

NBER WORKING PAPER SERIES

SELF-CONTROL AND DEMAND FOR PREVENTIVE HEALTH:
EVIDENCE FROM HYPERTENSION IN INDIA

Liang Bai
Benjamin Handel
Edward Miguel
Gautam Rao

Working Paper 23727
<http://www.nber.org/papers/w23727>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2017

We thank Child Relief International, the Center for Equitable Growth, and the Berkeley Population Center for funding. We thank Allyson Barnett, Fenella Carpena, Carson Christiano, Jen Kwok, and Deepak Saraswat for excellent research assistance. We are also grateful to seminar audiences at U.C. Berkeley, Harvard University, Stanford, CEU, Wharton, UPF and numerous others for useful suggestions. All errors remain our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Liang Bai, Benjamin Handel, Edward Miguel, and Gautam Rao. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Self-Control and Demand for Preventive Health: Evidence from Hypertension in India
Liang Bai, Benjamin Handel, Edward Miguel, and Gautam Rao
NBER Working Paper No. 23727
August 2017
JEL No. D91,I12

ABSTRACT

Self-control problems constitute a potential explanation for the under-investment in preventive health care observed in low-income countries. A commonly proposed policy tool to solve such problems is offering consumers commitment devices. We conduct a field experiment to evaluate the effectiveness of different types of theoretically-motivated commitment contracts in increasing preventive doctor visits by hypertensive patients in rural India. We document varying levels of takeup of the different commitment contracts, but find no effects on actual doctor visits or individual health outcomes. Thus, a substantial number of individuals pay for commitments, but then fail to follow through on the specified task, losing money without experiencing any health benefit. We develop and structurally estimate a pre-specified model of consumer behavior under present bias with varying levels of naivete. The results are consistent with a large share of individuals being partially naive about their own self-control problems: in other words, they are sophisticated enough to demand some commitment, but overly optimistic about whether a given commitment is sufficiently strong to be effective. The results suggest that commitment devices may in practice be welfare diminishing, at least in some contexts, and serve as a cautionary tale about the role of these contracts in the health care sector.

Liang Bai
School of Economics
University of Edinburgh
Liang.Bai@ed.ac.uk

Benjamin Handel
Department of Economics
University of California, Berkeley
508-1 Evans Hall #3880
Berkeley, CA 94720
and NBER
handel@berkeley.edu

Edward Miguel
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720
and NBER
emiguel@econ.berkeley.edu

Gautam Rao
Department of Economics
Harvard University
Littauer M-30
Cambridge, MA 02138
and NBER
grao@fas.harvard.edu

1 Introduction

Preventive health care is an important determinant of health outcomes, and may reduce overall health spending by reducing the incidence of serious conditions. Yet preventive care is widely thought to be under-utilized in both wealthy and poor countries (Kenkel 2000, Dupas 2011). In this paper, we study the adoption of preventive health behaviors in the management of hypertension in India. Hypertension is a chronic health condition and an increasingly important public health problem in low-income regions. In India, it is estimated that nearly half of adults aged 45 years and above are at risk, and rural India has seen an eight-fold increase in prevalence over the past six decades (Association of Physicians of India [API], 2013a; Mohan et al. 2013).

This paper focuses on the role of self-control problems in preventive health behavior. The intuition for why self-control problems might generate under-investment in preventive care is simple: preventive care often requires behaviors – such as visiting a doctor or changing one’s diet – which involve *utility costs* in the present, while the returns are in the distant future. An individual with limited self-control, say due to present bias, might procrastinate on engaging in such preventive behaviors, especially if that person is naive or overconfident about their future levels of self-control (O’Donoghue and Rabin 1999). While this intuition is clear, the role of present bias in driving health choices in low-income countries remains poorly understood and controversial. Two recent handbook chapter reviews on health in developing countries reach somewhat different conclusions about its empirical importance, with Glennerster and Kremer (2011) writing that current evidence “provide[s] considerable support for a present bias model”, while Dupas and Miguel (2017) argue that the empirical evidence remains inconclusive, although they do write that “the role of present bias in health decision-making may ... grow in developing countries as the burden of disease tilts towards non-communicable diseases (NCDs) like diabetes and hypertension”.

The most common way researchers have attempted to tackle self-control problems is through the provision of commitment devices or contracts (CCs). CCs allow individuals to voluntarily restrict their future choice set, or increase the costs of certain potential future actions. An individual without self-control problems should not have a demand to create such incentives for their future selves. Yet, researchers have documented demand for commitment in a number of contexts, such as savings commitments (Ashraf et al. 2006), smoking cessation (Gine et al. 2010), alcohol consumption (Schilbach 2017), fertilizer use (Duflo et al. 2011) and work effort (Kaur et al. 2015), among others.¹ These studies have provided ‘smoking gun’ evidence on the existence of self-control problems, and have suggested commitment contracts as a promising policy option.

However, theory provides an important warning about the effectiveness and welfare implications of commitment contracts: they are predicted to work well only when agents are sufficiently sophisticated. Sophisticated agents, who have accurate beliefs about their future levels of self-control,

¹Other recent studies of commitment devices include Royer et al (2015) on exercise commitments and Alsan et al. (2014) on saving commitments for education. Bryan et al. (2010) provide a review of the earlier literature.

will correctly predict how the incentives embedded in a particular CC will lead them to act in the future. They will choose commitment wisely to reach the first best from a long-run perspective. In contrast, fully naive agents, who (incorrectly) believe that their future selves will have *no* self-control problems, will have no demand for commitment. Such naive agents will be neither helped nor harmed by being offered commitment. The intermediate case is that of *partially naive* agents, who realize that they will continue to have limited self-control, but are overconfident about the level of their future self-control. Such agents might unwisely purchase CCs which provide “too little” commitment due to underestimating the magnitude of their future self-control problem, with potentially adverse welfare consequences.² While the literature provides little systematic evidence on the distribution of sophistication and naivete in the population, the best existing evidence suggests a high degree of naivete (Augenblick and Rabin 2015), consistent with the modest overall take-up of commitment highlighted by Laibson (2015). In summary, while CCs have the potential to improve individual health outcomes and wellbeing, it remains an open empirical question whether the risks highlighted here outweigh the benefits for agents who are induced to obtain more preventive care by CCs.

The current study was designed to explore these issues in the case of preventive care for hypertension in rural Punjab, India. The experimental interventions consisted of different types of commitment contracts for attendance at village “Hypertension Day” health camps conducted by a growing private-sector health care provider (henceforth the Provider). The health camps should be thought of as high-quality health care in the context of rural India, featuring a consultation with trained and licensed medical personnel, including doctors and nurses, and state of the art diagnostics and access to standard-of-care medications. A central goal of the interventions was to boost preventive care for hypertension among mature adults in our sample villages who had been identified (through screening) to either have hypertension or to be at high risk of developing the condition. The CCs we studied asked individuals to make an up-front payment, reducing the marginal cost of health camp visits in the future, similar to the contracts modeled in DellaVigna and Malmendier (2004). Some contracts required up-front payments larger than the visit fee, such that participants received some money back each time they made a recommended bi-monthly preventive visit to a health camp. We also subsidized the visits for a cross-cutting random subset of participants. These price discounts served two roles: to assess the role of affordability or liquidity constraints in dampening take-up, and to encourage even naive agents to sign up for commitment contracts, in anticipation of benefiting from the price discount. The control group instead pays for the preventive health care in a typical fee-for-service manner. The experiment was also designed to reduce the role played by poor consumer information by making sure that all treatment and control participants had at least some basic level of awareness about the health condition and recommended

²This point has been made theoretically by Heidhues and Koszegi (2010) and DellaVigna and Malmendier (2004), and empirically by DellaVigna and Malmendier (2006) in the context of gym attendance in the United States, and John (2017) in the context of savings plans in the Philippines.

preventive health behaviors.

