACTs were given, or because patients visited the clinic before the intervention officer arrived. Alternatively, some doctors or patients may have simply forgotten to apply the voucher. However, our model predicts that doctors may intentionally prescribe simple treatment without a voucher when their preferences disagree with those of the patient, and that they do so more often in *DV* (see prediction (5) in Appendix A). Intuitively, a doctor who prefers to prescribe something different from simple treatment, but suspects that her patient has a strong preference for it, may prefer to prescribe simple malaria treatment without the voucher. In this manner, she avoids the gatekeeping costs of refusing the patient’s preferred option, while ensuring patients with weaker preferences do not purchase simple treatment.³⁴ The assumption that the doctor cannot observe patient beliefs, π , is critical for obtaining this result – otherwise there would be no reason to hedge. Consistent with prediction 5, Appendix Table B7 shows that patients in *DV* were three percentage points more likely than patients in *PV* to purchase simple malaria treatment without a voucher (statistically significant at the 10 percent level using the two-sided test).

Alternative Explanations. An important alternative hypothesis for the *PV-DV* differences we see is limited attention: it could be that doctors in *DV* simply forgot about the vouchers in some cases, whereas in *PV* they were reminded by patients. Such a mechanism could produce the *DV* versus *PV* reduction in voucher use as well as the lower substitution into treatment that we observe in Table 3. However, this mechanism cannot explain the differences we observe in severe malaria prescription behavior. Specifically, with limited attention we would expect to see *less* substitution and *more* use of severe malaria treatment in *DV* than in *PV*.

Another possibility is that knowledge of the voucher led patients to exaggerate their symptoms, either intentionally or subconsciously (e.g. if the vouchers made the threat of

³⁴They may buy no treatment, or, if the doctor offers a menu of simple and severe, severe treatment.

malaria more salient), in order to affect the doctor’s assessment of their malaria risk. If doctors used patient-reported symptoms for diagnostics without accounting for this change, this would amount to a form of patient-driven demand that operates through changing doctors’ direct utility, rather than their gatekeeping costs.³⁵ However, the fact that doctors almost never violated the rules of the vouchers in *DV*, but did in *PV*, strongly suggests that gatekeeping costs matter in our context.

A final concern is that our results are driven by differential selection into care seeking between *PV* and *DV* (e.g. more people visit the clinic on *PV* days in order to take advantage of the vouchers). However, Table 2 shows no *PV-DV* difference in patient volume per clinic day. Patient observables are also balanced across treatment arms.

Hawthorne and Experimenter Demand Effects. It is also important to ask how our results might be influenced by Hawthorne or experimenter demand effects. This could be especially relevant for doctor behavior; although all clinics were formally told that the study was not a performance review and that collected data would be kept strictly confidential, clinics were informed that the study was being conducted in collaboration with the National Malaria Control Program and the National Directorate of Health. Clinics also knew that the pre-intervention training was part of the study; this training emphasized the importance of adhering to official Malian policy when treating malaria, meaning that only patients with confirmed positive test results should receive an antimalarial prescription. We therefore expect that Hawthorne effects would increase the likelihood of doctors adhering to official treatment guidelines, which would in turn bias us away from finding *PV-DV* differences in treatment outcomes. However, given that only 56 percent of patients prescribed an anti-malarial on Control days received a malaria test (testing before treatment was a major focus of the training), we believe doctors were not very concerned about the implicit monitoring that took place during survey days.

A related worry is that *PV-DV* differences are not driven by inherent patient preferences, but rather patient preferences that have been “enhanced” by experimenter demand effects (e.g. if patients interpreted receiving the voucher as a signal that “buying an antimalarial is encouraged by an outside NGO”). While this would not invalidate our tests of gatekeeping costs and reluctant compliance by doctors, understanding the scope of experimenter demand effects is important for interpretation and assessing external validity. We believe experimenter demand effects were minimal in our setting for two reasons. First, study staff were carefully trained to explain that the vouchers were not valid without a doctor’s signa-

³⁵A related possibility would be if doctors perceived more patients in *PV* to be “likely compliers” with prescribed simple malaria treatment, if for example more patients explicitly requested simple malaria treatment.

ture, and that the doctor was the one who would make the malaria diagnosis. Second, data from a cross-cut patient information treatment (*PI*) suggests patients were not influenced by experimenter demand effects. On half of all clinic observation days, acutely ill patients were shown a short video, which emphasized the importance of confirming all suspected malaria cases with a malaria test before taking an antimalarial (see Appendix Figure 2 for the randomization schedule). The video also informed patients that RDT tests should be available for free at the clinic, and described how different types of malaria tests worked and how long they took.³⁶

The video clearly communicated that testing suspected malaria cases was the “right” thing to do. However, Appendix Table B8 shows that *PI* actually *decreased* the share of patients taking a malaria test, while leaving the share of patients unchanged who were prescribed an antimalarial. Panel B estimates heterogeneity in *PI* effects by voucher treatment – the negative effect of *PI* is evident on non-voucher days, which suggests that effects are not simply driven by patients refusing tests in the hope of redeeming a voucher. These results are inconsistent with experimenter demand effects changing patient preferences, but are consistent with the hypothesis that some patients come to the clinic with a strong preference for an antimalarial, and decline testing when they are informed of Mali’s official policy of withholding treatment from patients who test malaria negative. At the same time, we find that excess malaria prescriptions on Patient Voucher days are concentrated on days *without* the Patient Information video (Panel B, column 1). This suggests that the video did reduce physician gatekeeping costs for some patients.

5.3 What Can We Learn About Doctor Preferences?

The results so far make it clear that doctors do not simply internalize patient utility when making prescription decisions. But how, exactly, do doctor preferences diverge from patient preferences?

Relative to the Control group, we find patients in the *DV* condition are just as likely to receive a prescription for some form of antimalarial (Table 3), but doctors substitute prescriptions away from severe malaria and into simple malaria (Figure 5). In general, differences between Control and *DV* are hard to interpret, as they may be due to relaxing the constraint on sales imposed by patient preferences, or changed doctor preferences.

The exception is a situation where doctors are purely revenue motivated. As we show in Appendix A (prediction (6)), the vouchers will lead to an increase in clinic revenue in this case. Appendix prediction (7) says that revenue-maximizing doctors will use the informa-

³⁶Appendix B describes this treatment in more detail

tional advantage in the *DV* treatment to increase clinic revenues relative to *PV*.³⁷ Appendix Table B9 shows that there is no significant difference in (per-patient) clinic revenues across treatment arms, while the average patient in both *PV* and *DV* pays roughly CFA 500 *less* than in the Control group (significant at the 5 percent level for *DV* only).

All told, our results do not paint a picture of doctors primarily motivated by revenue, at least on the margin. Rather, doctors seem to have a preference for withholding medically unnecessary treatment, even if this does not further their financial interests, but only if they do not have to endure a large gatekeeping cost to do so.

5.4 Quantifying Contributions to Overtreatment

What are the implications of doctors’ gatekeeping failures for overuse of antimalarials? Ideally, we would use the home-based RDTs to directly assess the match between prescriptions/purchases and the patient’s underlying malaria status.³⁸ Unfortunately, as discussed earlier, patients who consented to the home-based test were significantly sicker and more likely to have malaria, based on our predicted risk metric, than patients who refused the test (Appendix Table B3). Since *PV-DV* differences in malaria prescriptions and purchases are concentrated among those with the lowest predicted malaria risk, relying on the home-tested subsample could understate *PV-DV* differences in appropriate treatment.

As an alternative, we focus on “expected match quality” based on predicted malaria risk (denoted $\tilde{\pi}$), which we decompose into two parts. First, the probability of a “correct positive”, i.e. that an antimalarial is given to a truly sick patient, is given by $mp_{ict} = \tilde{\pi}_{ict} \times antimal_{ict}$, where $antimal_{ict}$ is a dummy variable equal to one if patient i at clinic c on day t was either prescribed or purchased an antimalarial of any type. Second, the probability that an antimalarial is withheld from a patient without malaria – a “correct negative” – is $mn_{ict} = (1 - \tilde{\pi}_{ict}) \times (1 - antimal_{ict})$. We consider treatment effects on both components separately, as well as the “expected match”, which for the sake of illustration puts equal weight on both positive outcomes: $match_{ict} = mp_{ict} + mn_{ict}$. In practice, the social planner may put very different weights on these two outcomes.

Table 5 shows treatment effects on the expected match and its two components, for both antimalarial prescriptions and purchases. While the voucher treatments did not improve

³⁷The greater number of patients on voucher days versus control days complicates our test of prediction (6), so we place more weight on prediction (7).

³⁸For the purposes of this discussion we abstract away from the decision to prescribe simple versus severe malaria treatment, since it is difficult, even with our rich data, to reliably identify which patients should get which type of malaria treatment. Assuming most patients in our sample should get simple malaria treatment, this abstraction should work to understate improvements in match quality associated with the Doctor Voucher condition.

correct positives (mp_{ict}) for prescriptions, they did for purchases (significant at the 1 percent level). The voucher treatments increased the correct positive purchase rate by 1.8-2.3 percentage points, which amounts to an 15-19 percent improvement over the mean in the Control group. This underscores the rationale of subsidizing health care: some truly sick patients will not purchase treatment due to credit constraints or because they value treatment below its cost.

This improvement comes at a cost, however. While mn does not get worse for prescriptions in DV , it does in PV (column 2); and mn significantly declines for both voucher conditions when considering purchases. Column 5 shows that an estimated 45 percent of patients do not have malaria *and* do not purchase an antimalarial in the Control group. PV decreases this share by 11 percentage points, or 24 percent, while DV decreases this share by 6.3 percentage points, or 14 percent. This underscores the concern that subsidies lead to overuse. The PV - DV difference is significant at the 1 percent level.

Turning to the expected match, we estimate that 48 percent of patients are prescribed the proper (malaria-related) treatment in the control group, and 57 percent purchase it. PV significantly decreases this share for both prescriptions (3.2 percentage points) and purchases (8.3 percentage points), while DV only decreases this share for purchases (3.9 percentage points). The 8.3 percentage point decline in the match rate amounts to a 19 percent *increase* in misallocation, which is substantial, and driven entirely by patient behavior: doctor preferences (captured by the effect of DV on the expected match for prescriptions) have no impact on misallocation, while distortions in prescription behavior from patient-driven demand and the decline patients' non-purchase rate each account for roughly half of the additional misallocation introduced by the ACT subsidy.

A weakness of our expected match measure is that it assumes no private information on the part of doctors. Doctors may have access to private signals – in particular malaria tests – that improve upon $\tilde{\pi}$ when making prescription decisions.³⁹ If this were the case, our expected match measures would be biased down, because private signals would improve the match of treatment to illness *conditional* on $\tilde{\pi}$. Assessing the consequences for estimated treatment effects is more difficult, since this depends on what drives “marginal” prescriptions in the different experimental conditions. For example, suppose all marginal prescriptions in PV were driven by malaria negative patients. In this case, we would falsely estimate a positive effect on mp (since marginal patients will have $\tilde{\pi} > 0$) and underestimate the negative effect on mn .

We can use the home-tested RDT subsample to get an idea of how much private information skews our expected match estimates. Appendix Table B10 shows treatment effects on

³⁹This is captured by η in our model.

the expected and actual match for the home-tested subsample. The estimates are strikingly similar: the expected match for prescriptions in the control group is 0.48, while the actual match is 0.51. Analogous numbers for purchases are 0.56 and 0.61. While these numbers are consistent with some private information, the impacts on match estimates are very modest, likely because doctors often do not use, or do not comply with malaria tests. Estimated treatment effects on the expected and actual outcomes are also very similar.

An important question is if over-prescription and over-purchasing of medication actually led to overtreatment, or if we captured a partial equilibrium effect, where patients learned about the vouchers and sought to “stockpile” medication for future illness episodes. Differences between PV and DV due to stockpiling would be consistent with patient-driven demand, but may not occur in a long-run equilibrium where subsidies are permanent. We check for stockpiling using our home survey data, where we cycled through the list of all medications purchased at the clinic and asked if the patient was currently taking the medication.⁴⁰ Appendix Table B11, column 2, shows that 94 percent of individuals who purchased ACTs for simple malaria at the clinic reported taking ACTs during the home survey, with no significant differences by voucher condition. Barring the possibility of poor self-reporting, this suggests that stockpiling is not a significant driver of the higher purchase rates in PV , and that higher sales also led to greater overtreatment.

Of course, the general equilibrium effects of a real world subsidy policy would likely still differ from the ones we estimate here: in the long run (potential) patients will learn of subsidies regardless of whether or not they are widely advertised, and this will change selection into the clinic; the direction of this selection is *a priori* unclear. Regardless of selection, the scale of our patient-driven demand effects has important implications. If even the best-trained doctors are susceptible to patient requests for unnecessary medication, then subsidizing care in the public sector will likely deliver the benefit of getting needed medication to the truly sick, but at the cost of overtreatment, wasted subsidy dollars, and, in the case of vital medications like antimalarials and antibiotics, an elevated risk of disease resistance.

