

Does Patient Demand Contribute to the Overuse of Prescription Drugs?

Carolina Lopez
Anja Sautmann
Simone Schaner



WORLD BANK GROUP

Development Economics
Development Research Group
November 2020

Abstract

This study conducted an experiment in Mali to test whether patients pressure doctors to prescribe medical treatment they do not necessarily need. The experiment varied patients' information about a discount for antimalarial tablets and measured demand for both tablets and costlier antimalarial injections. The study finds 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. Providers did not use the information advantage to sell injections—their use fell in both information conditions.

This paper is a product of the Development Research Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://www.worldbank.org/prwp>. The authors may be contacted at asautmann@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Does Patient Demand Contribute to the Overuse of Prescription Drugs?

Carolina Lopez and Anja Sautmann
and Simone Schaner*

JEL Codes: I12, I18, O12

Keywords: malaria treatment, demand for prescription drugs, health care overuse, doctor-patient interaction, Mali

*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, the Stanford Rosenkranz Global Health Policy Symposium, McGill, Brown, Microsoft Research, FGV Sao Paulo, PUC Rio, the University of Cape Town, the University of Gothenburg, and the Stockholm School of Economics 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 Facult de Mdecine de Pharmacie et d'Odontostomatologie de Bamako (FMPOS) (protocol No. 2016/47/CE/FMPOS) and registered in the AEA RCT registry (AEARCTR-0001707). The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. All errors are our own. Carolina Lopez: Brown University, carolina_lopez@brown.edu. Anja Sautmann: World Bank, asautmann@worldbank.org. Simone Schaner: University of Southern California, BREAD, and NBER, schaner@usc.edu.

Low-income countries must address two opposing challenges when crafting health policy: many patients do not seek care when they are ill (Whitehead, Dahlgren and Evans, 2001; Sautmann, Brown and Dean, 2017), but those who do often consume a great deal of unnecessary treatment (Das and Hammer, 2014; Li et al., 2012). One solution could be to subsidize critical treatments in the public sector, while staffing facilities with trained health-care providers to act as gatekeepers to subsidized care.

Yet providers 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 document large shares (36-81 percent) of malaria-negative patients receiving antimalarials at public health centers (Reyburn et al., 2004; Hamer et al., 2007; Bisoffi et al., 2009; Ansah et al., 2010). Providers' gatekeeping ability is key for healthcare policies that increase demand, such as subsidies: only if healthcare providers are good gatekeepers will such policies expand access to those who need treatment while limiting overuse among those who do not need treatment.

Why are so many providers such poor gatekeepers? Most existing literature studies *doctor-driven* reasons for overtreatment. These include misaligned financial incentives (e.g. Iizuka, 2012; Currie, Lin and Meng, 2014; Das et al., 2016) and knowledge gaps, compounded by low motivation (e.g. Banerjee, Duflo and Glennerster, 2008; Das, Hammer and Leonard, 2008).¹ Doctors and other healthcare providers 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 provider's recommendation. But by a similar argument, there is scope for *patient-driven* demand for unnecessary treatment: if the provider cannot persuade the patient that a medication is not needed, she may write a prescription just to keep the patient satisfied. Anecdotal evidence suggests patient demand matters: for example, in 2016 the U.S. pharmaceutical industry spent \$6 billion on direct-to-consumer marketing of prescription drugs, 22 percent of total marketing spending (Schwartz and Woloshin, 2019), and explicit or perceived patient demand is often cited by providers as a reason for overtreatment (Kotwani et al., 2010; Linder et al., 2014). Yet there has been much less rigorous research on this channel.

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

Our paper aims to fill this gap. We begin by developing a theoretical model of doctor-patient interactions, where doctors and patients may have different preferences about the best course of medical treatment for an illness, and may in particular disagree about whether a certain treatment is truly needed. *Patient-driven demand* occurs when the patient consumes more treatment than the doctor would like. In our model, this can happen when doctors change their prescription behavior to avoid “gatekeeping costs”. Patients inflict this utility cost (e.g., in time, hassle, or lost future business) on doctors when they learn that the doctor’s prescription does not allow them to obtain their preferred treatment option.² We then conduct an experiment focused on malaria treatment. Within the experiment, the model generates predictions for when and how a provider will selectively share information with patients to reduce gatekeeping costs and steer treatment towards her own preferences. The theory guides our analysis, allowing us to (1) test for the presence of gatekeeping costs, (2) assess to what extent patients drive increased use of antimalarials in response to a price subsidy, and (3) assess the extent to which providers drive demand for more expensive antimalarial injections. Finally, we estimate how these changes in demand impact the misallocation of treatment.

The experiment was conducted with 2,055 patients visiting 60 public-sector health clinics in Bamako, Mali. Mali’s policies for malaria control are in line with international consensus and dictate that patients should only be prescribed an antimalarial if they have a positive malaria test (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. We document widespread overuse of antimalarials even though diagnostic tests are readily available: in our control group, 58 percent of malaria-negative 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, by giving a voucher for a free course of ACT tablets. This subsidy increased the

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

attractiveness of simple malaria treatment relative to both no treatment and severe treatment. The experiment varied (within clinic, across different days) whether this discount was a provider’s private information or known to both provider and patient. On “Patient Voucher” days, all acutely ill patients visiting the health clinic were given a voucher. On “Doctor Voucher” days, the same vouchers were given to providers to dispense at their discretion. This means the healthcare provider could *choose* to make simple treatment more attractive to the patient by revealing the voucher.³

Our model shows that treatment outcomes will not vary across the Doctor and Patient Voucher conditions *unless there are gatekeeping costs*. Intuitively, the patient-facing costs and clinic revenues associated with different treatments are the same in the Doctor and Patient Voucher conditions; consequently a doctor (or other healthcare provider’s) preferred action based on these factors will also be the same in both conditions. The doctor will sometimes deviate from her preferred prescription to satisfy the patient and avoid gatekeeping costs. However, relative to the Patient Voucher condition, the doctor’s information advantage in the Doctor Voucher condition may help her avoid making such compromises. 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 the case where patients prefer more simple malaria treatment than doctors (“patient-driven demand”).

By contrast, suppose the provider would like to induce as much treatment as possible, for example because she wants 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 the case when providers have a stronger preference for severe malaria treatment than patients (“doctor-driven demand”). Note that our primary tests of doctor- and patient-driven demand rely on different margins (simple versus no malaria treatment for patient-driven demand; simple versus severe malaria treatment for doctor-driven demand). It is therefore *ex ante* possible to find support for one, both, or neither type of misalignment.

³Vouchers in both conditions had identical appearances, terms, and conditions. We 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.

We find robust evidence that gatekeeping costs affect providers' 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 provider and patient preferences are frequently misaligned. We also find consistent evidence that excess demand for vouchers on Patient Voucher days is driven by *patients' preferences for simple treatment*. First, providers 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. Second, the excess demand for antimalarials on Patient Voucher days is driven by patients with the fewest malaria symptoms – antimalarial purchases by this group are 9 percentage points (23 percent) higher. This suggests providers prefer to withhold treatment from these marginal low-risk cases, but accede to demand when patients know that treatment is free.

In contrast, we find no evidence of doctor-driven demand for severe malaria treatment. There are no significant differences in rates of severe malaria treatment across Patient and Doctor Voucher days, and we find no evidence that providers use their information advantage in the Doctor Voucher condition to increase clinic revenues.⁴

Next, we assess the match between underlying illness and treatment with data from follow-up home-based malaria tests with patients. In the Control group, we estimate that 57 percent of patients received “appropriate” malaria treatment, in that they either were 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, compared to only 4 percentage points in the Doctor Voucher condition.⁵ Over half of the worsening of the match in the Patient Voucher condition is driven by a change in providers' prescribing behavior that is not observed in the Doctor Voucher condition; in other words, patient-driven demand for ACTs leads to overtreatment.⁶

⁴Although doctors in our setting are salaried, they have an incentive to increase revenue because this funds their salaries, and bonuses are often tied to revenues. See Section 1 for more detail.