An important aspect of the research design is the use of both *fixed commitment contracts*, where individuals are asked to accept or reject a particular contract (with a given up-front payment and future money-back amounts) as well as *personalized commitment contracts*, where the individual is able to choose the amount of up-front payment, and thus the strength of the commitment. In other words, individuals offered a personalized CC can choose (within bounds) the amount of up-front payment, and therefore the later amount returned to them at each health visit they make. This latter contract provides rich information on the amount of commitment that individuals desire, shedding light on underlying preferences. If individuals are sophisticated, such personalized contracts should boost take-up and health care usage. In contrast, if individuals are partially naive, allowing them to choose their level of commitment could lead them to make systematic mistakes, choosing commitment amounts too small to ensure follow-through.

We also offered a contract which can potentially assist naive and partially naive consumers by enticing them to choose (stronger) commitments, by bundling the CC with a price discount. The intuition is that a naive consumer, who mispredicts that he will visit the doctor without commitment, might sign up for such a contract simply to avail themselves of the price discount. Ex post, such a consumer might follow through only due to the marginal incentives provided by the CC. For comparison, we can compare the performance of such a subsidized CC with simply offering (uncommitted) price discounts for the health camp visits. This discount-only condition comprises our final treatment.

These different contract options were offered to random subsets of participants, allowing us to credibly assess impacts on preventive health care utilization as well as later health outcomes, including blood pressure itself as well as diet and exercise behaviors that could reduce hypertension, all of which were collected in the endline survey six months later. The study benefited from a rich data environment, combining longitudinal individual survey data (collected at baseline and endline) on a range of health attitudes, behaviors and outcomes, together with health-camp administrative data on actual usage of preventive health visits over the six month follow-up period, for over 1,700 participants. The study featured close collaboration with the Provider, and the intervention was successfully carried out with minimal logistical disruptions or design contamination. We pre-specified the reduced-form econometric tests in a pre-analysis plan on the AEA RCT registry (#AEARCTR-0000062), and in a methodological innovation of the current study, also pre-specified the theoretical model that forms the basis of the structural estimation described below.

The empirical results indicate that the CC interventions were disappointing in terms of boosting preventive health care usage and improving health outcomes. In the non-subsidized treatment arms, take-up of personalized and fixed commitment contracts was just 13.7 and 14.1%, respectively, mirroring the relatively low take up of other health CC's in several recent studies (Laibson 2015; Dupas and Miguel 2017). This low level of demand suggests that relatively few people in this

high-risk population both value the hypertension health camp treatment *and* believe that CC's will help them overcome present bias. The level of commitment chosen in the personalized CC arm was generally low (lower than in the fixed contracts), with most respondents choosing to pre-pay just the health camp attendance fee with no additional incentive to be returned at the visit.

Actual attendance at the health camps was also disappointingly low, at less than 10% making at least one visit (out of the recommended three over six months) for those not receiving discounts, with unconditional attendance rates nearly identical and not statistically distinguishable in the fixed CC (9.5%), personalized CC (9.9%), and control (8.9%) groups. Take-up of the CC's was considerably higher in the group offered a 50% price subsidy, especially in the discounted personalized CC arm, where it reached 38%. Yet even in this case, only 13.7% of respondents made at least one visit to a health camp, a rate slightly *lower* than those who were offered a straight half-off price discount but no commitment (14.5%), and only 6.2% attended all three visits that they paid for. Given the low levels of attendance at the preventive health visit, it is perhaps not surprising that we do not find statistically significant treatment effects on respondents' endline health outcomes in terms of blood pressure and body weight.

Across the different CCs, between 62-77% of those who paid for a commitment contract failed to make even one visit to a health camp. Overall, between 8% (in the fixed contract with no discount) and 30% (personalized contracts with discounts) of all individuals chose to purchase a commitment contract but then failed to attend any health camps. Under reasonable assumptions, this suggests that a substantial fraction of the participant population experienced reduced welfare due to the commitment contract offers. Through the lens of our theory, these are likely to be individuals who were partially naive (i.e., moderately over-confident) about the extent of their own time inconsistent preferences. They appear to understand their own present-bias problem enough to demand some commitment, but ultimately purchase too little to actually overcome procrastination.

We develop and estimate a model of consumers with present bias and variation in naivete, providing estimated distributions of present bias, (β , as in Laibson (1997)), and naivete / sophistication ($\hat{\beta}$, as in O'Donoghue and Rabin (1999)) in our population. Given these estimates, we quantify the welfare implications of the various commitment contracts offered in the project, as well as derive counter-factual welfare impacts of alternative contracts and policies.

The structural results imply that a large share of individuals are in fact partially naive about their own time inconsistency problems, rationalizing the large gap between contract purchase and preventive health care usage that we document in the data. In particular, the estimated mean of β is 0.365, while the estimated mean of $\hat{\beta}$ is 0.795. Under the assumptions in the structural model, and adopting a utilitarian perspective, *consumer welfare* is considerably lower in the undiscounted CC treatment arms than in the control group. In theory, bundling subsidies with commitment induces respondents who might not otherwise use CCs to do so, boosting or hurting welfare depending on whether CCs are effective for such individuals. Consistent with the reduced form findings of

low follow-through, we find that consumer welfare (net of the discount) is reduced also in these treatments. In contrast, *social welfare* increases due to modest increases in utilization. This masks losses in consumer welfare due to failed commitment contracts, which create profits for the firm. Finally, we conduct counter-factual simulations using the structural model, to assess whether welfare would be improved by offering CCs with different features. Given the degree of naivete regarding present bias that we estimate in this population, the simulations imply that providing (marginally) greater upfront commitment would lead to even greater consumer welfare losses.

It is important to consider alternative explanations of our results, discussed in detail in Section 4.2. First, might consumers have simply forgotten that they signed up for a commitment contract? To address this concern, we cross-randomized the entire sample to receive reminders of the health camps (and the CCs, if they had signed up) through a home visit two weeks before the health camps ended. These reminders had no effect on health-camp attendance, suggesting that our results are not driven by limited memory or attention. Second, can shocks to the costs or benefits of attending the health camps explain the results? A large majority of those who signed up for commitment failed to follow through, despite having many opportunities (one per week for six months) to attend. Thus, idiosyncratic or mean-zero shocks (such as weather, competing events, or ill health) are unlikely to explain our results. In addition, our endline surveys and continuous monitoring of the experiment revealed no evidence of aggregate shocks to costs of attending, such as political or economic shocks. The average level and distribution of attitudes towards the provider did not change over time, and individual changes were uncorrelated with attendance, providing evidence against systematic negative updating about the quality of the provider and the health camps. A final possibility is that those who signed up did so simply due to social pressure exerted by the enumerators who explained the commitment contracts. We note that the stakes are substantially higher than typically estimated levels of social pressure (e.g. DellaVigna et al. 2012). Moreover, only 12% of those who signed up described persuasion by the enumerators as a reason for doing so in the endline.³

In this study, even with a forceful intervention targeting an at-risk population, deploying commitment, subsidies, information provision, a high-quality provider and home-visit reminders, overall utilization of preventive care remains very low. At the least, our results suggest that commitment contracts are not a panacea for low usage of preventive health care in low income settings, such as rural India. In fact, in our setting there are good reasons to believe that offering these contracts may even reduce consumer welfare. The amounts at stake in our experiment are moderate, on the order of US\$1-4 (25-100% of the daily agricultural wage). However, while it is tempting to think that a “stronger” commitment contract would be more effective (and it may well be), such

³Of course, the under-utilization of preventive care in general likely involves multiple factors, including liquidity constraints and poor information. We conduct a quite detailed information intervention for the entire sample, including measuring hypertension and educating individuals about the condition, making it less likely that a simple lack of information drives our results. In our setting, low information or liquidity constraints can explain low demand for care and take-up of commitment, but not the poor follow-through on commitment.

contracts might have lower take-up due to liquidity constraints and due to the demand for flexibility in the presence of uncertainty (Amador et al. 2006; Karlan and Linden 2017), and also open up the risk that partially naive consumers will be further harmed, if they do not follow through. It remains possible that CCs would have more beneficial impacts in other settings, for instance, where individuals are more sophisticated about their present bias, or where individuals have more experience with the health services provided.