6 Conclusion

A critical challenge for developing country health systems is to craft policy that simultaneously avoids undertreatment and overtreatment. Undertreatment remains a major problem; for example, treatable conditions like pneumonia, diarrhea, and malaria account for 48 percent of global postnatal child mortality (WHO, 2018). Overtreatment is problematic for a

⁴⁰ACT treatment courses last for three days and individuals were interviewed the day after the clinic visit, so the treatment course should be ongoing at the time of the home survey.

range of reasons: poor households with very high marginal utilities of income waste precious resources, overuse of vital drugs like antimalarials and antibiotics promotes disease resistance that will cost lives in the future, and overuse of government-subsidized care strains already-tight public budgets.

Doctors (and other health workers) are the first line of defense against overtreatment when care is subsidized. Yet a growing body of evidence makes it abundantly clear that doctors in both the public and private sectors are imperfect gatekeepers, prescribing unnecessary treatment at very high rates. It is therefore essential to understand the factors that limit doctors' ability to target treatments to the right patients.

The existing literature has largely focused on doctor-driven overtreatment. The implicit assumption is that doctors *prefer* to overtreat, for example because it is in their financial interest, or because they do not care to exert effort to arrive at a better diagnosis. Our paper broadens the scope of investigation by asking whether patients' preferences also drive overtreatment: while the doctor would prefer not to prescribe a medication, she is willing to do so in order to avoid the "gatekeeping costs" of dealing with an unhappy or distrustful patient. While this motivation is often anecdotally reported by doctors, there is limited rigorous work on its relevance.

We begin by developing a model that is channel agnostic, in that it can generate both doctor- and patient-driven overtreatment, and delivers testable implications for both channels. We test the model's predictions with an RCT that we conducted with 60 public health clinics in Bamako, Mali. Our context is well-suited to studying this issue, as misallocation of malaria treatment in public clinics is both verifiable (via follow-up RDTs), and rampant: 75 percent of patients with confirmed RDTs tested negative, yet 58 percent of these malaria-negative patients were prescribed some form of malaria treatment.

The results of our empirical investigation are clear: vouchers for free malaria treatment are much more likely to be used when both patients and doctors (as opposed to just doctors) know about their availability, indicating that gatekeeping costs are empirically important. These marginal voucher redemptions are driven by patients with the lowest levels of malaria risk, which is consistent with patient-driven demand. At the same time, we find no evidence that doctors in the Doctor Voucher treatment exploit their informational advantage to upsell patients more expensive severe malaria treatment, or increase clinic revenues more generally.

While our results by no means rule out the possibility that doctor preferences are driving a share of the baseline misallocation we see in our sample, they show that marginal misallocation from price subsidies is attributable to patient preferences and behavior. In fact, doctors are no more likely to write an antimalarial prescription when only they know about the price discounts; this margin only moves when both doctors *and* patients are informed. This

evidence is not just important for enriching our understanding of doctor preferences, patient preferences, and how they interact during the consultation process, our results also contribute to the ongoing debate over how to price health products in the developing world. We show that cost sharing in public sector settings can reduce overtreatment by reducing patient demand, but at the cost of fewer malaria patients getting appropriate treatment. Policies and tools that improve physicians' gatekeeping capabilities could help sustain higher subsidies while limiting overuse; additional research on the design and efficacy of such interventions is an important area for future work.

Our results are critically determined by the *mismatch* between doctor and patient preferences – patients prefer more treatment than doctors would like to prescribe. Doctor-patient preference gaps may be different in other contexts where doctors face different incentives and patients face different prices. Both our model and empirical approach can be adapted to other settings to generate customized evidence – and a more nuanced understanding of how gatekeeping failures mediate healthcare outcomes worldwide.

References

- Achan, J., A. O. Talisuna, A. Erhart, A. Yeka, J. K. Tibenderana, F. N. Baliraine, P. J. Rosenthal, and U. D'Alessandro (2011). Quinine, an Old Anti-Malarial Drug in a Modern World: Role in the Treatment of Malaria. *Malaria Journal* 10(1), 144.
- Ansah, E. K., S. Narh-Bana, M. Epokor, S. Akanpigiabiam, A. A. Quartey, J. Gyapong, and C. J. Whitty (2010). Rapid testing for malaria in settings where microscopy is available and peripheral clinics where only presumptive treatment is available: A randomised controlled trial in Ghana. *BMJ* 340, c930.
- Ashraf, N., J. Berry, and J. M. Shapiro (2010). Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia. *American Economic Review* 100(5), 2383–2413.
- Banerjee, A. V., E. Duflo, and R. Glennerster (2008). Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System. *Journal of the European Economic Association* 6(2-3), 487–500.
- Belloni, A., V. Chernozhukov, and C. Hansen (2014, May). High-Dimensional Methods and Inference on Structural and Treatment Effects. *Journal of Economic Perspectives* 28(2), 29–50.
- Bisoffi, Z., B. S. Sirima, A. Angheben, C. Lodesani, F. Gobbi, H. Tinto, and J. Van den Ende (2009). Rapid malaria diagnostic tests vs. clinical management of malaria in rural Burkina Faso: safety and effect on clinical decisions. A randomized trial. *Tropical Medicine & International Health* 14(5), 491–498.
- CDC (2015). Malaria. <http://www.cdc.gov/malaria/about/disease.html>.
- Cohen, J. and P. Dupas (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *The Quarterly Journal of Economics* 125(1), 1–45.
- Cohen, J., P. Dupas, and S. Schaner (2015). Price Subsidies, Diagnostic Tests, and Targeting of Malaria Treatment: Evidence from a Randomized Controlled Trial. *American Economic Review* 105(2), 609–45.
- Currie, J., W. Lin, and J. Meng (2014). Addressing Antibiotic Abuse in China: An Experimental Audit Study. *Journal of Development Economics* 110, 39–51.

- Das, J. and J. Hammer (2014). Quality of Primary Care in Low-Income Countries: Facts and Economics. *6*, 525–553.
- Das, J., J. Hammer, and K. Leonard (2008). The Quality of Medical Advice in Low-Income Countries. *The Journal of Economic Perspectives* 22(2), pp. 93–114.
- Das, J., A. Holla, A. Mohpal, and K. Muralidharan (2016, December). Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India. *American Economic Review* 106(12), 3765–99.
- Djimde, A. A., H. Maiga, I. Sagara, K. Traore, H. Niangaly, A. Bamadio, A. Grivoyannis, M. Tekete, O. B. Traore, A. Kodio, et al. (2016). Field Assessment of SD Bioline Malaria Species Antigen Detection by Rapid Diagnostic Tests in Mali. *Journal of Parasitology and Vector Biology* 8(1), 1–9.
- Dupas, P. (2014). Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment. *Econometrica* 82(1), 191–228.
- Dupas, P. and E. Miguel (2017). Impacts and Determinants of Health Levels in Low-Income Countries. In *Handbook of Economic Field Experiments*, Volume 2, pp. 3–93. Elsevier.
- Fleming-Dutra, K. E., A. L. Hersh, D. J. Shapiro, M. Bartoces, E. A. Enns, T. M. File, J. A. Finkelstein, J. S. Gerber, D. Y. Hyun, J. A. Linder, et al. (2016). Prevalence of Inappropriate Antibiotic Prescriptions among US Ambulatory Care Visits, 2010-2011. *JAMA* 315(17), 1864–1873.
- Hamer, D. H., M. Ndhlovu, D. Zurovac, M. Fox, K. Yeboah-Antwi, P. Chanda, N. Sipilinyambe, J. L. Simon, and R. W. Snow (2007). Improved diagnostic testing and malaria treatment practices in Zambia. *JAMA* 297(20), 2227–2231.
- Humar, A., C. Ohrt, M. A. Harrington, D. Pillai, and K. C. Kain (1997). ParaSight® F Test Compared with the Polymerase Chain Reaction and Microscopy for the Diagnosis of Plasmodium Falciparum Malaria in Travelers. *The American Journal of Tropical Medicine and Hygiene* 56(1), 44–48.
- IHME (2016). Global burden of disease study 2015 (gbd 2015) results. Online. Accessed August 22, 2016.
- Iizuka, T. (2012, May). Physician Agency and Adoption of Generic Pharmaceuticals. *American Economic Review* 102(6), 2826–58.

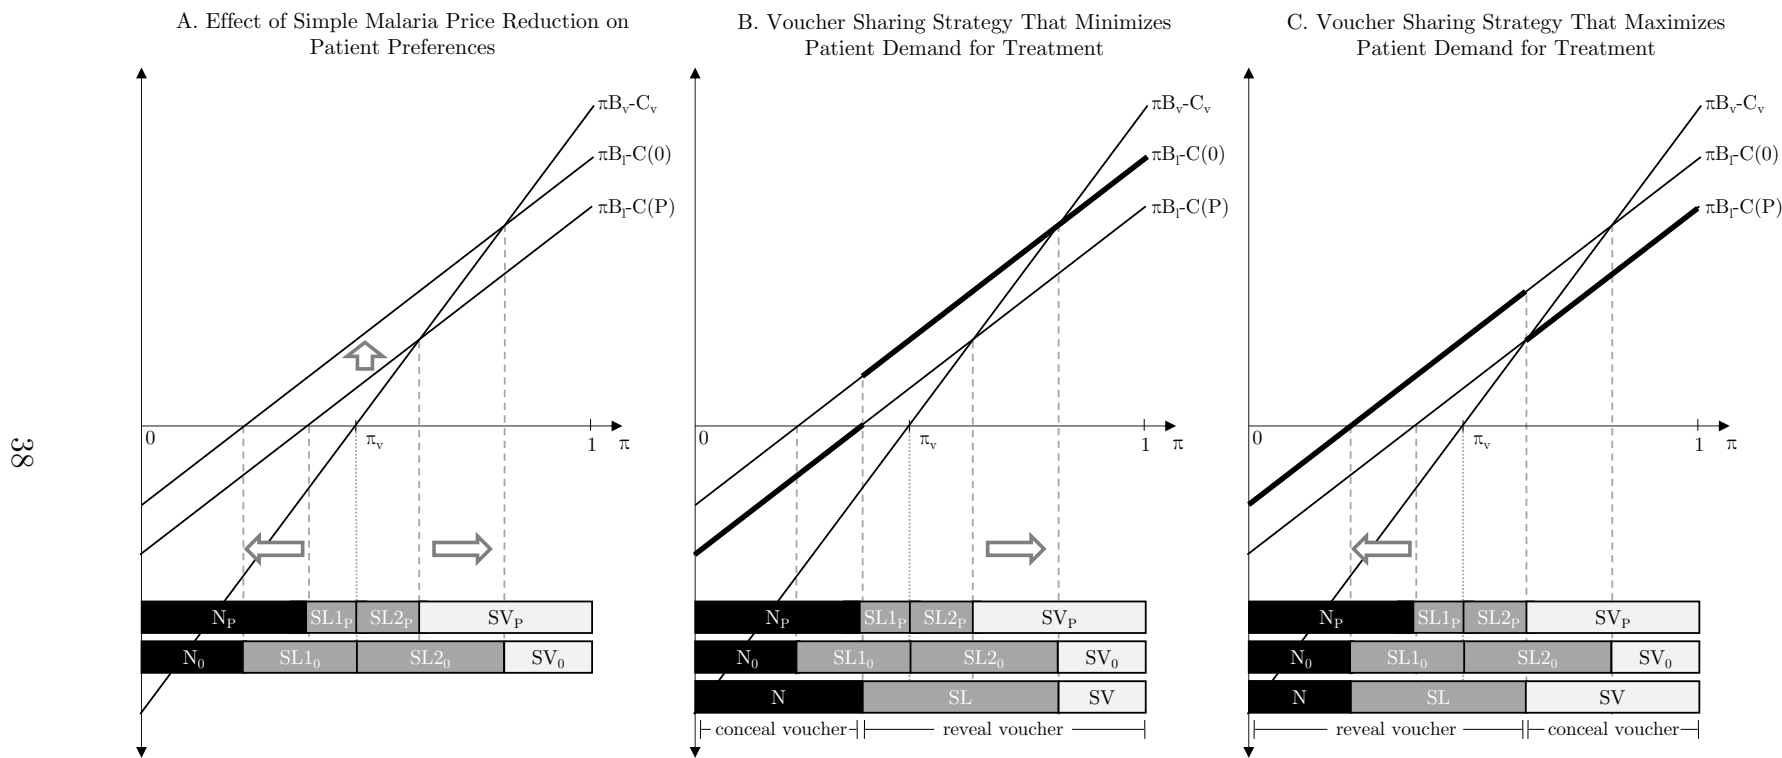
- Koné, D., D. Coulibaly, O. Doumbo, I.-S. Fall, E. Kibuchi, B. Mitto, G. Sang, D. Kyalo, R. W. Snow, and A. M. Noor (2015). *An Epidemiological Profile of Malaria in Mali*. Programme National de Lutte contre le Paludisme (PNLP).
- Kotwani, A., C. Wattal, S. Katewa, P. Joshi, and K. Holloway (2010). Factors Influencing Primary Care Physicians to Prescribe Antibiotics in Delhi India. *Family Practice* 27(6), 684–690.
- Kravitz, R. L., R. M. Epstein, M. D. Feldman, C. E. Franz, R. Azari, M. S. Wilkes, L. Hinton, and P. Franks (2005). Influence of Patients’ Requests for Direct-to-Consumer Advertised Antidepressants: A Randomized Controlled Trial. *JAMA* 293(16), 1995–2002.
- Kremer, M., J. Leino, E. Miguel, and A. Zwane (2011). Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions. *The Quarterly Journal of Economics* 126(1).
- Kyabayinze, D. J., J. K. Tibenderana, G. W. Odong, J. B. Rwakimari, and H. Counihan (2008). Operational Accuracy and Comparative Persistent Antigenicity of HRP2 Rapid Diagnostic Tests for Plasmodium Falciparum Malaria in a Hyperendemic Region of Uganda. *Malaria Journal* 7(1), 221.
- Li, Y., J. Xu, F. Wang, B. Wang, L. Liu, W. Hou, H. Fan, Y. Tong, J. Zhang, and Z. Lu (2012). Overprescribing in China, Driven by Financial Incentives, Results in Very High Use of Antibiotics, injections, and corticosteroids. *Health affairs* 31(5), 1075–1082.
- Linder, J. A., J. N. Doctor, M. W. Friedberg, H. R. Nieva, C. Birks, D. Meeker, and C. R. Fox (2014). Time of Day and the Decision to Prescribe Antibiotics. *JAMA Internal Medicine* 174(12), 2029–2031.
- Ministère de la Santé (2013, August). *Politique Nationale De Lutte Contre Le Paludisme*. Republique du Mali.
- PMI (2015). President’s Malaria Initiative Mali: Malaria Operational Plan FY 2016. Technical report, USAID.
- PMI (2016). President’s Malaria Initiative Mali: Malaria Operational Plan FY 2017. Technical report, USAID.
- PMI (2017). President’s Malaria Initiative Mali: Malaria Operational Plan FY 2018. Technical report, USAID.