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

⁶The remainder is driven by an increase in antimalarial purchases among individuals who would have received prescriptions even absent the price discount (some patients do not purchase prescribed medication, which in our context works to improve the match, since most patients do not actually have malaria).

Our experiment provides 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, Lin and Meng (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, Lin and Meng 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. Our contribution is to develop a method to identify whether this provider compliance with patient requests for medication is willing (the doctor prefers to prescribe) or reluctant (the doctor would rather not prescribe, but goes along to avoid gatekeeping costs). Our experiment also has the advantage that it uses real patients, allowing for natural doctor-patient interactions where patients themselves choose whether to express demand for treatment in response to the subsidy.

Our work also contributes to the 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, Dupas and Schaner (2015) find that subsidizing anti-malarials at private-sector “drug shops” dramatically improves access to life-saving medication, but making the subsidy too high generates large amounts of overtreatment with only marginal additional 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, Brown and Dean (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.

Our tests are specific to marginal changes in patient demand in response to the ACT subsidy and do not allow us to extrapolate how the different

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

mechanisms contribute to overall levels of overtreatment. Even so, public clinics are a prime delivery mechanism for acute malaria care across Sub-Saharan Africa, meaning this is precisely the margin of interest for policymakers deciding whether to subsidize malaria treatment. On this margin, we show subsidies induce patient-driven demand that leads to overtreatment. Thus, interventions that strengthen the provider’s authority and reduce gate-keeping costs (like patient communication tools) could help sustain subsidies while limiting overtreatment.

The rest of the paper is structured as follows: in the next section we provide background on health care and malaria treatment in Mali and present some stylized features of status quo malaria treatment. We develop a model of doctor-patient interactions in section 2. Section 3 describes our experimental design, section 4 presents and discusses our empirical results, including alternative explanations for our findings, and section 5 concludes.

1 Background

1.1 Health Care and Malaria in Mali

The public health system in Mali is organized around decentralized community-based primary care funded by user fees. At the foundation of this system are *Centres de Sant Communautaire* (CSCComs) – community-based primary care clinics managed by a local health association. A typical clinic in our sample is staffed with one or two physicians and around five medical trainees and five other staff who can prescribe malaria, along with additional technical and administrative staff. Clinics with a laboratory have a lab technician who conducts blood tests. Table B1 in the Appendix gives an overview of the staffing numbers in our sample. The local health association retains revenues from sales of medications and other user fees, which are used to fund the operations of the CSCCom. Qualitative research suggests this system leads staff to prioritize profit generation, since clinic revenues can be used to fund end-of-year physician bonuses (Touré, 2015; Escot, 2012). Moreover, pharmacy sales often directly fund staff salaries – this is the case in 87 percent of the clinics in our sample.

Similar community-funded public healthcare is found all over West Africa, and in other countries like China, where high rates of antibiotic use in the public sector have been attributed to physicians’ incentives to generate bonus

revenue (Currie, Lin and Meng, 2014). This structure – where medical staff are salaried but still have an 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.

CSCComs are one of the most important sources of care in Mali: 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 CSCCom. The quality of care across CSCComs 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)). All clinics in our sample have at least one salaried physician on staff, and in over 90 percent of clinics that have a pharmacy, there is a physician present when the pharmacy is open (see Appendix Table B1).

One of the most commonly-treated illnesses in CSCComs is malaria. Despite recent progress, malaria is the second largest cause of mortality in Mali, accounting for roughly 14 percent of all deaths and 19 percent of deaths among children under five (IHME, 2018). 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 complications, 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

⁸This recommendation was first made in the five-year strategic plan for 2007-2011.

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 These policies are in line with WHO recommendations (WHO, 2014) and are designed to limit the spread of drug resistance, which is hastened by overtreatment. Limiting overuse of antimalarials is particularly important, as drug resistance has rendered past generations of antimalarials ineffective across much of Sub-Saharan Africa and Asia (Arrow, Panosian and Gelband, 2004). This externality is one of the largest (potential) costs to overtreatment, though overtreatment also strains government budgets for drug subsidies, wastes resources of poor families, causes side effects, and can slow diagnosis and treatment of non-malarial illnesses.

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 (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.

1.2 Status Quo Malaria Treatment in the Study Sample

We use data from our Control group to quantify overuse of malaria treatment at the 60 clinics 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 3. We stationed surveyors at each clinic for

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.

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 provider 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.¹¹

Table 1 uses the 628 patient observations in the Control group to give an overview of how acute illnesses are treated in clinics in Bamako. Panel A shows average indicators regarding malaria beliefs and treatment outcomes. Panel B reports treatments that often accompany malaria medications: antibiotics, which are often given for fevers but are not indicated for malaria, as well as injections/IV, which are used in severe malaria treatment. Panel B also shows how often patients were referred to a hospital, the recommended course of action for severe malaria. Finally, Panel C shows average costs in West African CFA (at the time of the experiment CFA 610 was approximately USD 1), split by the type of treatment received.

Looking at means in the whole sample of 628 (column 2), 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 (Panel A).¹² 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

¹¹RDTs detect antigens for malaria, which remain in the bloodstream after the infection is cured. 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 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 occurred less than 1 percent of the time, and the invalid test rate was 0 percent (WHO, 2017).

¹²We classify treatment as indicating severe malaria treatment if the patient received artemisinin monotherapy or quinine, (1) as an injectable or (2) recorded in the survey as tablets, *and* received an ACT. These cases are consistent with guidelines for severe malaria treatment which require administering monotherapy/quinine into the bloodstream or rectally.

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 (Panel C).

We quantify overtreatment among the subset of patients who received a valid RDT result 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, even though 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. Finally, only 56 percent of Control group patients prescribed an antimalarial took a test, even though all clinics in our study had malaria testing capability. These facts suggest gatekeeping failures, especially given the very straightforward treatment guidelines for malaria (prescribe only for patients with a positive test).

There are some indications that overprescription might be driven by providers. The patterns of severe malaria treatment we observe match anecdotal reports that doctors in our sample treat positive malaria tests 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.¹³ The revenue especially from severe treatment provides an additional motive if doctors are concerned with the clinic’s financial health.¹⁴

Moreover, 25 percent of patients do not purchase their antimalarial pre-

¹³We find indications for this in our sample as well. Appendix Figure B1 shows that the match between malaria status in our at-home test and receipt of an antimalarial prescription is better among patients who received an RDT at the clinic as compared to microscopy. Specifically, we find higher rates of overtreatment after microscope-based tests, which suggests clinics may see more false positives with microscopy.

¹⁴While we do not observe profits, higher revenue from an individual patient will indicate higher profits provided profit is positively correlated with revenue. This likely holds in our case, especially since the higher cost of severe treatment includes service fees e.g. for injections. Our estimates suggest service fees are 35-47 percent higher for severe (as compared to simple) malaria. Since there was very limited patient queuing in our setting, the short-run marginal cost of providing such services is close to zero.

scription, with the prescription-purchase gap mostly driven by individuals declining to purchase available treatment (77 percent of cases) rather than stock outs.¹⁵ 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. First, patients over-estimate their malaria risk: 55 percent of respondents report suspecting malaria before they see the provider; among those with a valid home-based RDT, this proportion is 61 percent, even though only 25 percent actually test positive. Second, the share of patients prescribed treatment exceeds providers’ average estimate of true malaria prevalence (48 percent) by 13 percentage points (see Appendix Table B2). Finally, 57 percent of health workers report feeling pressure from patients to prescribe unnecessary medication, with over half mentioning antimalarials specifically.¹⁶

2 Theory

In order to investigate the role of patients pressuring health workers to prescribe, we use a simple model of a doctor-patient interaction 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 and doctor’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. In the second condition (“Doctor Voucher” or “*DV*”), only the doctor knows about the discount going into the consultation, and she can decide whether to reveal this information to the patient. In section 3, we explain how we operationalized these pricing treatments. The remainder of this section introduces the model and outlines the key testable predictions that identify patient- and doctor-driven demand. We leave the full model and formal analysis for Appendix A.

¹⁵Our 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.

¹⁶Appendix B gives additional detail on the health worker survey.