The remainder of the paper is organized as follows. Section 2 discusses the empirical setting of the study, the research design, and the data. Section 3 presents the main experimental results on contract take-up, preventive health care utilization, and health outcomes. Section 4 lays out a theoretical model of commitment contracts and consumer health care utilization, and discusses estimation issues, while section 5 describes the structural estimation results and counter-factual simulations. The final section concludes and discusses implications for health policy.

2 Study Setting and Research Design

Hypertension, otherwise known as high blood pressure, is one of the most prevalent chronic illnesses. In 2008, approximately 40% of adults aged 25 and over had been diagnosed with hypertension worldwide, and the condition accounted for at least 9.4 million deaths globally each year (World Health Organization 2013). In low-income countries such as India, where the public health system is plagued by low service quality (Banerjee et al. 2004), the disease burden from hypertension is especially high. According to the Association of Physicians of India, the prevalence of hypertension over the past six decades has grown almost 13-fold nationally in urban areas and almost 8-fold in rural areas. This has meant that among individuals aged 45 year and older, 45% are considered to be at risk of the disease (API, 2013a).

This study was carried out in four rural villages in the state of Punjab in Northern India. The study was implemented in partnership with an organization that delivers primary medical care services and clean drinking water to rural markets using community health clinics. In particular, the Provider conducts “Hypertension Day” health camps wherein an experienced doctor from a nearby city visits each village every week to treat hypertension patients.⁴ The consultation fee to see the doctor during these weekly clinics is Rs. 30 (excluding the cost of medicines and lab tests).⁵ During the visit, the doctor takes health measurements (blood pressure, BMI, and waist circumference), provides the patient with information about hypertension, and prescribes an appropriate treatment plan. The doctor also encourages the patient to make dietary and lifestyle changes such as decreasing salt intake and maintaining a healthy weight.

⁴Note that all individuals, even those without hypertension, are able to see the doctor during the camps. However, hypertension patients receive priority given that the program was launched specifically to address high blood pressure.

⁵As a point of reference, Punjab’s current legal minimum wage for agriculture is Rs. 250 per day, or roughly 5 USD, thus the consultation fee is worth roughly one hour of labor at this rate.

The standard medical advice, for individuals who either have been diagnosed as hypertensive or are at high risk of developing the illness, is to undergo regular monitoring of their condition through bi-monthly consultations with a doctor (API 2013b). Despite the large potential benefits, very few patients adhere to this recommendation. Combined with a commonly expressed desire to manage their condition, a leading candidate explanation for this lack of follow-through is time inconsistency or present bias. To test this hypothesis, randomly chosen subsets of individuals from our study sample were offered various types of commitment contracts, as well as price discounts, designed to lower the marginal cost of future consultations.

2.1 Sample Selection

Since the health camps are targeted towards patients with high blood pressure, our study sample consists of individuals above the age of 30 who either have hypertension or are at high risk of developing the condition. We follow widely accepted medical guidelines and define hypertensive patients as those with systolic blood pressure above 140 or diastolic blood pressure above 90.⁶

To identify such individuals, in 2012 we carried out a census in the four villages where the weekly camps were to be held, during which a team of enumerators first screened all members of a particular household by taking their blood pressure readings using an automatic blood pressure measurement device (see Figure 1).⁷ If the systolic or diastolic blood pressure reading is above the thresholds previously described, the enumerator immediately invited the individual to participate in the study and to complete the baseline survey. In the event that more than one household member had hypertension, the member with the more severe stage of the condition was invited to take part in the study. Furthermore, in the event that more than one household member was at the same stage (i.e., Stage 1 or Stage 2), the member with the highest systolic blood pressure reading was invited. Finally, non-hypertensive but high-risk individuals were also identified using a score algorithm based on age, gender, family history of hypertension and diabetes, tobacco use, physical activity, and waist circumference.⁸

Across the four sample villages, a total of 20,824 individuals from 4028 households were screened in the census activity. From this initial pool, 2004 households with at least one hypertensive member and an additional 276 households with at least one high-risk member were selected for the study, yielding a total of 2280 households. Of these 2280 households, 1725, or 75.7%, accepted our invitation to participate in the study. The main sample for this paper thus consists of these 1725 individual respondents who completed the baseline survey (Figure 1).

⁶Both the Association of Physicians of India and the NIH define hypertension in this manner.

⁷Enumerators were trained in operating the device, the Citizen CH-452 model, which has been validated by the ESH protocol and was selected for the project in consultation with a local medical doctor.

⁸This 100-point hypertension risk score algorithm is based on the current literature and was developed in consultation with doctors at the Provider.

2.2 Research Design

Immediately after administration of the baseline survey, the respondent was offered a commitment contract or discount coupons (or both) to visit the weekly health camp in their village for 3 times in 6 months, a frequency in line with the medical guidelines described above. We randomized the type of contract offered to each household, stratified by hamlet (a geographic cluster within the village which is relatively ethnically homogeneous) and household head's education, by using a computer random number generator prior to the enumerator's visit. Specifically, households were either offered a fixed contract, a personalized contract, or no contract, all with equal probability:

Group 1: No Commitment Contract This group was not offered any commitment contracts. Each respondent only received information about managing hypertension and a flyer with the times and location of the health camps, which were provided to all participants in the study irrespective of treatment arm.

Group 2: Fixed Commitment Contract This group was offered a commitment contract for 3 visits to the health camps during a 6-month period. As part of the contract, the respondent was required to pay in advance for all 3 doctor visits (Rs. 90, or Rs. 30 per visit). The respondent was also asked to pay an additional commitment amount of Rs. 30, which she receives back in equal installments of Rs. 10 each time she visits the doctor. In other words, the respondent pays a total of Rs. 120 up front, and receives Rs. 10 on each of the 3 visits. See Figure 2 for a graphical illustration.

Group 3: Personalized Commitment Contract While the commitment amount is fixed at Rs. 30 in Group 2, respondents in Group 3 can choose their own commitment amount beginning as low as Rs. 0.⁹ As above, the respondent receives this amount back in 3 equal installments every time she visits the doctor. The respondent is also required to pay in advance for 3 visits, so the total upfront payment is Rs. 90 for consultation fees plus the selected personalized commitment amount.

Each of these three groups was cross-cut with a price discount treatment in order to compare the effectiveness of commitment contracts with that of simple price incentives. Specifically, treatment individuals received 3 coupons that entitled them to half-price consultations (Rs.15 instead of Rs.30) during the 6-month program period. To sum up, our baseline randomization yielded a total of 6 arms (5 treatment and 1 control), as shown in Figure 1.

The bundling of commitment contracts with discount coupons also has the useful feature of potentially making the former more attractive to naive individuals. In particular, time-inconsistent

⁹In practice, the respondent's chosen commitment amount is rounded up or down so that it is divisible by 3.

respondents who do not regard themselves as suffering from such issues may be persuaded to sign up for a commitment contract in order to take advantage of the price incentives. These individuals could then potentially benefit from the contracts later on, through increased attendance at the health camps.