- PNLP, INSTAT, INFO-STAT, INRSP, and ICF International (2016). Enquete sur les Indicateurs du Paludisme au Mali (EIPM) 2015. Technical report, INSTAT, INFO-STAT and ICF International, Rockville, Maryland, USA.
- Prendergast, C. (2003). The Limits of Bureaucratic Efficiency. *Journal of Political Economy* 111(5), 929–958.
- Reyburn, H., R. Mbatia, C. Drakeley, I. Carneiro, E. Mwakasungula, O. Mwerinde, K. Saganda, J. Shao, A. Kitua, R. Olomi, B. Greenwood, and C. Whitty (2004). Overdiagnosis of malaria in patients with severe febrile illness in Tanzania: A prospective study. *British Medical Journal* 329(7476), 1212.
- Sautmann, A., S. Brown, and M. Dean (2017). Subsidies, Information, and the Timing of Children’s Health Care in Mali.
- Tarozzi, A., A. Mahajan, B. Blackburn, D. Kopf, L. Krishnan, and J. Yoong (2014). Microloans, Insecticide-Treated Bednets and Malaria: Evidence from a randomized controlled trial in Orissa (India). *The American Economic Review* 104(7), 1909–41.
- Taylor, S. M., M. E. Molyneux, D. L. Simel, S. R. Meshnick, and J. J. Juliano (2010). Does This Patient Have Malaria? *JAMA: The Journal of the American Medical Association* 304(18), 2048–2056.
- Trampuz, A., M. Jereb, I. Muzlovic, and R. M. Prabhu (2003). Clinical Review: Severe Malaria. *Critical Care* 7(4), 315.
- Whitehead, M., G. Dahlgren, and T. Evans (2001). Equity and Health Sector Reforms: Can Low-Income Countries Escape the Medical Poverty Trap? *The Lancet* 358(9284), 833–836.
- WHO (2014). Emergence and Spread of Artemisinin Resistance Calls for Intensified Efforts to Withdraw Oral Artemisinin-Based Monotherapy from the Market. Technical report, World Health Organization.
- WHO (2017). Malaria rapid diagnostic test performance: results of who product testing of malaria rdts: round 7 (2015-2016). Technical report, World Health Organization.
- WHO (2018). MCEE Methods and Data Sources for Child Causes of Death 2000-2016.
- Wongsrichanalai, C., M. J. Barcus, S. Muth, A. Sutamihardja, and W. H. Wernsdorfer (2007). A Review of Malaria Diagnostic Tools: Microscopy and Rapid Diagnostic Test (RDT). *The American Journal of Tropical Medicine and Hygiene* 77(6_Suppl), 119–127.

World Health Organization (2015). Malaria Rapid Diagnostic Test Performance: Results of WHO product testing of malaria RDTs: round 6 (2014–2015). Technical report.

Main Figures and Tables

Figure 1: Patient Preferences and Doctor Voucher-Sharing Strategies



Notes: Diagonal lines graph patient utility net of costs. Shaded bars with P and 0 subscripts indicate patient preferences when the price of simple malaria treatment is P , and 0 respectively. N denotes "no antimalarial purchase", $SL1$ denotes "simple malaria treatment, unwilling to purchase severe malaria treatment", $SL2$ denotes "simple malaria treatment, willing to purchase severe malaria treatment", and SV denotes "severe malaria treatment". The bottom shaded bars in panels B and C show the outcome when the doctor adopts the strategy specified in the panel title.


Figure 2: Within-CSCOM Randomization Design

CSCOM Number	WEEK 1							WEEK 2						
	Mon	Tues	Weds	Thurs	Fri	Sat	Sun	Mon	Tues	Weds	Thurs	Fri	Sat	Sun
1	C	--	PV	--	DV	--	--	PI	--	PI-PV	--	PI-DV	--	--
2	DV	--	C	--	PV	--	--	PI-DV	--	PI	--	PI-PV	--	--
3	PV	--	DV	--	C	--	--	PI-PV	--	PI-DV	--	PI	--	--
4	C	--	DV	--	PV	--	--	PI	--	PI-DV	--	PI-PV	--	--
5	DV	--	PV	--	C	--	--	PI-DV	--	PI-PV	--	PI	--	--
6	PI	--	PI-PV	--	PI-DV	--	--	C	--	PV	--	DV	--	--
7	PI-DV	--	PI	--	PI-PV	--	--	DV	--	C	--	PV	--	--
8	PI-PV	--	PI-DV	--	PI	--	--	PV	--	DV	--	C	--	--
9	PI	--	PI-DV	--	PI-PV	--	--	C	--	DV	--	PV	--	--
10	PI-DV	--	PI-PV	--	PI	--	--	DV	--	PV	--	C	--	--
11	--	C	--	PV	--	DV	--	--	PI	--	PI-PV	--	PI-DV	--
12	--	DV	--	C	--	PV	--	--	PI-DV	--	PI	--	PI-PV	--
13	--	PV	--	DV	--	C	--	--	PI-PV	--	PI-DV	--	PI	--
14	--	C	--	DV	--	PV	--	--	PI	--	PI-DV	--	PI-PV	--
15	--	DV	--	PV	--	C	--	--	PI-DV	--	PI-PV	--	PI	--
16	--	PI	--	PI-PV	--	PI-DV	--	--	C	--	PV	--	DV	--
17	--	PI-DV	--	PI	--	PI-PV	--	--	DV	--	C	--	PV	--
18	--	PI-PV	--	PI-DV	--	PI	--	--	PV	--	DV	--	C	--
19	--	PI	--	PI-DV	--	PI-PV	--	--	C	--	DV	--	PV	--
20	--	PI-DV	--	PI-PV	--	PI	--	--	DV	--	PV	--	C	--

LEGEND	
--	No data collection or interventions at CSCOM
C	Data collection at CSCOM, no interventions
DV	Doctor vouchers and data collection at CSCOM
PV	Patient vouchers and data collection at CSCOM
PI	Patient information and data collection at CSCOM
PI-DV	Patient information, doctor vouchers, and data collection at CSCOM
PI-PV	Patient information, patient vouchers, and data collection at CSCOM

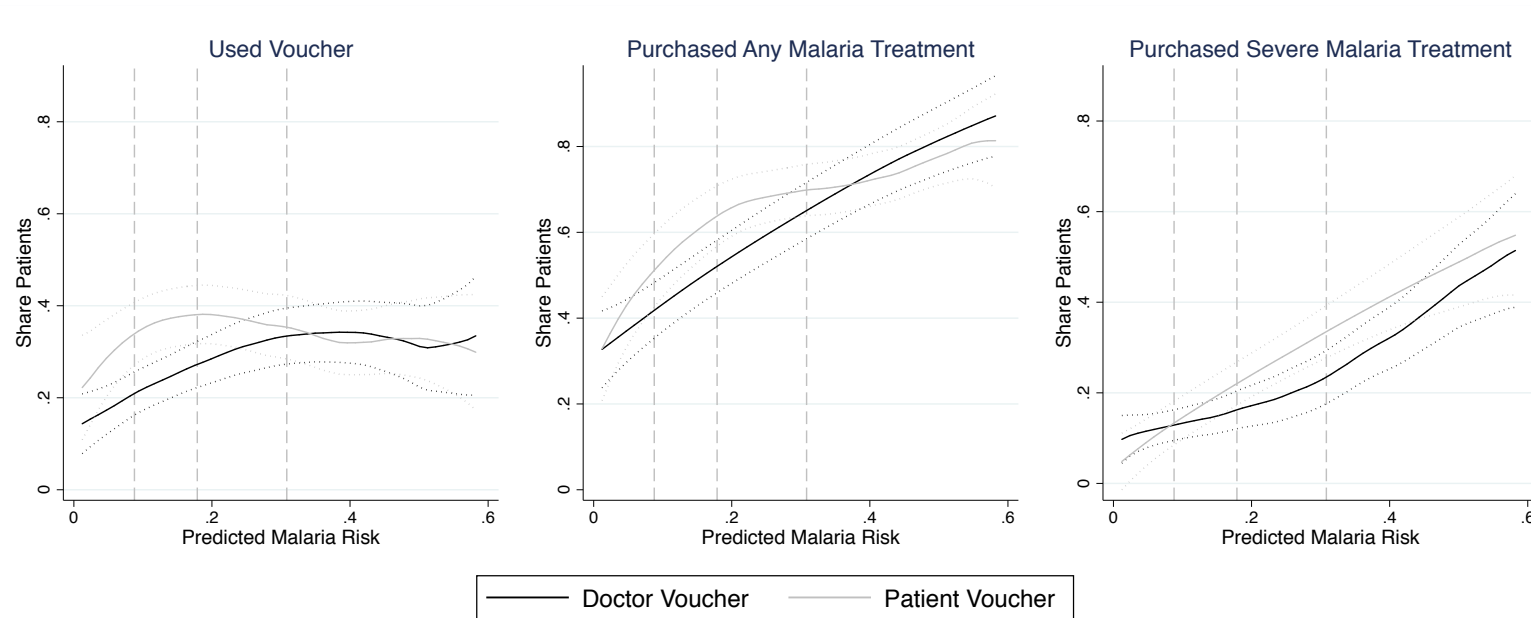
Notes: The interventions listed above ran between November 14-December 30 2016 in three two-week blocks, with 20 CSCOMs active in each two-week block.

Figure 3: Treatment Outcomes by Predicted Malaria Risk in Control Group




Notes: Results from local linear regressions. Regressions are run on the full sample, but graphs omit results for top and bottom 2.5 percent of malaria risk distribution to avoid influence of outliers. Vertical dashed lines indicate 25th, 50th, and 75th percentiles of predicted malaria risk respectively.

Figure 4: Treatment Outcomes by Predicted Malaria Risk and Voucher Group



Notes: Results from local linear regressions. Dotted lines give 90 percent confidence intervals. The standard errors used to calculate confidence intervals are bootstrapped by re-sampling clinics and recalculating predicted malaria risk on each of 1,000 bootstrap replications. Regressions are run on the full sample, but graphs omit results for top and bottom 2.5 percent of malaria risk distribution to avoid influence of outliers. Vertical dashed lines indicate 25th, 50th, and 75th percentiles of predicted malaria risk respectively.

Figure 5: Use of Vouchers for Simple and Severe Malaria by Treatment Arm



Notes: Each panel graphs the average value of the specified outcome by voucher condition.

Table 1: Descriptive Statistics from Control Group

	(1)	(2)	(3)	(4)	(5)
			Has Valid RDT (Home Survey)		
	Full Sample	N	All Patients	RDT Positive	RDT Negative
<i>Panel A. Malaria Beliefs and Treatment Outcomes</i>					
Respondent Suspects Illness is Malaria (Pre-Consultation)	0.551	628	0.61	0.791	0.55
Simple Malaria Treatment: Prescribed	0.313	627	0.295	0.163	0.338
Simple Malaria Treatment: Purchased	0.206	627	0.197	0.14	0.215
Severe Malaria Treatment: Prescribed	0.303	627	0.327	0.593	0.238
Severe Malaria Treatment: Purchased	0.255	627	0.283	0.535	0.2
<i>Panel B. Other Treatment Outcomes</i>					
Purchased Antibiotics	0.426	627	0.442	0.43	0.446
Received Injection or IV	0.398	628	0.454	0.674	0.381
Patient Referred to Hospital or Placed Under Observation	0.139	628	0.179	0.395	0.108
<i>Panel C. Patient Costs</i>					
Total Cost No Malaria Treatment (CFA)	2793	301	–	–	–
Total Cost Simple Malaria (CFA)	4475	126	–	–	–
Total Cost Severe Malaria (CFA)	10353	147	–	–	–

Notes: Variables measured in CFA top-coded at the 99th percentile. CFA610 \approx USD1. Sample sizes for columns 3, 4, and 5 are 346, 86, and 260 respectively.