2.1 Model Setup

Malaria Risk and Patient Preferences. Patient and doctor observe the patient’s symptoms and assign probability π that the patient has malaria.¹⁷ Patients receive different utilities from taking either simple malaria treatment ($t = L$) or severe malaria treatment ($t = H$). The expected utility of $t \in \{L, H\}$ is the expected benefit minus the cost:

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

where B_t denotes the benefit from treatment t to a malaria-positive patient, and P_t is the price of treatment to the patient, with $P_L \in \{0, P\}$, depending on whether the patient has a voucher. The full cost of treatment, $C_t(P_t)$, includes both monetary and non-monetary costs such as side effects. We assume $C_L(0) < C_L(P) < C_H(P_H)$. We normalize the expected utility of no treatment ($t = N$) to zero.

Figure 1, Panel A illustrates patient preferences with and without the simple malaria discount. Assuming that $\frac{C_L(\cdot)}{B_L} < \frac{C_H}{B_H}$, there are three regions of the malaria probability π from low to high: individuals in N prefer no treatment; patients in L prefer simple treatment; and those in H prefer severe treatment over simple or no treatment. This captures that the benefits of both simple and severe treatment – and thus the patient’s willingness to buy them – are increasing in the severity of symptoms, but that the benefits of severe treatment only outweigh the higher costs when symptoms are most pronounced. Panel A also shows that when the voucher lowers the price P_L from P to 0, L expands, so that more patients prefer simple treatment over both no and severe treatment.

Note that there is an area within L , denoted $L2$, where patients prefer simple treatment but would be willing to purchase severe (whereas in $L1$ they would rather get no treatment at all). This means doctors could “oversell” severe treatment to at least some patients.

Doctor Utility. Doctors may disagree with patients’ preferences. This is expressed by a wedge Δ_t , making the doctor’s direct utility from the patient taking treatment t

$$V_t(\pi, P_t) = U_t(\pi, P_t) + \Delta_t.$$

¹⁷The full model in the appendix allows for private signals about malaria risk; see also below.

The wedge may be positive, for example if the doctor does not care about financial costs paid by the patient, or if the doctor’s utility factors in clinic profits. A positive Δ_t also captures non-financial reasons for preferring to treat, such as risk aversion or fear of missing a diagnosis. On the other hand, Δ_t could be negative, meaning the doctor gets less utility from treatment t than the patient. This may reflect internalized costs of public provision of health care, anticipated side effects, concerns about creating disease resistance, or discomfort with disregarding official treatment policy.

From low to high π , preferences V_t generate regions \hat{N} , \hat{L} , and \hat{V} where doctors prefer no, simple, and severe treatment, respectively (assuming $\frac{C_L}{B_L + \Delta_L} < \frac{C_H}{B_H + \Delta_H}$). Unless doctors are completely altruistic and $\Delta_L = \Delta_H = 0$, these regions will be shifted relative to N , L , and V . When Δ_t is positive, the threshold for treatment t shifts to the left, and doctors prefer more of this treatment than patients. When Δ_t is negative, the threshold for treatment t shifts to the right and doctors have a lower preference for t than patients.¹⁸

The cases of disagreement are our object of interest. In what follows, we say patients drive (marginal) 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. The extent to which this drives misallocation will depend on whether doctor versus patient preferences are better aligned with the social planner. We can also assess misallocation empirically.

Gatekeeping Costs. If the doctor only cares about her own utility V , patients have no way to influence her prescription. To allow for a deviation 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 utility 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. We assume that the gatekeeping cost the doctor experiences is a share g of the patient’s lost utility relative to his

¹⁸It may be natural to assume that doctors prefer “more treatment” or “less treatment” than patients across both treatment options, but it is in principle possible to have wedges with different signs, e.g. $\Delta_L > 0$ and $\Delta_H < 0$, which would imply that doctors want to start treating with simple at lower levels of π , but prefer to switch to severe at a higher π .

first-best treatment choice.¹⁹

Patient-driven demand for simple treatment can occur for example when the doctor intrinsically prefers no treatment, but the patient strongly prefers simple, so that the gatekeeping costs exceed the utility loss the doctor incurs by writing a prescription: $gU_L(\pi, P_L) > |U_L(\pi, P_L) + \Delta_L|$.²⁰ In a similar manner, there may be patient-driven demand for severe treatment over simple or no treatment. Doctor-driven demand for severe treatment occurs when the patient prefers simple treatment, but the doctor prefers, and prescribes, severe treatment (note that this can only lead to a purchase when patients are prepared to buy severe treatment, i.e. $\pi \in L2$). Not only are preferences reversed between patients and doctors, but the utility the doctor gains from prescribing severe must outweigh the gatekeeping cost. This requires a higher wedge in doctor preferences in favor of severe over simple: $\Delta_H > \Delta_L$.

The range of beliefs π for which prescribing no or severe treatment incurs a gatekeeping cost increases when patients know that simple treatment is free (in the Patient Voucher condition). This reflects the fact that patient demand for simple treatment is higher when the price is lower; i.e. L expands. However, in the Doctor Voucher condition the doctor has the option to selectively conceal the voucher to reduce her gatekeeping costs from prescribing no treatment or severe treatment. We use this below to identify patient-driven and doctor-driven demand.

In the full model, we assume that patients and doctors have private signals about the patient's probability of being malaria positive. Unobserved signals allow for uncertainty on both sides: patients do not know in advance what prescription will be given, and the doctor does not know whether the patient will buy the prescription she makes, as we observe frequently in the data. This also means the doctor cannot be completely sure what gatekeeping cost she will incur when prescribing treatment t . In the face of this uncertainty, we permit doctors to avoid gatekeeping costs completely by letting the patient *choose* between simple and severe treatment. If the doctor takes this option, treatment is completely determined by patient preferences.²¹

¹⁹Note that when a patient prefers not to receive treatment, he can always choose not to purchase the prescription. This form of overprescription therefore incurs no gatekeeping cost.

²⁰Note that this requires $\Delta_L < 0$ since $U_L(\pi, P_L) > 0 > U_L(\pi, P_L) + \Delta_L$.

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

2.2 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.*

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 conceal the voucher from patients to whom she does not want to prescribe simple treatment and thereby avoid the increased gatekeeping costs when simple treatment is free. In PV , by contrast, these additional gatekeeping costs 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 gatekeeping costs are so high that the doctor is never willing to override the patient's preferences. Our predictions hold more generally; we formally establish these results in Appendix A.

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 for (simple) treatment.*

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). 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 high, this bar also reflects prescription and purchase outcomes under PV . The lowermost bar shows outcomes per the doctor's optimal strategy under DV .

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 patients with $\pi \in L1_P \cup L2_0$. 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 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. Note that we formulate this as a prediction about the demand for any treatment vs. no treatment, because, as the next prediction shows, the change in the rate of simple treatment is ambiguous.

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 severe 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 $\pi \in L1_0$, 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 ($\pi \in L2_0 \cap H_P$), thereby increasing the share of patients purchasing severe treatment relative to PV . The appendix establishes the general case: that any shift from simple in PV 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.*

This prediction highlights that our experiment can only detect patient-driven demand at the margin between no treatment and simple treatment,

and doctor-driven demand at the margin between severe treatment and simple treatment. Within our model, doctor-driven demand at the margin of any versus no treatment is not possible because doctors cannot force patients to buy unwanted treatment (although doctor-driven demand induced through influencing the patient’s beliefs, which is independent of the experimental interventions and therefore not part of our model, might still occur).

While our predictions work on different margins, finding evidence of demand driven by one side suggests that this side has a stronger preference for treatment than the other. For example, in the presence of patient-driven demand as in prediction (2), doctors have a weaker preference for simple treatment than patients, and so there would not be doctor-driven demand at this margin. Similarly, if there is doctor-driven demand for severe over simple treatment as in prediction (3), it implies that patients have a weaker preference for severe and so there would not be patient-driven demand for severe 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.

3 Experimental Design and Sample

3.1 Experimental Design

Sampling Frame. At the outset of our study, we obtained an administrative list of all CSComs in the city of Bamako and 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 were invited to attend a refresher training that covered Mali’s official malaria guidelines and a hands-on training on how to administer an RDT. 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. The trainings were conducted by five trainers from the PNL P and one trainer from the regional health directorate. Clinics were trained in six groups in November 2016; all clinics sent at least two and on average 3.9 providers to the training.