Respondents in the commitment contract groups could sign up for their respective contracts in several ways. First, they could accept the contract on the spot with the enumerator during the baseline survey, who subsequently collected payment. Second, respondents could sign up with the Provider’s village health workers (VHW) and health coordinators (HC), both of whom were well-known in the village since they often go door-to-door to assess the community’s health needs. Specifically, around 3 to 4 days after the enumerator offered the commitment contract to a particular household, the VHW and HC visited households who had not yet signed up for the contract. The VHW and HC then asked these respondents whether they would like to take up the contract on offer, as well as reminding them about the health camp schedule. Note that the VHW and HC visited all households in the study to remind them about the camps, including those in the control group, to hold constant any effect the VHW and HC’s visit may have. Lastly, respondents in the commitment contract groups were also able to sign up for the contracts directly at the clinic at any time during the course of the study.¹⁰

A final set of treatments were implemented two weeks before the conclusion of the 6-month program. In each village, half of the respondents were randomly selected to receive a short reminder about the hypertension camps. These respondents were personally visited by our team of enumerators, and were informed that there were 2 weeks left until the contracts or coupons would expire, if applicable. The other half of the respondents served as control, and did not receive the reminder. This intervention was designed to test whether inattention is an important factor in dampening attendance and health care utilization.

3 Experimental Results

In this study, we use three main data sets in the analysis. First, a baseline survey was conducted prior to presenting the treatments to respondents. This survey collected information on respondent and household characteristics, as well as the respondent’s health status, health-seeking behavior and knowledge, as well as time and risk preferences.

¹⁰Although the hypertension camps are only held once a week, the clinic is open Mondays through Fridays to sell medicine and conduct lab tests. Each respondent could sign up only for the commitment contract she was originally offered. While both the commitment contracts and price discount coupons covered 3 health camp visits, respondents were given the opportunity to renew these contracts and coupons at the clinic for the remainder of the 6-month program. These renewals were described to respondents when the contracts and discount coupons were initially introduced by enumerators. In the case of commitment contracts, for example, respondents who completed 3 visits in the first 2 months of the program could take up another commitment contract for 3 visits in the remaining 4 months. Similarly, for discount coupons, respondents who used up all 3 coupons in the first 2 months of the program could ask for another set of 3 coupons, which were valid for the remaining 4 months.

Second, we collected data on attendance at the hypertension camps for all study participants. For the 6-month period in which commitment contracts and discount coupons were valid, a member of our field staff was present during the weekly camps in each village to record the household ID number and names of all study participants who came to see the doctor. Furthermore, we collected such attendance data for one month after contracts and coupons expired, which allows us to examine treatment effects in a setting where commitment contracts were no longer available.

Finally, an endline survey was conducted in each village one week after the contracts and coupons expired. This survey asked questions similar to those in the baseline. In addition, it asked information on the doctor visits and other health care providers. This survey also included the respondent's weight and waist circumference measurements, as well as self-reported dietary and exercise changes.

3.1 Summary Statistics

Baseline characteristics for our sample are shown in Table 1 (Panels A and B). Our respondents come from households that have 5.5 members on average, with a mean annual household income of Rs.102k. 59% of our sample is female. Among household heads, the most common occupation is self-employment in agriculture (at 37%) and 45% can both read and write.

A large portion of our respondents also have characteristics that place them at risk for hypertension. For instance, the average age in our sample is 53.7, and the risk of high blood pressure increases with age. 51% of our sample is overweight, defined as having a Body Mass Index (BMI) of over 25.

Despite the relatively poor health status of individuals in our sample, few respondents visit a doctor for reasons other than acute illnesses. Based on self-reports at baseline, while 76% say that they "always" or "frequently" seek health care when they are feeling sick, only 6% of respondents visit the doctor for preventive care. While 71% of our sample reported to knowing they had hypertension, only 50% of these individuals are currently taking medication to manage their hypertension. These summary statistics suggest that behavioral barriers may play a significant role in explaining the low demand for health services in our setting. Indeed, a majority of respondents (74%) strongly agree that they are often impatient.

A randomization check of our commitment contract and discount treatments do not show any systematic, statistically significant differences across key variables (Table A1).

Before discussing our structural estimates of the model from Section 3, we first present experimental findings for (i) the take-up rates of commitment contracts, (ii) utilization of health care services, and (iii) health outcomes and behavior.

3.2 Contract Take-Up

Table 2 (Panel A) reports commitment contract take-up by treatment group. A number of patterns emerge. First, take-up without a discount is 13.7% (39 out of 284) for the “fixed” contract, and 14.1% (40 out of 283) for the “personalized” contract. So at the undiscounted price point, restricting the flexibility of the contracts does not reduce demand.

Second, discounts have a substantial impact on take-up: for the fixed contract with discount, 25.9% (72 out of 278) take up while for the personalized contract 38.6% (112 out of 290) take up. These results suggest that both (i) discounts have a marked impact on take-up and (ii) that the personalized contract and the discount are complementary. People are more likely to take up the contract with discount when they have the option to specify the amount that they commit. In general, those who take up the personalized contract choose a lower commitment amount than that specified by the fixed contract (see Figure 5), which suggests that many consider the fixed contract to be “too strong”.¹¹ The pattern of results is consistent with the idea that a bundled discount nudges consumers with low demand for commitment to sign up, particularly when they can choose small commitment amounts.

Table 3 (Panel A) presents results on contract take-up for a particular sub-sample. This “ideal” sample is comprised of individuals who both believe “it is possible to be healthy with hypertension if blood pressure is frequently monitored” *and* who trust the service provider. Clearly our contracts should be more attractive to these individuals, compared to those who think otherwise. There is indeed evidence consistent with this hypothesis. In particular, contract take-up is uniformly higher for the 439 respondents in this sub-sample. For instance, 19% of such individuals took up the undiscounted personalized contract, compared to 14% in the overall sample. Similarly, 49% of “ideal” sample individuals took up the discounted personalized contract, compared to 38% in the overall sample.

In terms of gender heterogeneity, Figure 4 presents the take-up rates differentially for our male and female respondents respectively. While there are no discernible differences in their take-up of the undiscounted commitment contracts, men did have significantly higher take-up rates of the discounted contracts than women in our sample. For instance, while 35% of our male respondents signed up for a fixed contract with discount, only 20% of our female respondents did so.

3.3 Doctor Visits

The next step in our analysis is to assess the effect of being offered a commitment contract or discount coupon on the attendance at health camps. As Figure 3 makes clear, we find a uniformly low rate of attendance across all treatment groups, varying between 5 and 10% of the recommended number of visits during our study period.

¹¹Over 80% of participants who signed up for a personalized contract (both with and without discount) opted to pay only for the 3 consultations upfront, without any additional commitment.

In particular, being offered a commitment contract by itself does not significantly increase health camp attendance on average. Compared to the 4.6% proportion of attendance in the control group, those offered the full-price fixed commitment contract made 5.4% of the recommended visits. Attendance for those offered a full-price personalized contract is a similar 6.6%. Furthermore, bundling discounts with commitment contracts does not increase attendance relative to simple price discounts. The patterns are identical if we measured utilization using a dummy for visiting at least once instead of the proportion of visits.

Analyzing health camp visits in the “ideal” sample, as defined above, indicates that our interventions may be slightly more effective for those in this latter group. Specifically, both the simple discount and the discounted personalized commitment contract treatment led to discernible increases in utilization when compared to the control group. As shown in Table 3, these effects can be sizable - among those in the “ideal” sample, 19.5% of individuals offered a personalized contract with discounted consultations went to the health camp at least once, compared to 6.5% of control individuals.

Another result is that conditional on signing up for a commitment contract, consumers in contracts bundled with discounts do worse than those in contracts without discounts. Figure 3 shows that consumers in discounted contracts (Panels E and F) are more likely to fail to visit the doctor. This is because some partially naive consumers may be attracted by the discounted consultations, but the commitment isn’t strong enough to increase their attendance. As such, there are more individuals losing out when they are offered a discounted contract.

Unlike the differential take-up rates of discounted contracts among men compared to women noted above, there are no differences between our male and female respondents in their attendance at the health camps. This is consistent with male individuals being more over-confident than their female counterparts.

Finally, the low rate of health camp attendance cannot be explained by simple inattention. This is because our randomized reminders had no effect on doctor visits across all treatment groups (Figure 6).