Table 2: Demographic Characteristics and Randomization Verification

	(1)	(2)		(3)	(4)		(5)	(6)
	Control Mean	Regression Coefficients		Doctor Voucher	Joint Test PV=DV	Joint Test PV=DV=0		N
<i>A. Sample Frame (Clinic × Day-Level Observations)</i>								
Number Eligible Logged Patients	5.23 [3.32]	0.725* (0.369)		0.7** (0.313)	0.94	0.062*		360
<i>B. Patient Characteristics (Individual-Level Observations)</i>								
Number of symptoms	3.56 [1.58]	0.033 (0.098)		-.142 (0.106)	0.036**	0.101		2055
Fever	0.801 [0.4]	0.04 (0.024)		0.004 (0.026)	0.066*	0.106		2055
Chills or Excessive Sweating	0.298 [0.458]	-.035 (0.032)		-.055** (0.025)	0.415	0.094*		2055
Nausea, Vomiting, or Diarrhea	0.497 [0.5]	0.004 (0.03)		-.028 (0.03)	0.296	0.527		2055
Poor Appetite	0.505 [0.5]	-.028 (0.032)		-.055* (0.029)	0.3	0.153		2055
Headache	0.621 [0.486]	0 (0.025)		-.014 (0.024)	0.514	0.753		2055
Cough	0.363 [0.481]	0.033 (0.024)		-.004 (0.029)	0.205	0.299		2055
Weakness/Fatigue	0.475 [0.5]	0.019 (0.028)		0.01 (0.03)	0.735	0.791		2055
Duration of Illness in Days	4.3 [4.28]	0.109 (0.228)		-.041 (0.292)	0.585	0.819		2055
Age	16.6 [15.1]	1.19 (0.773)		1 (0.743)	0.84	0.193		2055
Under 5 Years Old	0.314 [0.464]	-.041* (0.024)		-.023 (0.02)	0.464	0.216		2055
Male	0.395 [0.489]	0.029 (0.03)		0.05* (0.029)	0.515	0.229		2055
Patient is Pregnant	0.105 [0.307]	0.008 (0.027)		-.009 (0.021)	0.477	0.755		1139
Positive RDT (home)	0.249 [0.433]	-.045 (0.034)		0.005 (0.03)	0.128	0.271		1126
Predicted Malaria Risk	0.212 [0.16]	0.008 (0.009)		0.002 (0.008)	0.459	0.596		2055
<i>C. Respondent and Household Characteristics (Individual-Level Observations)</i>								
Patient Answered Clinic Survey	0.47 [0.499]	0.004 (0.027)		-.001 (0.024)	0.862	0.983		2055
Male	0.264 [0.441]	0.015 (0.023)		0.026 (0.024)	0.663	0.541		2055
Bambara	0.377 [0.485]	0.011 (0.028)		0.013 (0.027)	0.937	0.881		2053
Speaks French	0.518 [0.5]	-.004 (0.03)		-.001 (0.026)	0.896	0.988		2055
Literate (in French)	0.264 [0.441]	-.006 (0.029)		-.017 (0.027)	0.649	0.802		2055
Primary School or Less	0.424 [0.495]	0.032 (0.032)		0.043 (0.029)	0.735	0.324		2055
Household Size ⁺	10.3 [8.11]	0.413 (0.536)		0.655 (0.607)	0.682	0.539		1491
Share HH Under 15 ⁺	0.412 [0.191]	0.015 (0.013)		0.009 (0.013)	0.67	0.511		1485
Share HH Members Working ⁺	0.258 [0.188]	-.004 (0.01)		0 (0.012)	0.751	0.913		1485
Monthly income per capita ⁺	21988 [25298]	-2343* (1331)		-2907* (1631)	0.658	0.159		1432
Rental Value of Home ⁺	63823 [90722]	-3089 (5702)		-5055 (6085)	0.687	0.71		1469
Mosquito Nets Per Capita ⁺	0.491 [0.342]	-.01 (0.022)		-.006 (0.021)	0.795	0.893		1482

⁺ indicates that variable was recorded in the home survey only. Variables measured in CFA and duration of illness top-coded at the 99th percentile. CFA610 ≈ USD1. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table 3: Impacts on Malaria Treatment Outcomes

	(1)	(2)	(3)	(4)	(5)
		Any Malaria Treatment		Severe Malaria Treatment	
	Used Voucher	Prescribed	Purchased	Prescribed	Purchased
β_P : Patient Voucher	0.35*** (0.030)	0.052** (0.025)	0.14*** (0.027)	-0.046** (0.018)	-0.022 (0.020)
β_D : Doctor Voucher	0.26*** (0.024)	0.016 (0.025)	0.085*** (0.025)	-0.054** (0.021)	-0.050*** (0.019)
<i>P-values: Theory-Driven Tests</i>					
$\beta_P = \beta_D$: Two-sided	0.011**	0.079*	0.018**	0.735	0.193
$\beta_P = \beta_D$: One-sided	0.006***	0.040**	0.009***	0.632	0.903
One-sided test of:	GC	PD	PD	DD	DD
Mean (Control)	0	0.62	0.46	0.30	0.26
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. Eligible controls include clinic dummies, symptom dummies, duration of illness (topcoded at the 99th percentile) and its square, patient age and its square, a dummy for patients under 5, patient gender, a dummy to identify pregnant patients, a dummy to identify whether the patient (versus a caregiver) answered the survey, the gender of the survey respondent, an ethnicity (Bambara) dummy, a dummy for French speaking respondents, a dummy for literate respondents, a dummy for respondents with a primary education or less, a dummy to identify patients in the home-based follow up survey, and pairwise interactions between all previously-listed patient and respondent controls. Missing values are recoded to the sample mean and separately dummied out. These missing dummies are also used to construct pairwise interactions. GC, PD, and DD indicate tests of gatekeeping costs, patient-driven, and doctor-driven demand respectively. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table 4: Impacts on Malaria Treatment Outcomes - Heterogeneity by Predicted Malaria Risk

	(1)	(2)	(3)	(4)	(5)
		Any Malaria Treatment		Severe Malaria Treatment	
	Used Voucher	Prescribed	Purchased	Prescribed	Purchased
δ_{PH} : Patient Voucher \times High Risk	0.34*** (0.039)	0.035 (0.035)	0.12*** (0.037)	-0.055* (0.031)	-0.020 (0.033)
δ_{DH} : Doctor Voucher \times High Risk	0.32*** (0.038)	0.021 (0.034)	0.078** (0.036)	-0.12*** (0.036)	-0.12*** (0.033)
δ_{PL} : Patient Voucher \times Low Risk	0.36*** (0.041)	0.092** (0.036)	0.18*** (0.041)	-0.026 (0.027)	-0.016 (0.029)
δ_{DL} : Doctor Voucher \times Low Risk	0.21*** (0.033)	0.011 (0.035)	0.086** (0.034)	-0.0019 (0.030)	0.0085 (0.026)
θ : High Malaria Risk	0.0029 (0.041)	0.13*** (0.046)	0.13*** (0.048)	0.090* (0.048)	0.093** (0.043)
<i>P-values: Theory-Driven Tests</i>					
$\delta_{PH} = \delta_{DH}$: Two-sided	0.749	0.689	0.275	0.101	0.007***
$\delta_{PH} = \delta_{DH}$: One-sided	0.375	0.345	0.137	0.949	0.997
One-sided test of:	GC/DD	–	–	DD	DD
$\delta_{PL} = \delta_{DL}$: Two-sided	0.000***	0.030**	0.014**	0.481	0.435
$\delta_{PL} = \delta_{DL}$: One-sided	0.000***	0.015**	0.007***	0.241	0.218
One-sided test of:	GC/PD	PD	PD	–	–
Mean (Control, Low Risk)	0	0.49	0.33	0.15	0.12
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. Standard errors are based on 1,000 bootstrap replications, with re-sampling at the clinic level. Predicted malaria risk is re-calculated on each bootstrap replication. GC, PD, and DD indicate tests of gatekeeping costs, patient-driven, and doctor-driven demand respectively. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table 5: Impacts on Match Between Treatment and Illness

	(1)	(2)	(3)	(4)	(5)	(6)
	Expected Match: Prescribed			Expected Match: Purchased		
	Malaria Positive	Malaria Negative	Overall Match	Malaria Positive	Malaria Negative	Overall Match
β_P : Patient Voucher	0.0096 (0.0070)	-0.043** (0.021)	-0.032* (0.017)	0.023*** (0.0079)	-0.11*** (0.022)	-0.083*** (0.017)
β_D : Doctor Voucher	0.0073 (0.0076)	-0.0074 (0.021)	0.0049 (0.017)	0.018*** (0.0071)	-0.063*** (0.020)	-0.039** (0.016)
<i>P-values: Two-Sided Tests</i>						
$\beta_P = \beta_D$	0.718	0.034**	0.006***	0.537	0.008***	0.002***
$\beta_P = \beta_D = 0$	0.384	0.050*	0.015**	0.007***	0.000***	0.000***
Mean (Control)	0.15	0.33	0.48	0.12	0.45	0.57
N	2053	2053	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. The expected match for malaria positive is equal to predicted malaria risk times the relevant malaria treatment/purchase dummy. The expected match for malaria negative is equal to one minus predicted malaria risk times one minus the malaria prescription/purchase dummy. The overall expected match is the sum of these two variables. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

A Theoretical Framework

Malaria Risk and Beliefs

Recall that both doctor and patient receive their own unobservable signal of the patient's malaria risk (ϵ is the patient's signal, η is the doctor's signal). The presence of signal ϵ means that different patients with the same observed symptoms γ may respond differently to the same prescription (e.g. by purchasing or not purchasing the treatment). Since the doctor does not observe ϵ , she will take into account expected patient preferences but cannot fully predict if a given patient will purchase what she prescribes. This matches what we see in the data, where a considerable share of patients do not purchase all items in their prescription. For the same reason, we allow for an unobservable signal η : doctors make different prescription choices based on η for patients who look similar to the observer. Note that this nests the example in the main text, where the doctor observes the patients beliefs, as the variance of ϵ goes to zero.

In the most general setup, the patient may learn from the prescription he receives (and any additional messaging from the doctor) about the doctor's private signal η . In this case, ϵ describes the patient's belief before receiving the prescription, and the patient holds updated belief ϵ' after the prescription is written. Patients' and doctors' beliefs in equilibrium must be rational and mutually consistent, meaning that patients update ϵ correctly based on the average of all η that may lead to the chosen prescription. Given her own signal η , the doctor correctly predicts for each prescription level the patient's distribution of ϵ' and resulting purchasing probability, as well as her own expected utility from treatment and gatekeeping costs, and chooses the optimal prescription. This is reminiscent of informed-expert or cheap-talk models, although note that unlike in a pure cheap-talk setting, the prescription has real consequences, because it prevents the patient from buying options that were not prescribed.

Belief updating after observing the prescription introduces some additional complexity to the model, stemming from the need to derive updated equilibrium beliefs. Moreover, since our interventions have partial equilibrium effects, we have to assume that there will be patients who in the Doctor Voucher treatment have inaccurate expectations of the set of η who would prescribe no treatment or severe treatment (since, unlike the doctor, they are unaware of the existence of the discount). This in turn affects doctors' beliefs. These issues complicate the analysis, but do not provide additional insights into the problem we are considering, so we will assume in what follows that patients do not update their beliefs based on the doctor's prescription behavior. This could be the case if, conditional on observed symptoms, patients believe that doctor preferences are too far apart from their own, so that the prescription choice contains no useful information for them. As a result, we have $\epsilon = \epsilon'$.

Doctor's Expected Utility

Doctors have a range of prescription choices for malaria. They can prescribe simple treatment or severe treatment, or choose not to treat for malaria. In addition, we allow for the possibility that they may offer a menu and give the patient a choice between simple and severe treatment. This option avoids all gatekeeping costs; the only choice the doctor makes is whether she offers the discount to the patient.

We assume that gatekeeping costs are linear in the utility loss the patient experiences from not being able to buy her most preferred treatment, compared to what she will choose given the actual prescription (either buy nothing, or buy the prescribed treatment). Given the unobservability of ϵ , patient preferences are uncertain from the perspective of the doctor (as well as the researcher): two patients who exhibit the same symptoms and receive the same prescription may nonetheless respond differently, depending on their subjective belief about their malaria risk. As a result, the doctor decides based on the distribution of patient types she faces, $F(\pi | \gamma, \eta)$, not the actual type π .