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 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 and were stationed at a separate part of the clinic from the survey team. Intervention eligibility was not tied to survey participation.

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 provider 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 treating health worker, 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 provider(s) on staff, who could assign the vouchers to patients at their discretion.

Intervention officers always brought more than enough vouchers to cover a day's demand, so voucher rationing was not an issue. Both conditions used the same voucher design, which required the provider'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.

Information Treatments. The Doctor and Patient Voucher treatments were cross-cut with a third within-clinic treatment. This “Patient Information” (*PI*) was conveyed through a short video, which emphasized that all suspected malaria cases should be confirmed with a malaria test, noted that RDTs should be available for free at the clinic, and described recommended treatment for simple and severe malaria. The video also demonstrated an RDT test and how to recognize a positive test result. Finally, half the clinics were randomly selected to receive a more intensive training designed to increase providers’ trust in the diagnostic accuracy of RDTs. Both information treatments are analyzed in detail in a separate paper. We find limited evidence that the information treatments impacted the size of *PV* and *DV* treatment effects (see Appendix Tables B3 and B4).

The within-clinic randomization and associated data collection were conducted after the health worker trainings in November and December 2016, covering the end of the rainy season and therefore the period of highest malaria risk. We divided the 60 clinics into three cohorts of 20 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.

3.2 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. On clinic 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, weak-

ness, fatigue, or reduced consciousness. Before the doctors' consultation, enumerators recorded patients' basic demographic details, symptoms, and information on any prior treatment/diagnosis.²² After the consultation was complete, 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.²³ 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.²⁴ Appendix Table B5 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 administrative data on voucher redemption to quantify how informing patients about vouchers impacted redemption rates.

3.3 Background Characteristics and Randomization Verification

Table 2 presents summary statistics on the study population and verifies randomization balance. The first column of the table displays means and standard deviations of baseline characteristics in the Control group (patients

²²We interviewed the patient whenever s/he was aged 15 or older and well enough to take the interview; otherwise a guardian answered the survey.

²³From 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 where more than six patients were seen, until 60 patients per study day were sampled.

²⁴1,158 patients gave consent; 32 tests were inconclusive.

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.²⁵ 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 health workers versus patients will be of special focus in the analysis, 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 5-6 acutely ill patients with eligible symptoms visited the clinic on the average observation day. 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 of patients visiting the clinic changes in response to treatment. 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 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. Individuals in the Patient Voucher group were 6 percentage points *less* likely than those in the

²⁵We include date fixed effects because the proportion of *PV*, *DV*, and *C* days in the sample varies depending on the date.

Doctor Voucher treatment to have a positive RDT. Malaria-positive patients are more likely to get treated for malaria and especially severe malaria, so this difference 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 survey 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 B6 shows that the characteristics of individuals selected for the home survey versus not, and interviewed versus not (conditional on selection), are balanced. There is, however, significant selection into taking the home RDT: patients who consented to take the home test are sicker and more likely to have purchased an antimalarial than patients who did not give consent.

3.4 Predicting Malaria Risk

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

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

where pos_{ict} is a dummy variable equal to one if a patient tests RDT positive and \mathbf{x}_{ict} is a vector including dummy variables for symptoms, 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.²⁶ The results of this regression, shown in Appendix Table

²⁶We 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 also control for 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. Finally we control for the patient and respondent being different people. Results are similar if we

B7, are used to impute a malaria risk for each patient. Appendix Figure B2 shows that the distribution of predicted malaria risk by treatment group is very similar, and Table 2 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 that are informative about malaria status, the malaria risk of tested patients may be different, conditional on recorded symptoms, from 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 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 malaria risk is the running variable. The share of individuals prescribed an antimalarial steadily increases with risk; lower risk patients are more likely to receive simple treatment, while higher risk patients are more likely to receive severe. 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 B3 shows that these patterns hold for those with valid home-based RDT tests and those without.

4 Main Results

4.1 Empirical Approach

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, Chernozhukov and Hansen, 2014).²⁷ Appendix Tables B8 and B9 show that results with no additional

omit the basic demographic characteristics from the probit.

²⁷Candidate 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

controls are very similar. In addition, Appendix Tables B10 and B11 indicate that results are analogous when using covariates without the lasso procedure.

We estimate heterogeneous treatment effects with respect to predicted malaria risk using the following specification:

$$y_{ict} = \delta_0 + \delta_{PHPV} \times high_{ict} + \delta_{DHDV} \times high_{ict} + \delta_{PLPV} \times low_{ict} + (3) \\ \delta_{DLDV} \times low_{ict} + \theta high_{ict} + \mathbf{x}_{ict}' \alpha + \lambda_t + \nu_{ict}$$

where $high_{ict}$ and low_{ict} are dummy variables identifying patients with above/below median predicted malaria risk, \mathbf{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 regressors, we bootstrap standard errors, clustering at the clinic level.²⁸

We now turn to the data to assess our theoretical predictions empirically.

4.2 Impacts on Antimalarial Prescriptions and Purchases

Overall Impacts. Table 3 shows treatment effects on voucher use, antimalarial prescriptions, and antimalarial purchases. We report the coefficients on the treatment dummies β_P and β_D and the p-value of a two-sided test of $H_0 : \beta_P = \beta_D$. We also indicate the theoretical mechanism being tested (existence of gatekeeping costs, doctor-driven, or patient-driven overprescription – GC, DD, or PD) and whether we find significant evidence of that mechanism or not.

The first column looks at *voucher redemptions*. We find strong, significant evidence in favor of gatekeeping costs: voucher redemption is 9 percentage dummies, illness duration (top-coded at the 99th percentile), 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.

²⁸First 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.

points (35 percent) higher in *PV* compared to *DV*. 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* had no impact on prescriptions but increased purchases by 8.5 percentage points. Thus, providers used vouchers for patients to whom they were giving prescriptions already, but not to make new prescriptions. We reject equality of β_P and β_D at the 10 and 5 percent levels for prescriptions and purchases respectively, indicating significant evidence of patient-driven demand.

The last two columns turn to *severe malaria treatment* to test for evidence of doctor-driven demand. Both voucher treatments are associated with a decline in severe malaria treatment. While we cannot reject equality of β_P and β_D , point estimates are, if anything, larger in magnitude under *DV*. This is the opposite of our theoretical prediction for doctor-driven demand for severe treatment.

Impacts by Predicted Malaria Risk. If patients pressure providers to prescribe, higher rates of voucher use and antimalarial consumption under *PV* (relative to *DV*) should be driven by patients at the lower end of the predicted risk distribution. If providers drive demand for severe malaria treatment, we should find substitution from severe to simple treatment in *PV* when malaria risk is higher, because providers 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.²⁹ For reference, vertical lines demarcate the 25th, 50th, and 75th percentiles of the predicted risk distribution.

Again, our predictions for patient-driven demand for simple treatment 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

²⁹To account for clustering and the fact that predicted risk is a generated regressor, confidence intervals are bootstrapped using the procedure described in sub-section 4.1.

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. This cuts *against* the prediction for doctor-driven demand.

To quantify these patterns, Table 4 looks at heterogeneity in treatment effects by above versus below median predicted risk. While rates of voucher use among “high risk” patients in *DV* and *PV* are very similar, voucher use by “low risk” patients is 15 percentage points (71 percent) higher when patients know about the vouchers before consulting the doctor. We see similar patterns for any antimalarial prescription and purchase, where *PV* increases low-risk prescriptions by 8.1 percentage points (16 percent) and purchases by 9 percentage points (23 percent) relative to *DV*. These differences are significant at the 5 percent level. Finally, columns 4 and 5 explore effects for severe malaria treatment. While we formally reject $H_0 : \delta_{PH} = \delta_{DH}$ at the 1 percent level in column 5, the point estimates run counter to the theoretical predictions for doctor driven demand.