3.4 Health Outcomes

As part of the endline survey, we collected data on key health outcomes such as blood pressure (both systolic and diastolic) and body weight, as well as health-seeking behaviors such as exercising. Table 2 (Panel B) and Figures 7 and 8 present the average treatment effects for these outcomes.

Given the earlier results on doctor visits, it is perhaps not surprising that our treatments failed to significantly improve health outcomes. For instance, the proportion of respondents with hypertension at endline ranged from 49% in the control and discount groups to 54% in the personalized CC group, a statistically insignificant difference. The same pattern also holds for other measures of health (e.g. pre-hypertension, overweight and obesity status).

In terms of treatment heterogeneity (Table 3, Panel B), the estimated coefficients for the “ideal” sample are similar to those for the main sample. Therefore the experimental treatments did not lead to significant improvements in health outcomes even among individuals most likely to benefit from them in our sample.

3.5 Interpretation and Alternative Explanations

We interpret these results through the lens of a simple model of preventive health investments in the presence of limited self-control and partial naivete. We discuss this interpretation in much greater detail in Section 4, where we lay out the model, and Section 5, where we discuss the related structural estimates. We also consider alternative explanations, including limited memory or attention, uncertainty or shocks to the costs or benefits of attending, learning about the health camps, and social pressure to accept the contracts, in Section 4.2. In short, we find little evidence to support the alternative explanations. Instead, the evidence is consistent with individuals having some demand for commitment due to limited self control, but under-appreciating the strength of the incentives required to ensure that their future selves follow through on their commitment.

3.6 Simple Welfare Effects

Combining our results on contract take-up and frequency of doctor visits, we can document for each treatment group the required benefits of each visit in order to offset the losses incurred from purchased, but unattended, visits, adopting a utilitarian perspective. For instance, in the fixed CC group, a total of 47 visits were made, while 117 were purchased (at a price of Rs. 30 each). Recalling the standard commitment amount of Rs. 10 per visit, this implies that each realized visit needs to be worth at least Rs. 97 for consumer utility to be non-negative. Similarly, in the personalized CC group, a total of 57 visits were made, while 120 were purchased (at a price of Rs. 30 each). The vast majority of people who signed up for a personalized contract opted for zero additional commitment (Figure 5), which implies that each realized visit needs to be worth at least Rs. 63 for consumer utility to be non-negative among this group.

We can compute similar statistics for the combined CC and discount treatments as well. In the fixed CC plus discount group, for instance, a total of 52 visits were made, while 216 were purchased (at a price of Rs. 15 each). Given the fixed commitment amount of Rs. 10 per visit, each realized visit needs to be worth at least Rs. 110 for consumer utility to be non-negative. Finally, in the personalized CC plus discount group, a total of 78 visits were made, while 339 were purchased (at a price of Rs. 15 each). Once again, given the large proportion of participants opting for zero additional commitment, this implies that each realized visit needs to be worth at least Rs. 65 for consumer utility to be non-negative for this group.

While useful as a benchmark, these simple welfare calculations do not allow us to study the distribution of welfare changes as we move from one treatment to another. Nor do they allow us

to study the impacts of counterfactual contracts. To further investigate these welfare effects, we structurally estimate a simple, pre-specified, commitment model, which we now describe.

4 Commitment Model

It is clear from our analysis thus far that (i) a range of consumers demand commitment when offered to them, (ii) their take-up of commitment contracts is responsive to price discounts, and (iii) that many consumers do not follow through on their commitments once made. One explanation for why consumers demand commitment, and then do not follow through, is that consumers are present-biased but partially naive about that bias, in the sense that they underestimate how it will impact their future behavior.

In this section, we set up and estimate a micro-founded model of consumer demand for both commitment and preventive health services, with an emphasis on the potential roles that self-control and limited sophisticated about self-control play in these choices. Our primary goals are to (i) estimate key time preference, sophistication, and demand parameters and (ii) study the implications of those estimates for the effects of commitment contracts in our setting, and their welfare properties. Particularly, we use our estimates to study the welfare impacts of (marginally) more aggressive commitment contracts, which, in theory, have an ambiguous impact on consumer contract take-up, a positive impact on follow-through by consumers conditional on contract take-up, but potentially a negative impact on consumer *welfare* if many consumers still do not follow through conditional on contract take-up.

It is important to note that the model presented here is tailored to parsimoniously match the specific aspects of our empirical environment. The model focuses on the mechanisms of self-control and affordability as reasons for why consumers might under-consume preventive care. See our discussion earlier in Section 1 and later in this section for additional potential explanations for the consumer behavior we observe.

4.1 Model

We model consumer demand for commitment and preventive health care with three time periods, reflecting the key decisions that consumers make in our environment:

- **$t = 0$:** Consumers choose whether to enter a commitment contract if offered one, depending on the randomized intervention. If offered a personalized commitment contract, they also choose their commitment level if they enter a commitment contract.
- **$t = 1$:** Consumers choose whether or not to go to the doctor for recommended treatment. In our empirical setting, consumers have many opportunities to attend in the relevant time window. In

the model, for simplicity and tractability, we collapse these into a single time period and decision. Their cost / benefit of attending depends on whether they entered into a commitment contract or not, what the features of that contract are, and what their non-pecuniary costs of going are.

- **t = 2:** The health benefits of attending the doctor (or not attending) at $t = 1$ are realized.

We describe the model working backwards from $t = 2$ to $t = 0$. At $t = 2$ we assume that the benefit from treatment at $t = 1$ (relative to no treatment), denoted b_i is distributed normally conditional on observed variables X_i :

$$F(b_i|X_i) = \mathcal{N}(\alpha_b + \kappa_b X_i, \sigma_\epsilon)$$

Here, we allow for observable heterogeneity X_i to impact the mean benefit, with unobserved heterogeneity independent of X_i with variance σ_ϵ^2 .

At $t = 1$ consumers decide whether or not to go to the doctor for recommended treatment. We model consumer attendance as a binary decision to attend or not attend. In our descriptive analysis, most consumers who attend do so once in the six month time period over which we evaluate them, though some attend twice and others the recommended three times. In our structural model, we consider a consumer to have attended if they visited the doctor at least once during the study period, and assume that they realize the full benefit of treatment b_i in period two.¹²

Consumers' utility from going to the doctor during the six-month experimental timeframe ($t = 1$) is:

$$\begin{aligned} U_{i,Attend} &= \beta(X_i'')\delta b_i - C(X_i') - p_i \\ p_i &= f_i - d_i \text{ if no commitment contract} \\ p_i &= -m_i \text{ if commitment contract} \end{aligned}$$

Here, p_i is the marginal payment consumers make to attend at $t = 1$. In our environment, because consumers may have previously entered a commitment contract at $t = 0$, this price is person-specific and may be negative, i.e. consumers may be paid to go to the physician at $t = 1$. We break down p_i into three relevant components: (i) f_i , the standard per visit fee for a consumer with no commitment contract (ii) d_i , a per-visit discount given to consumer i (known in advance of the visit) and (iii) m_i the incremental payment a consumer receives above and beyond the typical per-visit fees if they entered into a commitment contract. We assume that monetary costs and rewards in period $t = 1$ translate into losses and gains in consumption utility in the same period – a plausible assumption given the small sums of money, and the fact that $t = 2$ represents the distant

¹²An alternative specification could explicitly model whether consumers make one, two, or three preventive visits during the study period, and make benefits b_i a function of the number of visits.

future where health benefits are realized. The $t = 1$ decision also depends on $C(X'_i)$, consumer i 's non-financial costs of going to the doctor. Empirically, we allow this cost to depend on a range of observable variables X'_i such that $C(X'_i) = \alpha_C + \kappa_C X'_i$.