Table A1 describes the doctor's expected utility from different prescription choices in different treatment conditions. Consider for example the doctor's expected utility net of gatekeeping costs in the Control from prescribing no treatment, as shown in Table A1, row (1):

$$EV_n(\gamma, \eta, P) = -g \int_{SL_P} U_l(\pi, P) dF(\pi | \gamma, \eta) - g \int_{SV_P} U_v(\pi) dF(\pi | \gamma, \eta).$$

The doctor experiences the share g of patient's lost utility as a gatekeeping cost. F is the cdf of possible malaria probabilities held by the patient, determined by the distribution of possible values of ϵ , given γ and η . We assume that $F(\pi | \gamma_1, \eta)$ is first-order stochastically dominated (FOSD) by $F(\pi | \gamma_2, \eta)$ if $\gamma_1 < \gamma_2$. This implies that the patient has on average a stronger preference for treatment if observable malaria symptoms are stronger, all else equal, and the mass of patients shifts from lower subjective malaria probabilities π to higher ones. The FOSD condition on F means that the gatekeeping costs from not prescribing any treatment are weakly increasing (in absolute terms) in γ for each ϵ , because the patient's utility from simple and severe treatment is increasing in π .² The expected utility from prescribing severe treatment in the Control is in row (3). The gatekeeping cost is lower at any γ, η than from not prescribing anything, as it only affects those patients who would like to buy simple treatment, but buy nothing or severe treatment instead. Gatekeeping costs of

²The composite function that is 0 on N , $gU_l(\pi, P)$ on SL , and $gU_v(\pi)$ on SV is weakly increasing, so its expectation is weakly increasing as γ increases.

prescribing severe are first increasing, then decreasing in π .

The expected utility from prescribing simple treatment in the control is given by row (5). The utility loss from gatekeeping increases in γ , because the expected gatekeeping cost of not prescribing severe rises as malaria symptoms worsen. Moreover, the gatekeeping costs in EV_l are always lower than in EV_n .

Finally, row (8) shows the expected utility from offering the patient a menu. This option avoids all gatekeeping costs, and provides utilities V_v and V_l according to the probability that the patient chooses severe or simple treatment, respectively.

The doctor can decide whether or not she wants to offer the voucher when prescribing simple malaria treatment. Doctor utility and gatekeeping costs are unchanged between the control and the doctor voucher treatment when the voucher is not used, per lines (1), (3), (5), and (8). This is because it is a weakly dominating strategy not to reveal the lower price of simple treatment in this situation. Thus, the patient's utility and beliefs are exactly the same as in the Control. By contrast, there *is* a difference between C and DV when offering simple treatment and the voucher is revealed (rows (5) and 9)). In these cases, utility in DV is the same as in PV . Rows (7) and (10) show the utility of offering simple, but not using the voucher in PV . This is the only instance where the doctor would incur a gatekeeping cost when offering the choice menu.

Finally, observe that gatekeeping costs are highest when no treatment is prescribed and malaria medications are subsidized (row (2)), and lowest when giving the patient the choice between simple and severe treatment, as long as the voucher is not withheld when the patient knows about it (rows (8) and (9)).

Recall that the doctor's preferences are described by areas \hat{N} ($V_v(\hat{\pi}) < V_l(\hat{\pi}, \cdot) < 0$), $\hat{S}L$ ($V_l(\hat{\pi}, \cdot) \geq 0$, $V_l(\hat{\pi}, \cdot) > V_v(\hat{\pi})$), and $\hat{S}V$ ($V_v(\hat{\pi}) \geq V_l(\hat{\pi}, \cdot) > 0$) across the range of malaria probabilities $\hat{\pi}$. While the patient's purchasing is probabilistic, at a given $\hat{\pi}$ and price, the doctor's innate preferences (excluding gatekeeping costs) are fully described by the ordering of V_l and V_v .

Analyzing the Model

Comparing Doctor Voucher and Patient Voucher Treatments. Recall that *patients drive demand* when the doctor prescribes (and the patient purchases) more aggressive treatment than the doctor herself thinks is optimal. Analogously, *doctors drive demand* when doctors lead patients to purchase more (powerful) treatment than patients find optimal.

Note first that gatekeeping costs increase unambiguously for all prescription choices except those that offer the patient simple treatment with the voucher when going from DV

to PV , because the patients learn that they are missing out on the discount. For no treatment, severe treatment, simple treatment without a voucher, or the choice menu without the voucher, it is straightforward to calculate that the relative utility loss from being in PV over DV is identically given by

$$\begin{aligned} & - \int_{SL_0} gU_l(\pi, 0)dF(\pi|\gamma, \eta) - \int_{SV_0} gU_v(\pi)dF(\pi|\gamma, \eta) \\ & + \int_{SL_P} gU_l(\pi, P)dF(\pi|\gamma, \eta) + \int_{SV_P} gU_v(\pi)dF(\pi|\gamma, \eta) < 0 \end{aligned}$$

(see e.g. row (8) vs. (10)). By contrast, the utility from prescribing simple with the voucher or the choice menu with the voucher remains exactly the same (see rows (6) and (9)). As a result, any change in prescription behavior between DV and PV must involve a switch from one of the options *without* voucher to one of the option *with* the voucher. We refer to this as observation (1), which immediately establishes prediction (1), that voucher use will be higher in PV than DV whenever there are gatekeeping costs and doctors and patients have different preferences over the optimal prescription.

Next, we want to establish prediction (2): that an increase in the overall rate of malaria treatment in PV versus DV indicates patient-driven demand. By observation (1), such a change can only be driven by a doctor who prescribes no treatment in DV , but simple treatment (with voucher) in PV . To complete the proof, we need to show that the doctor switches to prescribing simple treatment against her own preferences.

Not prescribing any treatment incurs high expected gatekeeping costs, and so it can only be preferred to simple treatment if the direct utility from simple treatment $V_l(\hat{\pi}, 0)$ is negative, and therefore also the utility from severe treatment, $V_v(\hat{\pi})$ (area \hat{N}). However, the doctor will switch for example to simple treatment in the patient-voucher arm if $P(\pi \in SL_0 \cup SV_0)V_l(\hat{\pi}, 0) \geq -g \int_{SL_0 \cup SV_0} U_l(\pi, 0)dF(\pi|\gamma, \eta)$, that is, the direct utility loss V_l is smaller in magnitude (less negative) than the expected gatekeeping cost from refusing the patient once he has learned about the discount. In short, the doctor is in \hat{N} , at the lower end of the belief spectrum, and would prefer not to treat, but is made to do so by the expected discontent of those patients who want the simple treatment at the lower price. This completes the argument.

Now we turn to prediction (3): that a higher rate of severe treatment in DV as compared to PV indicates doctor-driven demand. Marginal severe prescriptions in DV could either be from a doctor who prescribed (only) severe treatment, or who gave the patient a choice between simple and severe, but without revealing the voucher (the latter stems from observation (1) and the fact that more patients buy simple treatment in PV – this cannot

occur if all of SL_0 already purchased simple treatment in DV).

Giving a choice without revealing the voucher immediately indicates that $V_v(\hat{\pi}) > V_l(\hat{\pi}, 0)$, or else the doctor could have simply used the voucher to sell more simple treatment; that is, we are in $S\hat{V}_0$. Similarly, a prescription of severe treatment (only) over giving a choice indicates a strong preference for severe treatment over simple treatment, since the doctor can compel patients in $SL2$ to purchase, but at the cost of not selling simple to $SL1$, and gatekeeping costs from both types of patients. In short, whenever we observe the switch from severe to simple, it comes from doctors who prefer severe over simple, but patients who prefer simple over severe. This leads to prediction (3).

Last, prediction (4) follows from the fact that all switches from no treatment to simple under patient-driven demand occur when $\hat{\pi} \in \hat{N}_0$, but switches from severe treatment to simple occur when $\hat{\pi} \in S\hat{V}_0$.

Will Doctors Always Use Vouchers for Simple Treatment?

In our data, we observe patients in both PV and DV who purchase simple malaria treatment without a voucher. While this could be due to issues like doctor inattention, our model predicts that it is sometimes optimal for doctors to withhold vouchers. This can only be the case if the doctor feels compelled (by gatekeeping costs) to prescribe simple treatment, but would actually rather not sell it, either because she prefers to sell more severe treatment ($S\hat{V}_0$), or less treatment overall (\hat{N}_0). Concealing the voucher reduces patient demand for simple treatment. The doctor strikes a balance between gatekeeping costs (her strategy avoids gatekeeping costs for patients who buy simple at P) and prescribing her preferred treatment (her strategy ensures marginal patients who would only purchase ACTs when they are free will not take treatment). The utility from any prescription that involves simple treatment *without* a voucher shrinks from DV to PV , and doctors will substitute to an option that offers simple *with* the voucher. This leads to prediction (5):

Prediction (5) If the doctor prefers not to sell simple treatment, she may choose to prescribe and sell it without a voucher to some patients in DV. From DV to PV, the rate of prescribing simple without a voucher will decrease.

Thus, prescribing simple treatment without using an available voucher in DV can be another indicator of the presence of gatekeeping costs.

Doctors who Only Value Clinic Profits. A general issue in interpreting prescription and purchasing behavior, and the motivation for our experimental design, is that doctor and patient preferences are not observed. This makes it difficult to compare C and DV : a

doctor who changes her prescription from no to simple treatment from C to DV may do so because she preferred treatment all along, but was unable to sell it to the patient without the discount – or because her own preference changed based on the price change.

There is one exception, and this is the case of a doctor who intrinsically only values profits. This type of doctor has a fixed valuation of selling the patient severe vs. simple treatment, regardless of malaria probability or the price the patient pays: $V_v = V_v(\pi) > V_l(\pi, P) = V_l(\pi, 0) = V_l$. The doctor’s only restriction on malaria drug sales are patient preferences and gatekeeping costs – patients in N will not buy any treatment, and patients in $SL1$ will buy simple but not severe treatment; moreover, patients in $SL2$ will impose a gatekeeping cost on prescribing severe treatment. When comparing DV with C , the only change from the doctor’s perspective is that the constraint on sales that arises from patients’ willingness to pay for simple treatment is lifted. We have:

Prediction (6) For a revenue-maximizing doctor, per-patient revenue should be higher in DV than in C.

By contrast, the only difference when comparing DV with PV comes from the higher gatekeeping cost associated with prescribing severe treatment to $SL2_0$ patients under PV . Voucher use and per-patient revenue should be identical between both treatment arms if simple treatment generates the highest revenue. Otherwise, when severe is more profitable than simple, we have:

Prediction (7) For a revenue-maximizing doctor who is affected by gatekeeping costs, per-patient revenue should be higher in DV than in PV.

B Empirical Appendix

B.1 Additional Experimental Details

In addition to the doctor and patient voucher treatments, the experimental design included two other treatments designed to increase doctor and patient trust in RDTs. While accounting for these treatments has no impact on our main results, we describe them here in the interest of transparency.

Doctor Information (Across-Clinic Randomization). Half the clinics were randomly selected to receive the “Doctor Information” intervention. Clinics in this group received an enhanced refresher training that included the “basic information” referenced in the main text, plus an additional session on the diagnostic accuracy of RDTs. This training was

informed by our qualitative scoping work, which indicated that doctors had low levels of trust in RDTs and thought the tests were only capable of diagnosing malaria when parasite concentrations in the blood were very high. The session began by reviewing the sensitivity rate of the brand/make of RDTs used in clinics, per the most recent WHO quality assurance testing (World Health Organization (2015)). The trainer then introduced a validation study of the same brand/make of RDT conducted in Mali by a team of Malian researchers (see Djimde et al. (2016)). The trainees were shown a video in which one of the study’s principal investigators (a Malian M.D.-Ph.D.) described the results of the study. Key messages were: (1) Over 99 percent of true malaria blood samples tested RDT positive (the sensitivity of the test), (2) 73 percent of malaria negative blood samples tested negative (the specificity of the test) and (3) RDT sensitivity remained very high (89-92 percent) at low parasite loads (1-100 parasites/ μ L). The session closed by reviewing several other studies from sub-Saharan Africa and discussing why it is medically appropriate to refrain from prescribing ACTs to “suspect” malaria cases with a negative RDT.

Patient Information (Within Clinic Randomization). The patient information intervention was designed to improve patient and caregiver information about malaria treatment and diagnostic guidelines. The information was conveyed through a short narrative video, which depicted a mother taking her child to a clinic for a suspected malaria case. The video described the symptoms of malaria, emphasized that all suspected malaria cases should be confirmed with either an RDT or microscopy test, noted that RDTs should be available for free at the clinic, and described recommended treatment for simple and severe malaria. The video also showed a demonstration of an RDT test and described how to differentiate a positive versus negative test result. The main objective of the video was to inform patients about Mali’s official malaria treatment guidelines and give patients the information needed to request and verify the results of a malaria test if they so desired.

B.2 Doctor Surveys

In addition to the data analyzed in the main sections of this paper, we also collected data from health care providers at two points in time. First, we administered a post-training survey to doctors and other care providers who attended the refresher trainings that took place at the beginning of the study. The post-training survey tested providers’ knowledge of topics covered in the basic training (e.g. recommended malaria treatments, symptoms of severe malaria) and topics only covered in the extended “Doctor Information” treatment (e.g. sensitivity and specificity of RDTs). We also selected up to three care providers for a post-

intervention endline survey.³ In addition to topics covered in the post-training survey, the endline asked caregivers about perceived patient knowledge, demand for drugs, and personal preferences regarding malaria diagnosis and treatment.