Adherence to Treatment Guidelines. Why would use of severe malaria treatment – especially among high risk patients – be *lower* in *DV* relative to *PV*? Figure 5 offers an explanation: providers 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 experimental arm. In the Control group, none of the 26 percent of patients purchasing severe malaria treatment used a free ACT voucher. 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 provider 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 providers 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. Columns 1 and 2 in Appendix Table B12 verify that the rate of rule breaking is significantly higher in *PV* versus *DV*. This behavior is also consistent with patient-driven demand: providers 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*. Non-voucher purchases may have occurred because antimalarials other than ACTs were given, because patients visited the clinic when the intervention officer was absent, or because the provider or patient forgot to apply the voucher. Additionally, our model predicts that providers will sometimes intentionally prescribe simple treatment without a voucher, and that they do so more often in *DV* (see prediction (5) in Appendix A). Intuitively, a provider who prefers to prescribe something different from simple treatment, but suspects that her patient has a strong preference for it, may offer simple malaria treatment without the voucher. This “hedging” lets her avoid gatekeeping costs if the patient turns out to have a strong preference for simple treatment, while ensuring the patient will not purchase simple if his preferences are not strong. 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 B12 shows that patients in *DV* were three percentage points more likely than patients in *PV* to purchase simple malaria treatment without a voucher (significant at the 10 percent level).

Other Explanations. Alternatively, it could be that providers in *DV* simply forgot about the vouchers in some cases, whereas in *PV* they were reminded by patients. This could lower voucher use in *DV* relative to *PV* and explain why patients in *DV* are relatively less likely to receive/purchase malaria treatment (Table 3). This mechanism cannot, however, explain why the reduction in voucher use in *DV* is concentrated among low risk patients, while prescriptions for severe treatment are if anything lower in *DV* than in *PV*. With limited attention we would expect to see less use of vouchers across the risk distribution under *DV* versus *PV*, meaning *more* use of severe malaria treatment.

Another possibility is that knowledge of the vouchers led patients to exaggerate their symptoms, either intentionally or subconsciously. If this led providers to overestimate malaria risk and therefore prescribe more antimalarials, this would amount to a form of patient-driven demand that operates through changing providers’ direct utility, rather than gatekeeping

costs.³⁰ However, the fact that providers almost never violated the rules of the vouchers in *DV*, but did in *PV*, suggests that gatekeeping costs matter in our context.

A final question is whether differences between *PV* and *DV* reflect differential rates of patients “stockpiling” ACTs for future use. 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.³¹ Appendix Table B13, 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* relative to *DV*.

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 all 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 Hawthorne effects to have increased the likelihood of providers adhering to official treatment guidelines.

Since treatments rotated within clinic, this would impact behavior in *all* treatment arms, which would bias us away from finding *PV-DV* differences in outcomes. Consistent with this (and the hypothesis that Hawthorne effects decay across observation days, see e.g. Leonard and Masatu (2010)),

³⁰A related possibility would be if providers perceived more patients in *PV* to be “likely compliers” with prescribed simple malaria treatment, if for example more patients explicitly requested simple malaria treatment.

³¹ACT 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.

Appendix Figure B4 shows that *PV-DV* differences are larger after the first couple of clinic observation days. At the same time, we find that only 56 percent of patients prescribed an antimalarial on Control days received a malaria test, and there is no significant difference in this “test given prescription” rate across treatment arms, even though testing before treatment was a major focus of the training and a requirement for using the voucher. Overall this suggests the implicit monitoring that took place during the experiment had limited effects on behavior that, if anything, bias us away from detecting effects.

A related worry is that *PV-DV* differences are driven by experimenter demand effects on the patient side (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, it is important for interpretation. 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 provider’s signature, and that the doctor was the one who would make the malaria diagnosis. Second, data from the patient information treatment (*PI*) suggests patients were not influenced by experimenter demand effects. The *PI* video communicated that suspected malaria cases should only be treated if they received a positive test. However, Appendix Table B14 shows that *PI* actually *decreased* the share of total patients and antimalarial prescription-holders who took a malaria test, while leaving the share of patients who were prescribed an antimalarial unchanged. Panel B shows there is no significant interaction effect with the voucher treatment arms, 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 patients having preferences of their own about whether to get tested.

4.3 Do Doctors Use Vouchers to Increase Revenue?

Financial motives are one of the most cited reasons for doctor-driven demand for drugs. We argued above that doctor-driven demand for severe treatment would arise if (relative to patient preferences) providers get substantially more utility from severe over simple treatment ($\Delta_H > \Delta_L$). A health worker who cares mostly about revenue would be a plausible explanation for this, since severe treatment is much pricier and incurs more high-margin service

fees. Having said that, providers could opt to use the vouchers to prescribe simple treatment instead of severe because this strategy leads to an almost guaranteed purchase. In this case, providers might be revenue driven and motivated to oversell treatment, but they may not be “driving demand” (prescribing more treatment than patients would like) per our definition.

Fortunately we can test whether providers are revenue motivated independent of whether this leads doctors to drive demand for severe malaria treatment: prediction (6) in Appendix A shows that vouchers will lead to an increase in clinic revenue when doctors are purely revenue motivated, and prediction (7) says that revenue-maximizing doctors will use the informational advantage in the *DV* treatment to increase clinic revenues relative to *PV*.³²

In practice, we find no significant differences in (per-patient) clinic revenues across treatment arms (Appendix Table B15). Moreover, 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). Lastly, from the results we already saw, ACT prescriptions and purchases for low-risk patients, voucher use, and overall treatment rates in *DV* do not increase at the same rate as in *PV*. This is not consistent with providers who are simply trying to sell as much treatment as possible. All told, our results do not paint a picture of providers primarily motivated by revenue, at least on the margin.

4.4 Quantifying Contributions to Overtreatment

What do doctors’ gatekeeping failures mean 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. Unfortunately, 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 B6). 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,

³²The greater number of patients on voucher days versus control days complicates our test of prediction (6), so we place more weight on prediction (7).

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 desirable 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 increase the number of correct positives (mp_{ict}) for prescriptions, they did for purchases (1.8-2.3 percentage points, significant at the 1 percent level). This amounts to a 15-19 percent increase over the mean in the Control group and underscores the rationale of subsidizing health care: some truly sick patients may not purchase treatment due to credit constraints or because they value treatment below its cost.

The higher treatment rates for malaria-positive patients come at a price, 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. In PV , this share is 11 percentage points (24 percent) lower, in DV it is 6.3 percentage points (14 percent) lower. The PV - DV difference is significant at the 1 percent level, underscoring the concern that patient-driven demand at low drug prices leads to overuse.

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. In PV , this share is significantly lower for both prescriptions (3.2 percentage points) and purchases (8.3 percentage points), while it is lower for purchases only in DV (3.9 percentage points). The 8.3 percentage point lower 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 in patients’ non-purchase rate each account for roughly half of the additional misalloca-

tion we observe with the ACT subsidy.

Changes in mismatch rates relative to the Control need to be interpreted with caution, given that more patients visit clinics on *PV* and *DV* days. More generally, we cannot extrapolate our findings in the voucher treatments to mismatch rates we would observe in a long-term equilibrium with subsidized ACTs, since the composition of patients visiting the clinic would likely change. However, we have no evidence of differential selection in *PV* and *DV*, and the analysis indicates that patient-driven demand in our sample significantly worsens the treatment-illness match.

One 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.³³ This would bias our expected match measures downwards, 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*.

Appendix Table B16 uses the home-tested RDT subsample to get an idea of how much private information skews our expected match estimates. We find strikingly similar treatment effects on the expected and actual match for the home-tested subsample. Control group means are also similar, at 0.48 for the expected match and 0.51 for the actual match. 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 providers often do not use, or do not comply with malaria tests.

5 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 range of reasons:

³³This is captured by the private signal η in our model in the Appendix.

poor households with very high marginal utilities of income waste precious resources, overuse of vital drugs like antimalarials and antibiotics promotes disease resistance, and unnecessary consumption 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 providers 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 they do not bear the cost of treatment, because it is in their financial interest to increase revenue, 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 may 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.

We begin by developing a model that can generate both doctor- and patient-driven demand for treatment, and delivers testable implications for both channels (albeit at different margins of substitution). 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.