In our environment, when a consumer enters into a commitment contract, they never pay when going to the doctor at $t = 1$ and typically receive money in return.¹³ Given the specifics of our environment, the possible p_i for consumers at $t = 1$, denoted in rupees per visit, are:

$$p_i = \begin{cases} 30 & \text{if Control} \\ 15 & \text{if Discount Only} \\ -15 & \text{if Fixed Contract} \\ -m_i & \text{if Personalized Contract} \end{cases}$$

In the personalized contract, m_i is restricted to be greater than 0, such that if consumers enter into a personalized contract they either go to the doctor for free or receive a payment when going to the doctor. See Section 3 for more details on the distribution of m_i for those taking up personalized contracts.

We allow consumers to potentially be present-biased through the parameter β , which reflects the extent to which consumers more heavily discount all future periods relative to the present period. We allow β to vary across individuals, depending upon a set of observable variables X''_i and an unobservable component ς_i such that:

$$F(\beta_i^* | X_i) = \mathcal{N}(\alpha_\beta + \kappa_\beta X''_i, \sigma_\varsigma) \quad (1)$$

$$\beta_i = \max[\min[\beta_i^*, 1], 0] \quad (2)$$

A consumer chooses to visit the doctor if and only if $U_{i,Attend} \geq 0$ at $t = 1$. At $t = 1$, when consumers decide whether to go to the doctor or not, they are potentially subject to present bias since they have to incur the cost of visiting the doctor now, but don't receive the benefits from treatment until out in the future at $t = 2$. At $t = 0$, when consumers choose whether or not to enter into a commitment contract (if they have the option to do so) both $t = 1$ and $t = 2$ are in the future, so there is no issue of present bias moving between those latter two periods from the perspective of a $t = 0$ consumer. If they do not fully perceive the extent of their present bias, then at $t = 0$ they may overstate the likelihood they will actually follow through and visit the doctor at $t = 1$.

We now turn to the consumer decision of whether to enter into a commitment contract at $t = 0$. Consumers in our control group and in the treatment group that only receives a per visit discount

¹³This is why price is represented as a negative of an incremental commitment amount above per visit fees in our model: any case when a consumer signs up for a commitment contract involves them paying the entire standard per visit fees up front.

have no decision to make at $t = 0$ since they have no opportunity to enter into a commitment contract. Consumers who are offered a commitment contract decide whether or not to take up that contract at $t = 0$, based on their perceptions of their present bias and whether or not they believe the commitment contract will allow them to overcome present bias that prevents them from going to the doctor at $t = 1$.

Crucially, at $t = 0$, consumers' *perceptions* of their present bias ($\hat{\beta}$) are what matters for contract take-up, not their actual present bias. Consumers choose to take up the fixed commitment contract, where they commit the fees for three visits plus an additional 45 rupees up front (15 rupees per visit), if they perceive that the commitment has positive value. In the treatment group with no per visit discount, this occurs when a consumer believes at $t = 0$ that (i) they want themselves to go to the doctor at $t = 1$ (ii) they will not go at $t = 1$ without additional commitment and (iii) the commitment amount in the contract is strong enough to induce them to go at $t = 1$.¹⁴

Formally, a consumer with a fixed contract offer with no lump sum discount will take-up that offer if and only if the following conditions hold:

$$\delta b_i \geq C_i + 30 \tag{3}$$

$$C_i + 30 \geq \hat{\beta}(X_i''')\delta b_i$$

$$15 \geq C_i - \hat{\beta}(X_i''')\delta b_i \tag{4}$$

The conditions for taking up a fixed commitment contract with a discount are similar to those for a fixed contract with no discount, but because the discount is bundled with the commitment contract in our discount treatments, the condition is simply that someone believes they will go to the doctor at $t = 1$ with the contract (without the discount, they will not take up the commitment device if they think they will go to the doctor at $t = 1$ without the contract). Thus, they will take up the fixed discount contract if the following two conditions hold:

$$\delta b_i \geq C(X_i) + 15 \tag{5}$$

$$15 \geq C(X_i) - \hat{\beta}(X_i''')\delta b_i \tag{6}$$

If a consumer is in a treatment where they are offered a personalized commitment contract without a bundled up front discount, they take-up the contract if equation 3 is satisfied. This is true because

¹⁴It is important to note that the simplifying assumption in the model of a binary follow-through choice likely causes us to underestimate consumer losses from signing up for a commitment contract in our setting. The model does not allow for partial follow-through to commitment, which would involve introducing consumer learning or a stochastic cost of attendance component. As a result, consumers who attend once are counted as having followed-through, though they may experience higher losses in practice by paying for second and third visits that they don't follow through on. Since many consumers choose commitment but don't follow through at all, the model arguably approximates our observed environment well despite this simplification.

any demand for commitment is sufficient to take up a personalized contract, given that they can choose a commitment amount greater than or equal to 0.¹⁵ If they take up the contract, they commit a per visit amount m_i up front (on top of normal fees) such that:

$$m_i = \max[C_i - \hat{\beta}(X_i''')\delta b_i, 0] \quad (7)$$

Note that we assume that consumers choose the smallest permitted commitment amount they perceive as being necessary to ensure follow-through. This could result from, say, a small (unmodeled) liquidity cost. However, since consumers anticipate receiving this commitment amount back upon attending at $t = 1$, we assume that they do not cut back on other consumption due to the temporary hand-over of the commitment amount m_i . For the personalized contract with a discount, a consumer accepts if the equation 5 above holds, and, if so, they commit the same per-visit amount that is equal to what they would choose without a discount, as described above.

$\hat{\beta}$ represents consumers' *beliefs* about the degree of present-bias they will have when making the decision of whether to go to the doctor at $t = 1$. As is typical in the literature, we assume that $\beta \leq \hat{\beta} \leq 1$. When $\beta = \hat{\beta}$ a consumer is *sophisticated* about their present-bias, i.e. they exactly perceive the extent to which they will be present-biased at $t = 1$. When $\beta < \hat{\beta} = 1$ a consumer is *fully unsophisticated* or *fully naive*: they think they will have no present-bias in the future, though they actually will. When $\beta < \hat{\beta} < 1$ consumers perceive some but not all of the present-bias they will have at $t = 1$ and they are said to be *partially sophisticated* or *partially naive*.

A key goal of our empirical analysis is to estimate the joint distribution of β and $\hat{\beta}$. If most consumers are present-biased and sophisticated about those biases then commitment contracts are likely to increase consumer welfare by satisfying demand for commitment and increasing take up of preventive health care. If consumers are not present-biased, then there is no need for commitment and contracts will be welfare neutral, or welfare negative if they are costly to implement. If most consumers are present-biased but are partially or fully naive, then commitment contracts could reduce consumer welfare by getting consumers to commit money up front, but then lose that money when they don't follow through on their commitments. We are unaware of prior work estimating the joint distribution of present-bias and perceptions about present-bias in field data, nor do we know of work that then goes on to use these estimates to explicitly derive the welfare implications of a commitment contract product.

Our primary specification allows for unobservable heterogeneity, in addition to observable heterogeneity. We parametrize unobservable heterogeneity in $\hat{\beta}$ with a two parameter Beta distribution. Specifically, we assume that, across consumers in the population, the distance that $\hat{\beta}_i$ is between β_i and 1 is distributed Beta with parameters τ_1 and τ_2 . Formally:

¹⁵Even deposited amounts smaller than the show-up fee would, in practice, serve as a (weaker) commitment device. In practice, for logistical reasons, we could not implement such smaller commitment amounts, and required a minimum deposit equal to the consultation fee even in the personalized contract.

$$\hat{\beta}^* \rightarrow \mathcal{B}(\tau_1 + \kappa_{\hat{\beta}} X'''_i, \tau_2)$$

$$\hat{\beta}_i = \beta_i + (1 - \beta_i)\hat{\beta}^*$$

This statistical specification for heterogeneity in $\hat{\beta}$ flexibly allows for $\hat{\beta}_i$ to vary between β and 1 while remaining within those boundaries for each consumer. In an alternative specification that we investigate for robustness, we also allow for the distance that $\hat{\beta}^*$ is between β and 1 to be correlated with β to add flexibility in modeling the behavior of low β consumers relative to high β consumers.