B.3 Analysis Sample

In total, our enumerators logged 2753 clinic visits during the clinic survey. Our analysis sample includes patients/respondents who met the following criteria: consented to the survey (2 observations excluded), the patient was present at the clinic (0 observations excluded), the clinic visit was for an acute illness (neither preventive care nor follow-up visit for earlier treatment, 442 observations excluded), and the patient had at least one of the following symptoms: fever; chills and/or excessive sweating; nausea, vomiting or diarrhea; poor appetite, unwilling to eat or to breastfeed; headache; cough; weakness, fatigue, or reduced consciousness (31 observations excluded). In addition, we only include in the analysis those observations that satisfy the following: complete clinic intake interview (61 observations excluded), the name of the respondent from the intake interview was confirmed in the exit interview (5 observations excluded), and the respondent was available to continue with the clinic exit interview (157 observations excluded). This leaves us with a final clinic survey sample of N=2055.

B.4 Variable Construction

Administrative Records. To determine which patients received and redeemed a voucher, we asked intervention officers to keep notes on voucher delivery and redemption. When a patient received a coupon signed by the doctor, they went to the pharmacy with two copies of the coupon (original and copy). The pharmacist priced the prescribed ACT on both parts and countersigned each coupon, then gave the patient the ACT and the part of the coupon marked "copy". After completing the purchase of the other medicines prescribed in her prescription, the patient delivered the coupon to the intervention officer. At this stage, the intervention officer took notes of name and age of the patient, price of ACT, and the presence of signatures (to check validity). We merge these notes with the in-clinic survey by using name, age of the patient, clinic visit date, and name of the clinic.

- *Used voucher* - this variable is constructed by using records of vouchers redemption, and is equal to one if a patient or doctor voucher was redeemed.

³We always interviewed the head doctor at the CSCoM. Subject to the number and type of staff at a CSCoM we also randomly selected one other doctor and one other care provider (including nurses, health technicians, and midwives) for interview.

Selected Clinic Survey Variables. To construct malaria treatment variables, we recorded medications reported by individuals after the consultation at the clinic (during the exit interview). The respondent was asked to report all the medicines and equipment that were prescribed; we included a detailed list of medications (generics and brands) and equipment commonly prescribed at the clinics. We also allowed the enumerator to describe an item if it was not included in the list. We recoded items included in these descriptions and constructed dummy variables that indicate a medication or item used in a malaria treatment. In addition, we asked if the items were purchased, and which were the main reasons to not buy the item. We recoded the answers “free, donated” as "purchased".

- *Respondent Suspects Illness is Malaria (Pre-Consultation)* – equal to one if the respondent answered “malaria (uncomplicated, severe or unspecified)” to the question “What illness do you think you/the patient suffer(s) from?”
- *Duration of Illness in Days* – based on survey question “For how many days have you/has the patient had the illness?”. Top-coded at the 99th percentile.
- *Received Injection or IV* – equal to one if the respondent paid for one or more items that indicate the use of an injection or IV. This includes: fees paid to health workers to receive an injection, IV, perfusion set (épícrânién, épícranni), catheter, fluids via an IV infusion, perfusion, syringe, injection/perfusion, Ringer’s lactate solution, glucose serum, and saline serum.
- *Simple Malaria Treatment (Prescribed/Purchased)* – this variable was constructed from individuals’ reports of what medications were prescribed. First, we code the variable to one if the individual declares any of the following: ACT (brand/type not specified), specified ACT (Artekin, Artefan, Coartem, ACT for adolescents, ACT for children, ACT for adults, Malacur, Combiart, or Laritem), artemether+lumefantrine (we also set this variable equal to one if a voucher was used according to administrative records), amodiaquine (including Amoquin, Camoquin, Novaquin), artemether tablets (if tablet/injection was unspecified we assume tablet if ‘received injection or IV’ was equal to zero), artesunate tablets (if tablet/injection was unspecified we assume tablet if ‘received injection or IV’ was equal to zero), quinine tablets (if tablet/injection was unspecified we assume tablet if ‘received injection or IV’ was equal to zero), sulfadoxine/pyrimethamine (we also checked for the following combinations but all the observations were zero: artesunate+amodiaquine, artemether+amodiaquine, artemether+SP, artesunate+SP). Finally, this variable was set to zero if a severe malaria treatment was prescribed/purchased.

- *Severe Malaria Treatment (Prescribed/Purchased)* - dummy equal to one if an individual reports: quinine injection (if injection/tablet was unspecified we assume injection if ‘received injection or IV’ was equal to one), artemether injection (if injection/tablet was unspecified we assume injection if ‘received injection or IV’ was equal to one), or artesunate injection (if injection/tablet was unspecified we assume injection if ‘received injection or IV’ was equal to one). In addition, we set this variable to one if a monotherapy/quinine tablets and an ACT treatment were prescribed/purchased, as this is consistent with delivering monotherapies via suppository. Here, monotherapy/quinine includes quinine/artemether/artesunate, while ACT treatment is a dummy variable equal to one if an individual reported any of the following: unspecified ACT, specified ACT (constructed as above), artemether+lumefantrine. We also checked for artesunate+amodiaquine, artemether+amodiaquine, artesunate+SP, and artemether+SP, but all the observations were zero.
- *No Malaria Treatment (Prescribed/Purchased)* – is constructed as a dummy variable equal to one if an individual did not report a malaria treatment (simple or severe).
- *Expected Match, Malaria Positive (Prescribed/Purchased)* – This variable is equal to the predicted malaria risk (explained in main text) times a dummy variable equal to one if an individual purchased or was prescribed a severe or simple malaria treatment.
- *Expected Match, Malaria Negative (Prescribed/Purchased)* – This variable is equal to the predicted probability of no malaria (1-predicted malaria risk) times a dummy variable equal to one if an individual did not purchase or was not prescribed a severe or simple malaria treatment.
- *Overall Match (Prescribed/Purchased)* – is the sum of the two previous variables.
- *Simple/Severe Malaria Treatment and Used Voucher* – equal to one if the patient purchased a simple/severe malaria treatment and a voucher was used (according to administrative records), zero otherwise.
- *Simple/Severe Malaria Treatment, No Voucher* – equal to one if the patient purchased a simple/severe malaria treatment and a voucher was not used (according to administrative records), zero otherwise.
- *Purchased Antibiotics* – equal to one if a respondent reported the purchase of antibiotics. We included: Amoxicillin, Amoxicilline+Cla, Ampicilline, Cefadroxil, Cefixime, Ceftriaxone, Ciprofloxacin, Clamoxyl, Cotrimoxazol (Trimoprim), Diazole,

Erycin, Erythromycin, Flagyl, Gentamycin, Metronidazole, Oracefal, Oxacilline, Penicillin, Synozole, and unspecified antibiotics. We also checked for Amodix, Amoxitem, Augmentine, Azithromycin, Bactox, Binozyt, Cedrox, Oleandomycine, Uclaprim, and Unasyn but all the observations were zero.

- *Patient Referred to Hospital or Placed Under Observation* – equal to one if the respondent answered “yes” to the question “Were you/was the patient placed under observation at the CSCom?” or “Were you/was the patient sent to a CSRef or hospital?”.
- *Total Cost of Treatment (CFA)* - individuals were asked to report what total price they paid for the consultation and all treatments. We set this value equal to zero if the patient had no record of prescribed/purchased treatments or a bill, and we top-coded at the 99th percentile.
- *Clinic Revenues* – equal to total cost of treatment plus the amount reimbursed (based on administrative data) if an ACT voucher was redeemed.

Selected Home Survey Variables.

- *Household Size* – is the total number of household members, the sum of two questions “How many members has your household aged 14 years or younger?” and “How many members has your household aged 15 years or older?”
- *Share Household Under 15* – members aged 14 years or younger divided by total household size.
- *Share Household Members Working* – based on survey question “How many members of your household have a permanent job or own a steady business?” divided by total household size.
- *Monthly Income Per Capita* – each respondent was asked to estimate the total monthly income of her household, then we divided this amount by the household size. Top-coded at the 99th percentile.
- *Rental Value of Home* – based on survey question “How much rent does your household pay?” or “Could you estimate the rent you would pay if you rented this dwelling?” if the household owned the dwelling. We allowed for different rent periods, so we adjusted the amount to construct a monthly measure. The variable was divided by 12 if it was expressed in annual terms, or multiplied by 52/12 if it was weekly variable. Top-coded at the 99th percentile.

- *Mosquito Nets Per Capita* – based on survey question “How many mosquito nets do the people in your household own?” divided by total household size.
- *Taking All Purchased ACTs* – During the home survey, we asked if patients were taking the medications purchased at the clinic “Is the patient/Are you currently taking ‘name of medication’?”. This question was only recorded for medications coded as purchased during the clinic survey (a small share of medications given at a zero price were not coded as purchased due to enumerator error; this variable is missing for patients whose ACTs were coded this way). We constructed dummy variables equal to one if the patient was taking a purchased medication. To determine if a patient was currently taking an ACT, we created a dummy equal to one if a patient was taking at least one of the following medications (conditional on the purchase of an ACT): artemether+lumefantrine (tablet), Artefan, Artekin, Coartem, ACT for adolescents, ACT for adults, ACT for children, Malacur, Combiart, Laritem, or unspecified ACT.
- *Taking Purchased ACT for Simple Malaria* – constructed the same way as the previous variable but conditional on the purchase of a simple malaria treatment.
- *Positive RDT* – based on the enumerator’s report of the home-based RDT; “What was the RDT test result?” Equal to one if positive, zero if negative, missing if not taken or inconclusive.

Health Worker Post-Intervention Survey.

- *Malaria Prevalence: General Population* – this variable is the answer to the question “Consider an average day in November. In the general population (including those who do not visit a clinic and do not feel sick), out of 1000 people, how many have malaria on that day?” divided by 1000.
- *Malaria Prevalence: Clinic Patients* – we included the question “Assume you have 100 patients during this period. Among them, how many are children under 5?”, then “Among those X children, how many have malaria?” and “Among those (100-X) patients 5 and above, how many have malaria?”. This variable is the sum of the last two questions divided by 100.
- *Feels Pressure from Patients to Prescribe Unnecessary Medication* – this variable is equal to one if the health worker said yes to the question “Do you ever feel pressure from patients to prescribe certain medicines when you think they are not necessary?”