We find that vouchers for free malaria treatment are much more likely to be used when both patients and providers (as opposed to just providers) know about their availability, indicating that gatekeeping costs are empirically important. Patients with the lowest levels of malaria risk receive more prescriptions and buy more treatment, which is consistent with patient-driven demand for simple malaria treatment. We also find that this additional demand significantly worsens overtreatment. On the other hand, we find no evidence that physicians in the Doctor Voucher treatment exploit their informational advantage to upsell patients more expensive severe malaria treatment, or increase clinic revenues more generally.

Our approach is agnostic about the underlying factors that might drive doctors' or patients' preferences for administering treatment even when malaria risk is low; and our results do not rule out the possibility that doctor preferences are driving a share of the baseline misallocation we see in our sample. However, our findings do show that marginal misallocation from price subsidies is attributable to patient preferences and behavior. Doctors are no more likely to write an antimalarial prescription when only they know about the price discounts; this margin only moves when both providers *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. Our findings imply 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 model highlights how the effects of gatekeeping costs will be determined by the nature of *mismatch* between doctor and patient preferences – in our setting patients prefer more simple malaria treatment than providers would like to prescribe. Doctor-patient preference gaps may be different in other contexts where providers 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, Jane, Ambrose O Talisuna, Annette Erhart, Adoke Yeka, James K Tibenderana, Frederick N Baliraine, Philip J Rosenthal, and Umberto 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, Evelyn K, Solomon Narh-Bana, Michael Epokor, Samson Akanpibiam, Alberta Amu Quartey, John Gyapong, and Christopher JM 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.
- Arrow, K., C. Panosian, and H. Gelband,** ed. 2004. *Saving Lives, Buying Time: Economics of Malaria Drugs in an Age of Resistance*. Washington, DC:Institute of Medicine: National Academies Press.
- Ashraf, Nava, James Berry, and Jesse M. Shapiro.** 2010. "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia." *American Economic Review*, 100(5): 2383–2413.
- Banerjee, Abhijit V, Esther Duflo, and Rachel 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, Alexandre, Victor Chernozhukov, and Christian Hansen.** 2014. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives*, 28(2): 29–50.
- Bisoffi, Zeno, Bienvenu Sodiomon Sirima, Andrea Angheben, Claudia Lodesani, Federico Gobbi, Halidou Tinto, and Jef 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, Jessica, and Pascaline Dupas.** 2010. “Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment.” *The Quarterly Journal of Economics*, 125(1): 1–45.
- Cohen, Jessica, Pascaline Dupas, and Simone 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, Janet, Wanchuan Lin, and Juanjuan Meng.** 2014. “Addressing Antibiotic Abuse in China: An Experimental Audit Study.” *Journal of Development Economics*, 110: 39–51.
- Das, Jishnu, Alaka Holla, Aakash Mohpal, and Karthik Muralidharan.** 2016. “Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India.” *American Economic Review*, 106(12): 3765–99.
- Das, Jishnu, and Jeffrey Hammer.** 2014. “Quality of Primary Care in Low-Income Countries: Facts and Economics.” 6: 525–553.
- Das, Jishnu, Jeffrey Hammer, and Kenneth Leonard.** 2008. “The Quality of Medical Advice in Low-Income Countries.” *The Journal of Economic Perspectives*, 22(2): pp. 93–114.
- Dupas, Pascaline.** 2014. “Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment.” *Econometrica*, 82(1): 191–228.
- Dupas, Pascaline, and Edward Miguel.** 2017. “Impacts and Determinants of Health Levels in Low-Income Countries.” In *Handbook of Economic Field Experiments*. Vol. 2, 3–93. Elsevier.
- Escot, Fabrice.** 2012. “Bilan des diagnostics socio-sanitaires et des ressources humaines de la santé de la région de Sikasso.”
- Fleming-Dutra, Katherine E, Adam L Hersh, Daniel J Shapiro, Monina Bartoces, Eva A Enns, Thomas M File, Jonathan A Finkelstein, Jeffrey S Gerber, David Y Hyun, Jeffrey A Linder, et al.** 2016. “Prevalence of Inappropriate Antibiotic Prescriptions among US Ambulatory Care Visits, 2010–2011.” *JAMA*, 315(17): 1864–1873.

- Hamer, Davidson H, Micky Ndhlovu, Dejan Zurovac, Matthew Fox, Kojo Yeboah-Antwi, Pascalina Chanda, Naawa Sipilinyambe, Jonathon L Simon, and Robert W Snow.** 2007. “Improved diagnostic testing and malaria treatment practices in Zambia.” *JAMA*, 297(20): 2227–2231.
- Humar, Atul, Colin Ohrt, Mary Anne Harrington, Dylan Pillai, and Kevin 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.** 2018. “Global Burden of Disease Study 2017 (GBD 2017) Results.” *Online*, Accessed May 22, 2020.
- Iizuka, Toshiaki.** 2012. “Physician Agency and Adoption of Generic Pharmaceuticals.” *American Economic Review*, 102(6): 2826–58.
- Koné, Diakalia, Drissa Coulibaly, Ogobara Doumbo, Ibrahima-Socé Fall, Eliud Kibuchi, Bernard Mitto, Gilbert Sang, David Kyalo, Robert W. Snow, and Abdisalan M. Noor.** 2015. *An Epidemiological Profile of Malaria in Mali*. Programme National de Lutte contre le Paludisme (PNLP).
- Kotwani, Anita, Chand Wattal, Shashi Katewa, PC Joshi, and Kathleen Holloway.** 2010. “Factors Influencing Primary Care Physicians to Prescribe Antibiotics in Delhi India.” *Family Practice*, 27(6): 684–690.
- Kravitz, Richard L, Ronald M Epstein, Mitchell D Feldman, Carol E Franz, Rahman Azari, Michael S Wilkes, Ladson Hinton, and Peter Franks.** 2005. “Influence of Patients’ Requests for Direct-to-Consumer Advertised Antidepressants: A Randomized Controlled Trial.” *JAMA*, 293(16): 1995–2002.
- Kremer, Michael, Jessica Leino, Edward Miguel, and Alix Zwane.** 2011. “Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions.” *The Quarterly Journal of Economics*, 126(1).
- Kyabayinze, Daniel J, James K Tibenderana, George W Odong, John B Rwakimari, and Helen 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.

Leonard, Kenneth L., and Melkiory C. Masatu. 2010. “Using the Hawthorne effect to examine the gap between a doctor’s best possible practice and actual performance.” *Journal of Development Economics*, 93: 226–234.

Linder, Jeffrey A, Jason N Doctor, Mark W Friedberg, Harry Reyes Nieva, Caroline Birks, Daniella Meeker, and Craig R Fox. 2014. “Time of Day and the Decision to Prescribe Antibiotics.” *JAMA Internal Medicine*, 174(12): 2029–2031.

Li, Yongbin, Jing Xu, Fang Wang, Bin Wang, Liqun Liu, Wanli Hou, Hong Fan, Yeqing Tong, Juan Zhang, and Zuxun 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.

Ministère de la Santé. 2013. *Politique Nationale De Lutte Contre Le Paludisme.* Republique du Mali.

PMI. 2015. “President’s Malaria Initiative Mali: Malaria Operational Plan FY 2016.” USAID.

PMI. 2016. “President’s Malaria Initiative Mali: Malaria Operational Plan FY 2017.” USAID.

PMI. 2017. “President’s Malaria Initiative Mali: Malaria Operational Plan FY 2018.” USAID.

PNLP, INSTAT, INFO-STAT, INRSP, and ICF International. 2016. “Enquete sur les Indicateurs du Paludisme au Mali (EIPM) 2015.” INSTAT, INFO-STAT and ICF International, Rockville, Maryland, USA.

Prendergast, Canice. 2003. “The Limits of Bureaucratic Efficiency.” *Journal of Political Economy*, 111(5): 929–958.