Our primary specification includes five observable shifters in X , impacting the benefit from treatment b_i . These include (i) our hypertension severity index (ii) our general health index (iii) gender (iv) whether the consumer was already taking hypertension medication prior to the study and (v) whether they are literate. Gender is included as the main shifter in X' and X''' impacting β and $\hat{\beta}$ respectively. We also include an indicator of whether the consumer is employed or not. We present a range of specifications where different observable shifters impact different parameters, for comparison to our primary approach.

4.2 Discussion and Alternative Mechanisms.

The model makes a number of assumptions and simplifications in order to focus on the key issues of self-control, commitment, and affordability. It therefore excludes a number of potential alternative explanations of the results we presented in the previous section. We discuss these assumptions and the plausibility of the alternative channels below.

Uncertainty and learning about benefits. First, we assume that consumers are fully informed about the potential health benefits of preventive health care both when choosing commitment at $t = 0$ and when choosing attendance at $t = 1$. In our experiment, we implemented an extensive information campaign to all consumers at baseline, including those in the control group, to mitigate the role of limited information. That said, consumers in our environment are often unfamiliar with, and potentially suspicious of, western-style medical care, and the health provider we work with was relatively new to the study area. This could explain low overall demand for doctor visits, as well as low demand for commitment (if consumers would not like to attend even in the long run). However, it would only explain failure to follow through on commitment contracts if consumers update negatively about the benefits over time, e.g. by learning from the early attendees that the quality of the provider is lower than expected (or the costs of attending are higher than expected). To measure this, we included baseline and endline survey questions asking the respondents to report their trust in and satisfaction with the provider. We find no change in the average level or distribution of reported trust in the provider over time, and individual changes over time are uncorrelated with attendance. Nor do those who fail to follow through report reductions in trust

over time. Only 6.5% of those who failed to follow through reported dissatisfaction with the Provider as a reason for not attending in the endline survey.

Social pressure. A second related assumption is that consumers do not purchase a commitment contract simply due to social pressure exerted by surveyors – with no intention of actually attending. Scripts were carefully designed and implemented to explain what a commitment contract is, while maintaining neutrality about whether a consumer should purchase a contract or not. The approximately \$3 required is a meaningful share of the daily minimum agricultural wage (\$4) and of average daily household income (\$5.6)– a much higher fraction than typically estimated social pressure costs (DellaVigna et al. 2012). Only 12% of consumers mentioned such social pressure as a reason for signing up when debriefed during the endline survey.

Uncertainty in costs of attending. A third assumption is a lack of uncertainty in the environment, especially in the costs of going to the doctor. We assume for simplicity that the costs (including the non-pecuniary costs) are fixed and known in advance. Suppose, instead, that consumers understand their average cost of going to a doctor, but idiosyncratic shocks may increase or decrease their costs on a given day. We first note that, in practice, individuals have six months of weekly opportunities to attend, and are largely retired or self-employed adults. Moreover, a substantial majority of those who sign up for commitment never attend. Thus, mean-zero idiosyncratic weekly shocks are very unlikely to explain our results. Instead, could aggregate shocks at the village level, or permanent shocks affecting a large share of individuals be important? For example, unexpected crop failure might have induced a large share of individuals to temporarily migrate for work. Our endline survey finds no evidence of such aggregate shocks. Nor was this period marked by any large political or economic shocks in India or in Punjab.

Altogether, we find little evidence in favor of these alternative explanations for the high rates of failure of commitment contracts in our context.

4.3 Identification.

We now discuss econometric identification of the model, given these underlying assumptions. The key parameters to separately identify are (i) α_β , κ_β , and σ_ς for β (ii) τ_1 , τ_2 , $\kappa_{\hat{\beta}}$ for $\hat{\beta}$ (iii) α_b , κ_b , σ_ϵ for b (iv) α_C , κ_C for C and (v) δ .

The basis for our identification is the randomization of consumers into the six different experimental arms, including the control group and five treatments. Our identification arguments depend crucially on the multiple different treatments along with the experimental randomization, which provide cross-sectional variation in commitment contracts and pricing that impacts the distribution of decision paths.

First, we discuss identification of β from $\hat{\beta}$. As a result of our experimental randomization, the joint distribution of β and $\hat{\beta}$ for consumers in each treatment will be the same as sample size goes to ∞ . For fixed values of the other parameters, different combinations of β and $\hat{\beta}$ imply different

sequential decisions for consumers randomized into commitment contract treatments.

We illustrate this in Figure 9. In particular, it shows how the choice to take up a fixed commitment contract or not (when offered) at $t = 0$ identifies regions of $\hat{\beta}$ for a given consumer. Consumers with high $\hat{\beta}$ don't believe they have genuine commitment issues so are unlikely to demand commitment without a corresponding bundled discount. Consumers with medium $\hat{\beta}$ will demand commitment when they intrinsically value recommended medical care, because it will help them overcome their procrastination at $t = 1$. Consumers with low $\hat{\beta}$ will not take up a fixed contract, even if they want commitment to consume medical care, because they perceive the fixed commitment amount will not be enough to get them to visit the doctor at $t = 1$.

Figure 9 also examines what different sequences of decisions imply for $\hat{\beta}$ and β jointly. For example, Region 2 with medium $\hat{\beta}$ and low β is a region where consumers think commitment is valuable for them, but once they commit money up front they don't follow through on their commitment at $t = 1$ because their β is low. Similarly, Region 5 has medium β and high $\hat{\beta}$, showing consumers who think they will go regardless of whether they have commitment, so don't purchase a contract without a discount, but end up only going if they have a commitment contract. This group of consumers will benefit from being offered a commitment contract paired with a bundled discount, relative to our other treatments.

Crucially, our experimental design generates multiple kinds of variation to separately identify β and $\hat{\beta}$. In addition to partitioning the space of $(\beta, \hat{\beta})$ based on the sequence of choices made when offered a commitment contract (as shown in Figure 9), having cross-cutting treatments that pair discounts with commitment contract offers helps to identify the joint distribution of $(\beta, \hat{\beta})$ by shifting *anyone* who thinks they should go to the doctor into a commitment device (rather than just people who feel they need commitment). Additionally, having the control group and treatment with per visit discounts only helps identify the distribution of $\beta\delta b$ for consumers separately from $\hat{\beta}$. Finally, the personalized commitment contract treatments help identify $\hat{\beta}$ more precisely relative to β because each consumer in this arm makes an individual-specific choice signaling their $\hat{\beta}$.

Taken together, the randomization into the different treatments, paired with the subsequent sequence of decisions consumers make, identifies the joint distribution of $(\beta, \hat{\beta})$ (conditional on other parameters being fixed). It is also important to discuss how the remaining parameters are separately identified. δ is separately identified from β because δ describes relative $t = 1$ and $t = 2$ utility from a $t = 0$ decision perspective, but $\beta\delta$ describes that relative utility from a $t = 1$ perspective.

b is identified separately from β because at $t = 0$, when choosing commitment contracts, consumers choose based on $\hat{\beta}$ and δ , and b but not β . Then, given randomization across different treatments, choices made at $t = 0$ reflect b but not β , as long as consumers are not perfectly sophisticated with respect to their self-control. Cost is identified separately from other factors because cost occurs at $t = 1$, and thus is not indexed by δ , and is only multiplied by $\hat{\beta}$ at $t = 0$ but not by

β at any point.

Finally, it is important to note that δ is not non-parametrically identified separately from b since these factors always multiply each other in consumer utility at all time periods they enter. They are parametrically identified given the assumptions that benefits are linear in observables and have normally distributed unobserved heterogeneity, while we do not allow for heterogeneity in δ . In practice, this means that $\delta * b$ should be thought of as one identified quantity, rather than ascribing a specific proportion of this quantity to δ or b . We treat $\delta * b$ this way in our upcoming results discussion.