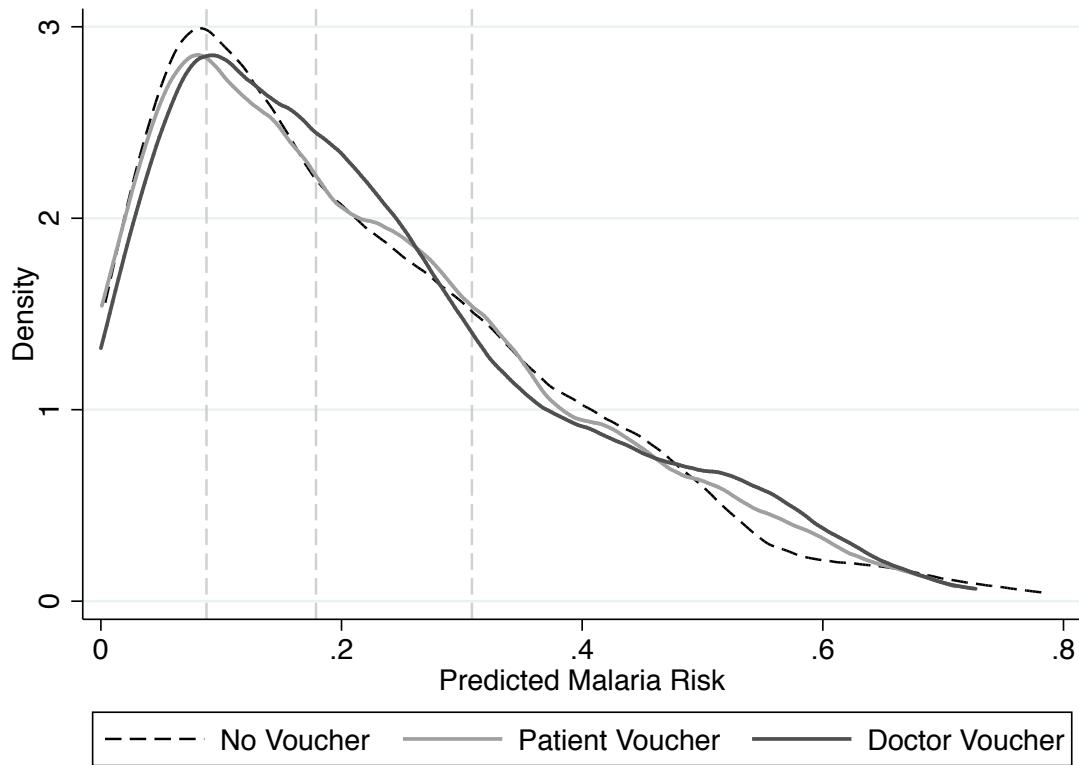
Appendix Figures and Tables

Table A1: Doctor's Expected Utility from Different Prescription Choices

	Utility from treatment use	Utility loss from gatekeeping cost
No treatment		
(1) C & DV	0	$-g \int_{SL_P} U_l(\pi, P) dF(\pi \gamma, \eta) - g \int_{SV_P} U_v(\pi) dF(\pi \gamma, \eta)$
(2) PV	0	$-g \int_{SL_0} U_l(\pi, 0) dF(\pi \gamma, \eta) - g \int_{SV_0} U_v(\pi) dF(\pi \gamma, \eta)$
Severe treatment		
(3) C & DV	$P(\pi \in SL_{2P} \cup SV_P) V_v(\hat{\pi})$	$-g \int_{SL_{1P}} U_l(\pi, P) dF(\pi \gamma, \eta) - g \int_{SL_{2P}} [U_l(\pi, P) - U_v(\pi)] dF(\pi \gamma, \eta)$
(4) PV	$P(\pi \in SL_{2_0} \cup SV_0) V_v(\hat{\pi})$	$-g \int_{SL_{1_0}} U_l(\pi, 0) dF(\pi \gamma, \eta) - g \int_{SL_{2_0}} [U_l(\pi, 0) - U_v(\pi)] dF(\pi \gamma, \eta)$
Simple treatment		
(5) C & (DV/no voucher)	$P(\pi \in SL_P \cup SV_P) V_l(\hat{\pi}, P)$	$-g \int_{SV_P} [U_v(\pi) - U_l(\pi, P)] dF(\pi \gamma, \eta)$
(6) (PV & DV)/voucher	$P(\pi \in SL_0 \cup SV_0) V_l(\hat{\pi}, 0)$	$-g \int_{SV_0} [U_v(\pi) - U_l(\pi, 0)] dF(\pi \gamma, \eta)$
(7) PV/no voucher	$P(\pi \in SL_P \cup SV_P) V_l(\hat{\pi}, P)$	$-g \int_{SL_{1_0} \setminus SL_{1P}} U_l(\pi, 0) dF(\pi \gamma, \eta) - g \int_{SL_{1P} \cup SL_{2_0}} [U_l(\pi, 0) - U_l(\pi, P)] dF(\pi \gamma, \eta)$ $-g \int_{SV_0} [U_v(\pi) - U_l(\pi, P)] dF(\pi \gamma, \eta)$
Patient choice between simple and severe treatment		
(8) C & (DV/no voucher)	$P(\pi \in SL_P) V_l(\hat{\pi}, P) + P(\pi \in SV_P) V_v(\hat{\pi})$	0
(9) C & (PV & DV)/voucher	$P(\pi \in SL_0) V_l(\hat{\pi}, 0) + P(\pi \in SV_0) V_v(\hat{\pi})$	0
(10) PV/no voucher	$P(\pi \in SL_P) V_l(\hat{\pi}, P) + P(\pi \in SV_P) V_v(\hat{\pi})$	$-g \int_{SL_0} U_l(\pi, 0) dF(\pi \gamma, \eta)$ $+g \int_{SL_P} U_l(\pi, P) dF(\pi \gamma, \eta) + g \int_{SV_P \setminus SV_0} U_v(\pi) dF(\pi \gamma, \eta)$

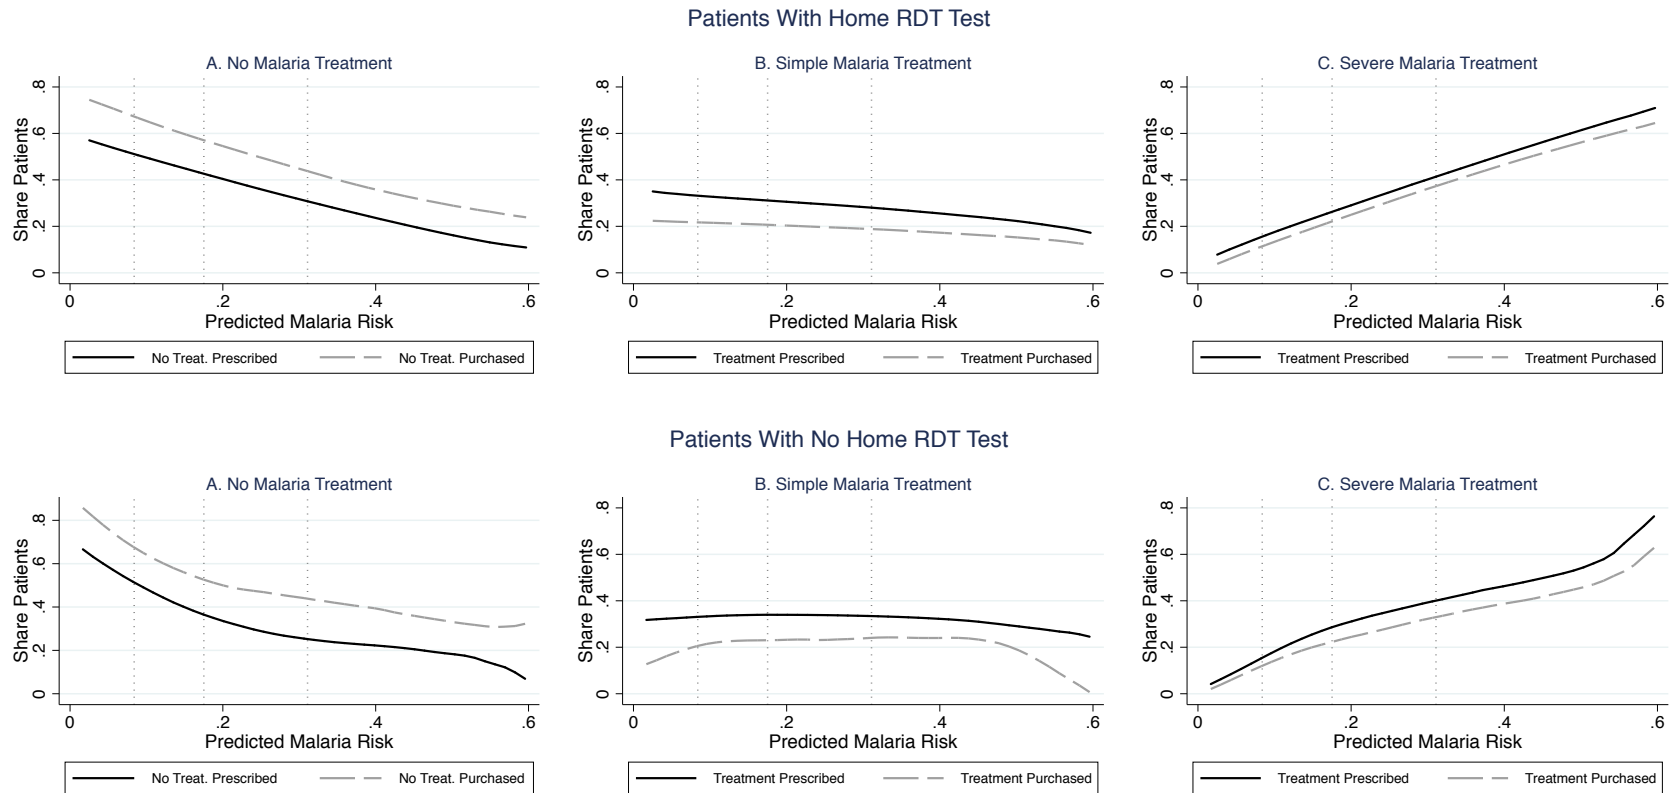
Notes: C denotes control, DV denotes voucher given to doctor, PV denotes voucher given to patient. "No voucher" and "voucher" indicate if the doctor offers the subsidy when prescribing simple treatment.

Figure B1: Distribution of Predicted Malaria Risk by Treatment Group



Notes: Kernel density estimates. Vertical dashed lines indicate 25th, 50th, and 75th percentiles of overall distribution respectively.

Figure B2: Treatment Outcomes by Predicted Malaria Risk in Control Group – By Home RDT Test Status



Notes: Results from local linear regressions. Regressions are run on the full sample, but graphs omit results for top and bottom 2.5 percent of malaria risk distribution to avoid influence of outliers. Vertical dashed lines indicate 25th, 50th, and 75th percentiles of predicted malaria risk respectively.

Table B1: Health Worker Beliefs from Post-Intervention Survey

	(1)	(2)	(3)
	Mean	SD	N
Malaria Prevalence: General Population	0.35	0.274	143
Malaria Prevalence: CCom Patients	0.482	0.194	143
Feels Pressure from Patients to Prescribe Unnecessary Medication	0.566	0.497	143
Feels Pressure: Antimalarials	0.519	0.503	81
Feels Pressure: Pain Killers	0.333	0.474	81
Feels Pressure: Antibiotics	0.21	0.41	81
Feels Pressure: Other Medicines	0.247	0.434	81

Notes: Results from post-intervention health worker survey. Sample includes doctors, nurses, and health technicians. A health worker is coded as feeling pressure to prescribe if s/he answers yes to the question: Do you ever feel pressure from patients to prescribe certain medicines when you think they are not necessary? Doctors answering yes were then asked to specify which medications. Antimalarial also includes quinine; pain killer includes analgesic, anti inflammatory, and sedative; antibiotic includes unspecified antibiotics and ciprofloxacin.

Table B2: Selection into Analysis Samples by Treatment

	(1)	(2)	(3)	(4)	(5)	(6)
		Regression Coefficients		P-Values		
	Control Mean	Patient Voucher	Doctor Voucher	Joint Test PV=DV	Joint Test PV=DV=0	N
<i>A. Whole Sample</i>						
Took Home Survey	0.734 [0.442]	-.008 (0.027)	-.011 (0.025)	0.94	0.889	2055
Took Home-Based RDT	0.551 [0.498]	-.001 (0.031)	-.007 (0.032)	0.861	0.972	2055
<i>B. Selected for Home Survey</i>						
Took Home Survey	0.86 [0.347]	-.004 (0.022)	0.008 (0.021)	0.558	0.829	1735
Took Home-Based RDT	0.646 [0.479]	0.003 (0.033)	0.007 (0.035)	0.889	0.978	1735
<i>C. Took Home Survey</i>						
Took Home-Based RDT	0.751 [0.433]	0.007 (0.033)	0.001 (0.035)	0.841	0.97	1495

*, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B3: Selection Into Home Survey and RDT Consent

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Whole Sample			Selected Home Survey			Took Home Survey		
	Mean: Not Selected	Diff: Selected	N	Mean: Survey Not Taken	Diff: Took Survey	N	Mean: Took RDT	Diff: Refused RDT	N
<i>A. Patient Characteristics</i>									
Number of symptoms	3.31 [1.63]	0.255 (0.158)	2055	3.48 [1.59]	0.09 (0.12)	1735	3.7 [1.57]	-.556*** (0.095)	1495
Fever	0.797 [0.403]	0.023 (0.031)	2055	0.838 [0.37]	-.021 (0.027)	1735	0.833 [0.373]	-.073** (0.029)	1495
Chills or Excessive Sweating	0.197 [0.398]	0.083** (0.034)	2055	0.25 [0.434]	0.034 (0.03)	1735	0.298 [0.457]	-.061* (0.035)	1495
Nausea, Vomiting, or Diarrhea	0.484 [0.501]	0.005 (0.037)	2055	0.429 [0.496]	0.07* (0.037)	1735	0.515 [0.5]	-.074* (0.037)	1495
Poor Appetite	0.444 [0.498]	0.038 (0.037)	2055	0.471 [0.5]	0.013 (0.038)	1735	0.495 [0.5]	-.05 (0.033)	1495
Headache	0.584 [0.494]	0.038 (0.048)	2055	0.579 [0.495]	0.05 (0.033)	1735	0.66 [0.474]	-.141*** (0.039)	1495
Cough	0.35 [0.478]	0.028 (0.028)	2055	0.425 [0.495]	-.055 (0.038)	1735	0.38 [0.485]	-.043 (0.029)	1495
Weakness/Fatigue	0.45 [0.498]	0.041 (0.037)	2055	0.492 [0.501]	-.001 (0.043)	1735	0.516 [0.5]	-.114*** (0.037)	1495
Duration of Illness in Days	4.09 [3.66]	0.272 (0.302)	2055	4.45 [4.58]	-.093 (0.31)	1735	4.34 [4.7]	0.033 (0.369)	1495
Age	15.9 [14.5]	1.75* (0.882)	2055	16.8 [15.7]	0.917 (0.906)	1735	18.5 [16.2]	-3.42*** (1.22)	1495
Patient Under 5 Years Old	0.325 [0.469]	-.04 (0.026)	2055	0.287 [0.454]	-.003 (0.029)	1735	0.246 [0.431]	0.174*** (0.035)	1495
Male (Patient)	0.434 [0.496]	-.014 (0.027)	2055	0.404 [0.492]	0.019 (0.031)	1735	0.408 [0.492]	0.067* (0.035)	1495
Patient is Pregnant	0.122 [0.328]	-.021 (0.022)	1139	0.12 [0.327]	-.022 (0.029)	967	0.103 [0.304]	-.021 (0.025)	834
Predicted Malaria Probability	0.205 [0.164]	0.013 (0.014)	2055	0.191 [0.147]	0.031*** (0.011)	1735	0.235 [0.164]	-.057*** (0.01)	1495
Purchased Malaria Treatment	0.522 [0.5]	0.031 (0.037)	2053	0.512 [0.501]	0.047 (0.037)	1735	0.577 [0.494]	-.075* (0.039)	1495
<i>B. Household Characteristics</i>									
Patient Answered Clinic Survey	0.466 [0.5]	0.006 (0.034)	2055	0.496 [0.501]	-.028 (0.032)	1735	0.496 [0.5]	-.123*** (0.031)	1495
Male	0.269 [0.444]	0.012 (0.029)	2055	0.237 [0.426]	0.05 (0.031)	1735	0.286 [0.452]	0.009 (0.034)	1495
Bambara	0.3 [0.459]	0.102*** (0.028)	2053	0.392 [0.489]	0.012 (0.03)	1733	0.412 [0.492]	-.04 (0.035)	1493
Speaks French	0.491 [0.501]	0.03 (0.038)	2055	0.521 [0.501]	0 (0.031)	1735	0.528 [0.499]	-.032 (0.038)	1495
Literate (in French)	0.234 [0.424]	0.026 (0.026)	2055	0.321 [0.468]	-.07** (0.027)	1735	0.247 [0.431]	0.017 (0.035)	1495
Primary School or Less	0.475 [0.5]	-.03 (0.036)	2055	0.433 [0.497]	0.013 (0.027)	1735	0.447 [0.497]	0.001 (0.037)	1495
Household Size ⁺							10.6 [8.19]	0.173 (0.669)	1491
Share HH Under 15 ⁺							0.417 [0.191]	0.015 (0.016)	1485
Share HH Members Working ⁺							0.252 [0.188]	0.02 (0.014)	1485
Monthly income per capita ⁺							19000 [21000]	4917** (1964)	1432
Rental Value of Home ⁺							57000 [77000]	21000*** (6885)	1469
Mosquito Nets Per Capita ⁺							0.481 [0.31]	0.018 (0.026)	1482