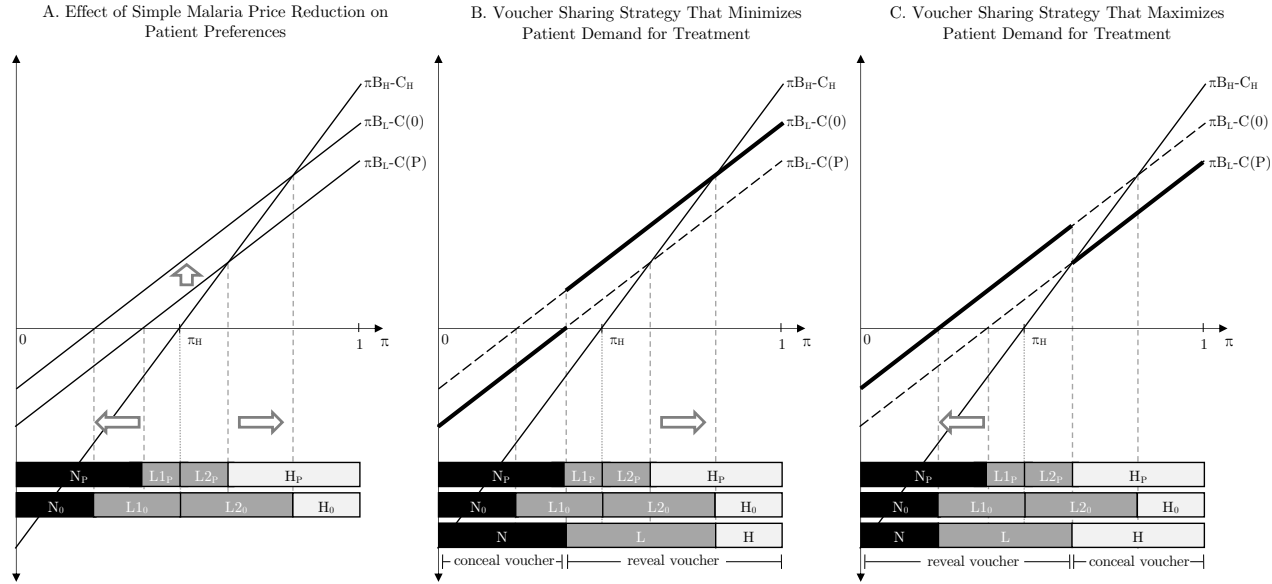
Reyburn, H, R Mbatia, C Drakeley, I Carneiro, E Mwakasungula, O Mwerinde, K Saganda, J Shao, A Kitua, R Olomi, BM Greenwood, and CJ Whitty. 2004. “Overdiagnosis of Malaria in Patients with

- Severe Febrile Illness in Tanzania: A Prospective Study.” *British Medical Journal*, 329(7476): 1212.
- Sautmann, Anja, Samuel Brown, and Mark Dean.** 2017. “Subsidies, Information, and the Timing of Children’s Health Care in Mali.”
- Schwartz, Lisa M., and Steven Woloshin.** 2019. “Medical Marketing in the United States, 1997-2016.” *JAMA*, 321(1): 80–96.
- Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong.** 2014. “Micro-loans, Insecticide-Treated Bednets and Malaria: Evidence from a randomized controlled trial in Orissa (India).” *The American Economic Review*, 104(7): 1909–41.
- Taylor, Steve M, Malcolm E Molyneux, David L Simel, Steven R Meshnick, and Jonathan J Juliano.** 2010. “Does This Patient Have Malaria?” *JAMA: The Journal of the American Medical Association*, 304(18): 2048–2056.
- Touré, Laurence.** 2015. “User fee exemption policies in Mali: sustainability jeopardized by the malfunctioning of the health system.” *BMC Health Services Research*, 15(3): S8.
- Trampuz, Andrej, Matjaz Jereb, Igor Muzlovic, and Rajesh M Prabhu.** 2003. “Clinical Review: Severe Malaria.” *Critical Care*, 7(4): 315.
- Whitehead, Margaret, Göran Dahlgren, and Timothy 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.” World Health Organization.
- WHO.** 2017. “Malaria rapid diagnostic test performance: results of WHO product testing of malaria RDTs: round 7 (2015-2016).” World Health Organization.
- WHO.** 2018. “MCEE Methods and Data Sources for Child Causes of Death 2000-2016.”

Wongsrichanalai, Chansuda, Mazie J Barcus, Sinuon Muth, Awaludin Sutamihardja, and Walther 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.

Figure 1: Patient Preferences and Doctor Voucher-Sharing Strategies

42



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”, $L1$ denotes “simple malaria treatment, unwilling to purchase severe malaria treatment”, $L2$ denotes “simple malaria treatment, willing to purchase severe malaria treatment”, and H 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.

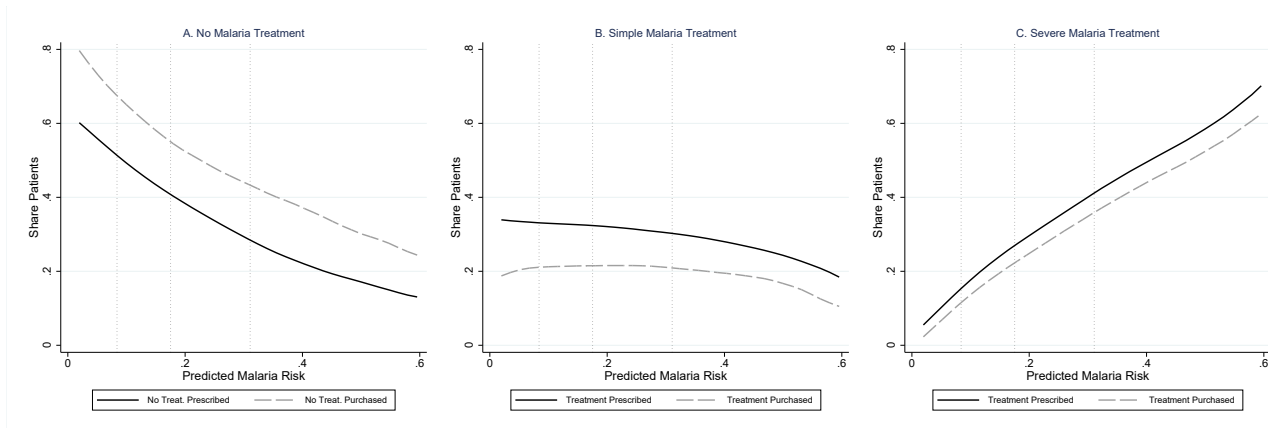
Figure 2: Within-CSCOM Randomization Design

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: Before the interventions launched, clinic staff received trainings on RDTs on November 2, 3, and 4 2016. 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



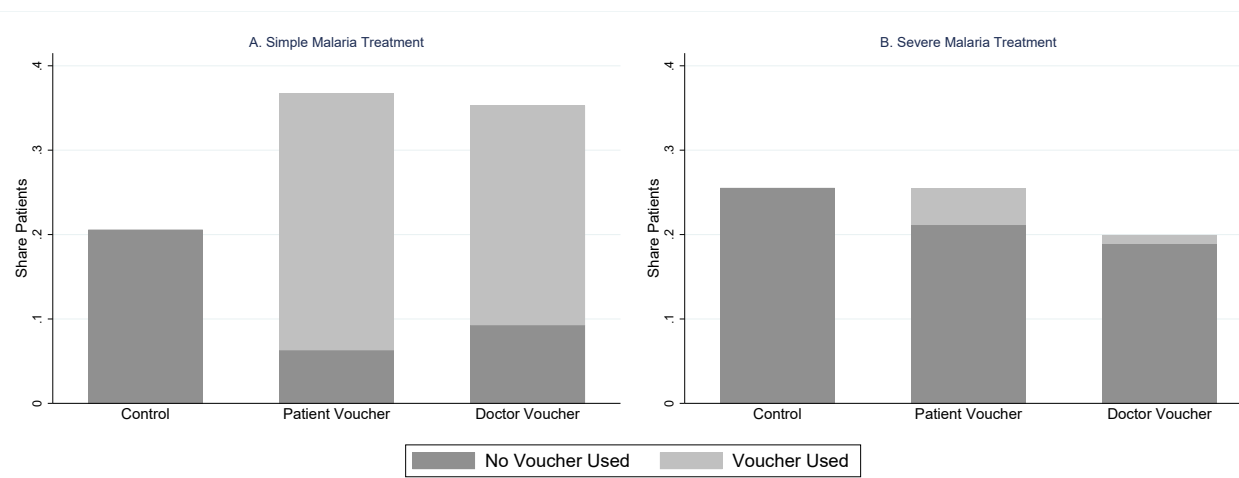
Notes: Results from local linear regressions. Regressions are run on the full sample, but graphs omit results for top and bottom 2.5 percent of malaria risk distribution to avoid influence of outliers. Vertical dashed lines indicate 25th, 50th, and 75th percentiles of predicted malaria risk respectively.

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 share of patients receiving the specified treatment by voucher condition.

Table 1: Descriptive Statistics from Control Group

	(1)	(2)	(3)	(4)	(5)
	Full Sample		Has Valid RDT (Home Survey)		
	N	Mean	All Patients	RDT Positive	RDT Negative
<i>Panel A. Malaria Beliefs and Treatment Outcomes</i>					
Respondent Suspects Illness is Malaria (Pre-Consultation)	628	0.551	0.610	0.791	0.550
Simple Malaria Treatment: Prescribed	627	0.313	0.295	0.163	0.338
Simple Malaria Treatment: Purchased	627	0.206	0.197	0.140	0.215
Severe Malaria Treatment: Prescribed	627	0.303	0.327	0.593	0.238
Severe Malaria Treatment: Purchased	627	0.255	0.283	0.535	0.200
<i>Panel B. Other Treatment Outcomes</i>					
Purchased Antibiotics	627	0.426	0.442	0.430	0.446
Received Injection or IV	628	0.398	0.454	0.674	0.381
Patient Referred to Hospital or Placed Under Observation	628	0.139	0.179	0.395	0.108
<i>Panel C. Patient Costs</i>					
Total Cost No Malaria Treatment (CFA)	301	2793.349	–	–	–
Total Cost Simple Malaria (CFA)	126	4474.873	–	–	–
Total Cost Severe Malaria (CFA)	147	10352.714	–	–	–

Notes: Column 1 reports the number of non-missing observations of variables among all individuals in the Control group. Column 2 reports means for this sample. Columns 3-6 limit the sample to Control group individuals who were administered a valid RDT test at home, with column 3 reporting means for all patients with a valid test, column 4 reporting means for patients with a positive test, and column 5 reporting means for patients with a negative test. Sample sizes for columns 3, 4, and 5 are 346, 86, and 260 respectively. Variables measured in CFA top-coded at the 99th percentile. CFA610 \approx USD1.

Table 2: Demographic Characteristics and Randomization Verification

	(1)	(2)	(3)	(4)	(5)	(6)
	Control Mean	Regression Coefficients		P-Values		N
		Patient Voucher	Doctor Voucher	Joint Test PV=DV	Joint Test PV=DV=0	
<i>A. Sample Frame (Clinic × Day-Level Observations)</i>						
Number Eligible Logged Patients	5.233 [3.322]	0.757** (0.348)	0.634** (0.312)	0.725	0.054*	360
<i>B. Patient Characteristics (Individual-Level Observations)</i>						
Number of symptoms	3.559 [1.578]	0.023 (0.081)	-0.090 (0.086)	0.114	0.270	2055
Fever	0.801 [0.400]	0.045* (0.024)	0.012 (0.022)	0.096*	0.130	2055
Chills or Excessive Sweating	0.298 [0.458]	-0.045 (0.027)	-0.051** (0.024)	0.782	0.092*	2055
Nausea, Vomiting, or Diarrhea	0.497 [0.500]	-0.004 (0.028)	-0.024 (0.026)	0.455	0.599	2055
Poor Appetite	0.505 [0.500]	-0.023 (0.031)	-0.040 (0.026)	0.501	0.292	2055
Headache	0.621 [0.486]	-0.006 (0.023)	-0.007 (0.023)	0.963	0.951	2055
Cough	0.363 [0.481]	0.027 (0.025)	-0.005 (0.028)	0.230	0.403	2055
Weakness/Fatigue	0.475 [0.500]	0.029 (0.028)	0.025 (0.027)	0.881	0.538	2055
Duration of Illness in Days	4.299 [4.276]	0.180 (0.233)	-0.007 (0.268)	0.444	0.648	2055
Age	16.603 [15.096]	1.545** (0.760)	1.282* (0.740)	0.760	0.078*	2055
Under 5 Years Old	0.314 [0.464]	-0.041 (0.025)	-0.034 (0.022)	0.746	0.179	2055
Male	0.395 [0.489]	0.031 (0.030)	0.058** (0.026)	0.377	0.102	2055
Patient is Pregnant	0.105 [0.307]	0.003 (0.022)	-0.016 (0.020)	0.331	0.567	1139
Positive RDT (home)	0.249 [0.433]	-0.037 (0.027)	0.028 (0.027)	0.015**	0.051*	1126
Predicted Malaria Risk	0.212 [0.160]	0.008 (0.007)	0.008 (0.007)	0.981	0.361	2055
<i>C. Respondent and Household Characteristics (Individual-Level Observations)</i>						
Patient Answered Clinic Survey	0.470 [0.499]	0.003 (0.026)	0.003 (0.024)	0.990	0.990	2055
Male	0.264 [0.441]	0.017 (0.024)	0.039* (0.022)	0.350	0.213	2055
Bambara	0.377 [0.485]	0.006 (0.023)	0.011 (0.027)	0.850	0.926	2053
Speaks French	0.518 [0.500]	-0.015 (0.031)	0.000 (0.027)	0.550	0.822	2055
Literate (in French)	0.264 [0.441]	-0.005 (0.030)	-0.012 (0.028)	0.772	0.901	2055
Primary School or Less	0.424 [0.495]	0.046 (0.033)	0.047 (0.030)	0.962	0.248	2055
Household Size ⁺	10.308 [8.111]	0.347 (0.506)	0.729 (0.602)	0.500	0.483	1491
Share HH Under 15 ⁺	0.412 [0.191]	0.017 (0.013)	0.013 (0.014)	0.777	0.397	1485
Share HH Members Working ⁺	0.258 [0.188]	-0.004 (0.010)	-0.002 (0.011)	0.886	0.937	1485
Monthly income per capita ⁺	21987.613 [25298.328]	-2143.183 (1388.442)	-2928.525** (1459.962)	0.522	0.135	1432
Rental Value of Home ⁺	63823.469 [90722.250]	-2452.218 (5457.584)	-5053.802 (5934.901)	0.607	0.694	1469
Mosquito Nets Per Capita ⁺	0.491 [0.342]	-0.016 (0.022)	-0.012 (0.021)	0.830	0.766	1482

Notes: Robust standard errors clustered at the clinic level in parentheses. All regressions include clinic visit date fixed effects. ⁺ 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 and theory-driven tests</i>					
$\beta_P = \beta_D$	0.011**	0.079*	0.018**	0.735	0.193
Test for mechanism:	GC	PD	PD	DD	DD
Significant evidence of mechanism:	Yes	Yes	Yes	No	No
Mean (Control)	0.000	0.616	0.461	0.303	0.255
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the clinic 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 and theory-driven tests</i>					
$\delta_{PH} = \delta_{DH}$	0.749	0.689	0.275	0.101	0.007***
Test for mechanism:	GC/DD	–	–	DD	DD
Significant evidence of mechanism:	No	–	–	No	No
$\delta_{PL} = \delta_{DL}$	0.000***	0.030**	0.014**	0.481	0.435
Test for mechanism:	GC/PD	PD	PD	–	–
Significant evidence of mechanism:	Yes	Yes	Yes	–	–
Mean (Control, Low Risk)	0.000	0.486	0.329	0.154	0.116
N	2055	2053	2053	2053	2053

Notes: Robust standard errors clustered at the clinic 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</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)	0.153	0.326	0.479	0.122	0.449	0.570
N	2053	2053	2053	2053	2053	2053

Notes: Robust standard errors clustered at the clinic 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.