⁺ indicates that variable was recorded in the home survey only. Variables measured in CFA and duration of illness top-coded at the 99th percentile. CFA610 \approx USD1. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B4: Predicting RDT Positivity With Observables

	(1) RDT Positive
Fever	0.442*** (0.170)
Chills or Excessive Sweating	0.198* (0.105)
Nausea, Vomiting, or Diarrhea	0.382*** (0.0955)
Reduced Appetite	0.00968 (0.0987)
Headache	0.238** (0.120)
Cough	-0.185** (0.0794)
Weakness, Fatigue, or Reduced Consciousness	0.125 (0.0979)
Duration of Illness in Days	-0.0189** (0.00904)
Age Patient	-0.00438 (0.00535)
Patient Under 5 Years Old	-1.473*** (0.236)
Under 5 × Age	0.266*** (0.0988)
Patient is Male	1.030** (0.414)
Patient is Pregnant	-0.357* (0.201)
Ethnic group: Bambara	0.153* (0.0865)
Respondent Speaks French	-0.219 (0.134)
Respondent is Literate in French	-0.454*** (0.145)
Respondent Has Primary Education or Less	-0.123 (0.119)
Patient Answered Clinic Survey	-0.383** (0.165)
Pseudo R-Squared	0.145
N	1126

Notes: Standard errors clustered at the CSCoM level in parentheses. Respondent refers to individual who answered clinic survey. ***, **, and * indicate significance at the 1, 5, and 10 percent significance levels respectively.

Table B5: Impacts on Malaria Treatment Outcomes, No Additional Controls

	(1)	(2)	(3)	(4)	(5)
		Any Malaria Treatment		Severe Malaria Treatment	
	Used Voucher	Prescribed	Purchased	Prescribed	Purchased
β_P : Patient Voucher	0.35*** (0.031)	0.067** (0.027)	0.15*** (0.031)	-0.033* (0.019)	-0.011 (0.021)
β_D : Doctor Voucher	0.27*** (0.025)	0.023 (0.026)	0.092*** (0.027)	-0.058** (0.023)	-0.051** (0.020)
<i>P-values: Theory-Driven Tests</i>					
$\beta_P = \beta_D$: Two-sided	0.015**	0.067*	0.027**	0.369	0.104
$\beta_P = \beta_D$: One-sided	0.008***	0.034**	0.014**	0.815	0.948
One-sided test of:	GC	PD	PD	DD	DD
Mean (Control)	0	0.62	0.46	0.30	0.26
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. GC, PD, and DD indicate tests of gatekeeping costs, patient-driven, and doctor-driven demand respectively. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B6: Impacts on Malaria Treatment Outcomes - Heterogeneity by Predicted Malaria Risk, No Additional Controls

	(1)	(2)	(3)	(4)	(5)
		Any Malaria Treatment		Severe Malaria Treatment	
	Used Voucher	Prescribed	Purchased	Prescribed	Purchased
δ_{PH} : Patient Voucher \times High Risk	0.34*** (0.041)	0.038 (0.037)	0.12*** (0.040)	-0.056* (0.033)	-0.021 (0.033)
δ_{DH} : Doctor Voucher \times High Risk	0.33*** (0.039)	0.029 (0.033)	0.088** (0.037)	-0.13*** (0.034)	-0.12*** (0.031)
δ_{PL} : Patient Voucher \times Low Risk	0.36*** (0.042)	0.095** (0.038)	0.18*** (0.043)	-0.013 (0.026)	-0.0063 (0.027)
δ_{DL} : Doctor Voucher \times Low Risk	0.21*** (0.032)	0.0067 (0.040)	0.087** (0.036)	0.00089 (0.030)	0.0096 (0.025)
θ : High Malaria Risk	0.0078 (0.028)	0.26*** (0.041)	0.26*** (0.042)	0.28*** (0.038)	0.26*** (0.036)
<i>P-values: Theory-Driven Tests</i>					
$\delta_{PH} = \delta_{DH}$: Two-sided	0.719	0.810	0.400	0.085*	0.008***
$\delta_{PH} = \delta_{DH}$: One-sided	0.360	0.405	0.200	0.957	0.996
One-sided test of:	GC/DD	–	–	DD	DD
$\delta_{PL} = \delta_{DL}$: Two-sided	0.000***	0.021**	0.011**	0.670	0.594
$\delta_{PL} = \delta_{DL}$: One-sided	0.000***	0.010**	0.006***	0.335	0.297
One-sided test of:	GC/PD	PD	PD	–	–
Mean (Control, Low Risk)
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. Standard errors based GC, PD, and DD indicate tests of gatekeeping costs, patient-driven, and doctor-driven demand respectively. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B7: Use of Voucher for Purchased Malaria Treatment

	(1)	(2)	(3)	(4)
	Severe Malaria Treat- ment and Used Voucher	Severe Malaria Treat- ment, No Voucher	Simple Malaria Treat- ment and Used Voucher	Simple Malaria Treat- ment, No Voucher
β_P : Patient Voucher	0.042*** (0.010)	-0.065*** (0.020)	0.31*** (0.029)	-0.14*** (0.021)
β_D : Doctor Voucher	0.0097* (0.0054)	-0.062*** (0.020)	0.26*** (0.024)	-0.11*** (0.021)
<i>P-values: Theory-Driven Tests</i>				
$\beta_P = \beta_D$: Two-sided	0.004***	0.886	0.119	0.061*
$\beta_P = \beta_D$: One-sided	0.002***	0.443	0.059*	0.030**
One-sided test of:	PD	DD	PD	PD
Mean (Control)	0	0.26	0	0.21
N	2053	2053	2053	2053

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. DD and PD indicates a test of doctor and patient-driven demand respectively. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B8: Impacts of Patient Information on Malaria Testing at the Clinic

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		All Patients			If Prescribed Antimalarial		
	Prescribed Antimalarial	Any Malaria Test	RDT Test	Microscopy Test	Any Malaria Test	RDT Test	Microscopy Test
<i>Panel A. Overall Effects</i>							
Patient Information	-0.015 (0.024)	-0.068*** (0.020)	-0.037* (0.021)	-0.015 (0.021)	-0.059** (0.027)	-0.015 (0.026)	-0.020 (0.028)
Mean (No PI)	.653	.571	.305	.265	.643	.295	.334
<i>Panel B. By Voucher Treatment Group</i>							
Patient Information	0.016 (0.036)	-0.072* (0.037)	-0.074** (0.029)	-0.0086 (0.033)	-0.068 (0.047)	-0.078** (0.036)	-0.041 (0.046)
Patient Information × Doctor Voucher	-0.023 (0.044)	0.030 (0.056)	0.052 (0.048)	0.016 (0.046)	0.014 (0.073)	0.10 (0.063)	0.0060 (0.063)
Patient Information × Patient Voucher	-0.074 (0.052)	-0.034 (0.048)	0.028 (0.041)	-0.061 (0.044)	0.0061 (0.059)	0.078 (0.057)	0.014 (0.053)
Doctor Voucher	0.026 (0.029)	0.013 (0.039)	0.0096 (0.034)	-0.0032 (0.031)	0.021 (0.053)	0.0013 (0.053)	-0.0046 (0.040)
Patient Voucher	0.094** (0.038)	0.078* (0.040)	-0.014 (0.039)	0.088** (0.039)	0.057 (0.043)	-0.014 (0.055)	0.029 (0.045)
<i>P-Values</i>							
Patient Voucher=Doctor Voucher	0.036**	0.178	0.525	0.010***	0.473	0.730	0.379
Mean (No PI, No Voucher)	0.60	0.52	0.30	0.23	0.60	0.30	0.30
N	2053	2055	2055	2055	1342	1342	1342

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B9: Impacts on Clinic Revenues and Patient Costs (CFA)

	(1) Clinic Revenues	(2) Patient Costs
<i>Panel A. No Interactions</i>		
β_P : Patient Voucher	-147.8 (326.3)	-511.6 (331.7)
β_D : Doctor Voucher	-207.0 (237.6)	-531.5** (238.8)
<i>P-Values: Two-Sided Tests</i>		
$\beta_P = \beta_D$	0.801	0.934
$\beta_P = \beta_D = 0$	0.669	0.084*
<i>Panel B. Interactions with Predicted Malaria Probability</i>		
δ_{PH} : Patient Voucher \times High Risk	-268.0 (488.1)	-633.1 (497.3)
δ_{DH} : Doctor Voucher \times High Risk	-734.6* (378.7)	-1185.6*** (379.0)
δ_{PL} : Patient Voucher \times Low Risk	-11.8 (368.3)	-405.1 (372.2)
δ_{DL} : Doctor Voucher \times Low Risk	384.6 (320.3)	166.3 (324.4)
θ : High Malaria Risk	453.7 (463.8)	492.2 (478.6)
<i>P-values: Two-Sided Tests</i>		
$\delta_{PH} = \delta_{DH}$	0.210	0.147
$\delta_{PL} = \delta_{DL}$	0.267	0.117
Mean (Control)	5098.9	5098.4
N	1864	1864

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. In Panel B, standard errors are based on 1,000 bootstrap replications, with re-sampling at the clinic level. Predicted malaria risk is re-calculated on each bootstrap replication. All variables measured in CFA top-coded at the 99th percentile. CFA610 \approx USD1. Malaria cases classified based on doctor prescriptions. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B10: Impacts on Match Between Treatment and Illness - RDT Sub-Sample

	(1)	(2)	(3)	(4)
	Expected Match		Actual Match	
	Prescribed	Purchased	Prescribed	Purchased
β_P : Patient Voucher	-0.041*	-0.090***	-0.030	-0.079**
	(0.024)	(0.022)	(0.033)	(0.032)
β_D : Doctor Voucher	-0.038*	-0.070***	-0.015	-0.096***
	(0.021)	(0.022)	(0.033)	(0.035)
<i>P-values: Two-Sided Tests</i>				
$\beta_P = \beta_D$	0.858	0.357	0.634	0.663
$\beta_P = \beta_D = 0$	0.158	0.000***	0.657	0.008***
Mean (Control)	0.48	0.56	0.51	0.61
N	1126	1126	1126	1126

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. In columns 3 and 4 match quality is equal to 1 if an individual is malaria positive and was prescribed/bought an antimalarial or is malaria negative and was not prescribed/did not buy an antimalarial and is zero otherwise. In columns 1-2 the value of one is replaced with either the probability an individual is positive (for antimalarial receipt) or the probability an individual is negative (for non-receipt). *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.

Table B11: Share of Patients Taking An ACT at Home Survey

	(1) All Prescribed ACTs	(2) Prescribed ACT for Simple Malaria
β_P : Patient Voucher	-0.028 (0.039)	0.0037 (0.036)
β_D : Doctor Voucher	-0.065 (0.048)	0.025 (0.031)
<i>P-values: Two-Sided Tests</i>		
$\beta_P = \beta_D$	0.405	0.562
$\beta_P = \beta_D = 0$	0.408	0.709
Mean (Control)	0.92	0.94
N	460	346

Notes: Robust standard errors clustered at the level in parentheses. All regressions control for clinic visit date fixed effects. We use double selection lasso to choose additional controls. See notes to Table 3 for a list of potential controls. The first column is limited to individuals who purchased an ACT treatment at the CSCoM as part of either simple or severe malaria treatment. The second column is limited to individuals who purchased an ACT as part of simple malaria treatment. *, **, and *** denote statistical significance at the 10, 5, and 1 percent levels respectively.