4.4 Estimation.

We estimate the model with a smoothed Accept-Reject simulated maximum likelihood methodology that, given the candidate parameters, matches the predicted decision paths for consumers in the population to their actual decision paths [see, e.g., Train (2009) for an econometric discussion and Handel (2013) for another applied example with a similar approach]. Define the set of parameters to be estimated as Θ .¹⁶ For a consumer with set of observables X in treatment T we match their sequence of decisions (which depends on T) to the predicted sequences of decisions for candidate parameters, and choose the parameters with the best match given choices across the control and all five treatments.

There are three types of decisions that could enter the decision path, and hence the likelihood function, for a given individual. The first is choice of commitment contract (if offered). The second is what amount they commit specifically if the contract offered is a personalized contract. The third is whether they visit the doctor or not at $t = 1$. For the random coefficient parameters representing unobserved heterogeneity on each dimension, we take 50 simulated draws, above which the estimation results are stable.

Appendix B provides substantial detail on the likelihood function and estimation process, which we omit here for parsimony. Next, we turn to our results and their implications.

5 Results

Table 5 presents our primary estimates for the model of consumer self-control and demand for preventive health care services. The key parameters are the joint distribution of β and $\hat{\beta}$. We present three specifications: (i) a baseline specification with limited observable heterogeneity (ii) this same baseline specification incorporating correlation between the unobserved components of β and $\hat{\beta}$ and (iii) our primary specification that incorporates observable heterogeneity on a range of potentially important dimensions.

¹⁶In our primary specification, these parameters include $\alpha_\beta, \kappa_\beta, \sigma_\zeta, \tau_1, \tau_2, \kappa_{\hat{\beta}}, \alpha_b, \kappa_b, \sigma_\epsilon, \alpha_C, \kappa_C$, and δ .

All three specifications show relatively similar results for the joint distribution of β and $\hat{\beta}$. Consumers have relatively low β on average, indicating a meaningful degree of present-bias when making the decision of whether or not to visit the doctor at $t = 1$. In our primary specification the average β is 0.365. There is meaningful dispersion in β as well, with an estimated population standard deviation of 0.395 around this mean (incorporating truncation at 0 and 1). 39.8% of consumers have $\beta = 0$ indicating that when it gets to $t = 1$ a meaningful portion of consumers feel that going to the doctor at $t = 1$ is not a valuable activity in the moment. On the flip side, 15.1% of consumers have $\beta = 1$, so no present-bias. Our primary specification also estimates a lower average β for males relative to females, though we find limited gender effects for $\hat{\beta}$. The mean and standard deviation of β are both slightly lower in the baseline specifications.

The estimated mean of $\hat{\beta}$ in our primary specification is 0.795 with a population standard deviation of 0.13, most of which comes from unobservable heterogeneity. Thus, when making decisions at $t = 0$, consumers perceive that they will have some present-bias at $t = 1$ but much less than they will actually have when making the decision to go to the doctor at $t = 1$. As discussed in our section on identification, these estimates are consistent with our earlier empirical results showing that a meaningful portion of consumers demand commitment at $t = 0$ but then don't follow through on those commitments for the most part at $t = 1$. Consumers perceive that they will have some present-bias in the future, and try to correct for that with commitment, but under-predict their degree of present-bias, causing many of them to not follow through on their commitments (and lose money in the process). Figure 10 presents the empirical joint distribution of β and $\hat{\beta}$, which can be compared to Figure 9 and the corresponding discussion of how the sequences of decisions consumers make identify β and $\hat{\beta}$.

Table 5 also presents a baseline specification that allows for correlation in the unobserved heterogeneity terms for β and $\hat{\beta}$. This is implemented by allowing for different distributions of $\hat{\beta}$ conditional on whether β is low ($\beta < 0.5$) or high ($\beta \geq 0.5$). There is only limited correlation in these unobserved heterogeneity components: the mean distance from β to 1 is 47% of the way towards 1 from β for ($\beta < 0.5$), and 55% of the way there for ($\beta \geq 0.5$).¹⁷ Both distributions of $\hat{\beta}$ have little dispersion, given β . Since these correlations have limited impact on predicted behavior, we omit them for parsimony in our primary specification, where $\hat{\beta}$ is on average 66% of the distance from β to 1, for all β .

Across our specifications, we estimate the perceived benefits and costs of going to the doctor based on consumers' revealed preferences. In our primary specification, which estimates both observed and unobserved heterogeneity in costs and benefits, the mean perceived benefit is 61.3 rupees while the mean perceived cost is 26.5 rupees.¹⁸ Consumers who have high blood pressure

¹⁷Note that β and $\hat{\beta}$ are mechanically correlated by the fact that $\hat{\beta}$ must be between β and 1. The correlation being estimated here is whether the proportional distance from β to 1 for $\hat{\beta}$ varies as a function of β .

¹⁸In the two baseline specifications mean estimated benefits are 30 while mean estimated costs are slightly above 30 in each (32.1 and 30.5 respectively).

and who are thus at high risk for hypertension-related medical issues have higher preferences for attending the doctor, equal to 40.34 rupees on average. Consumers high on the sickness index (which is not specifically related to hypertension) value going to the physician by 13 rupees less than other consumers for hypertension treatment. Consumers with prior hypertension medication value going to the doctor by 13 rupees less as well.¹⁹ Males value going to the doctor by 45.37 rupees more than females, while literate consumers value going to the doctor by 36.56 rupees more than illiterate consumers. For consumer costs, being employed increases costs very slightly, by 4 rupees on average, while our baseline specifications show that distance to the camp has a limited impact on preferences.²⁰ The unobserved component of perceived benefits, σ_ϵ has a standard deviation of 197.7 rupees, around the mean of 61, implying this is an important component of consumer demand in our setting.

It is important to note that, though the discount factor δ is technically identified from perceived benefit b , in practice with our sample these factors are not robustly estimated separately from one another across our specifications, due to the limited sample size.²¹ δ is estimated to have a value of 0.234 in our primary specification, but has estimated values of 0.780 and 0.687 in our two baseline specifications. Since choices at $t = 0$ are based on the perceived net benefit of attending the doctor $\delta * b - c$, we present the distribution of this net benefit in Table 5 since this is a quantity that is robustly identified. For our primary specification, the mean of this net benefit is -19.3, with 25th quantile equal to -55.50 and 75th quantile equal to 17.93. At the high end of the distribution, this net benefit is greater than 100. The estimated mean net benefit is similar to this in both baseline specifications, while the spread is similar but slightly larger. This suggests that, though it is difficult to separately identify δ and b in practice, we are robustly identifying the perceived net benefit of visiting the doctor $\delta * b - c$, which is the important sufficient statistic for identifying the joint distribution of β and $\hat{\beta}$.

Table 6 presents some statistics related to model fit. The model predicts take-up of 7.2% for the commitment contracts without bundled discounts, compared to 13.8% in the data. It predicts take-up of 26.7% for commitment contracts with bundled discounts, compared to 32.1% in our data. It predicts 4%, 8%, and 8% doctor attendance for people in the control, fixed contract with no discount, and personalized contract with no discount treatments respectively, compared to actual values of 8%, 9%, and 9% in the data for these treatment groups. These equivalent predicted values for the treatments with bundled discounts are 8.4%, 18%, and 14.3% while these moments in

¹⁹While at first these two findings may seem counter-intuitive, for the first finding it is likely that sick consumers are less concerned about preventive services for conditions they don't have. For the second, it is possible than once consumers have medication for hypertension they are then less likely to go for preventive visits (though they are still strongly recommended.)

²⁰15% of our observations have missing distance values, which we measure as either being less than 1 km from the doctor, between 1 and 5 km for the doctor, or greater than 5 km. As a result, we include this in our baseline specification, and once we verify that it has a very limited impact there (1.36 rupees per kilometer additional cost per visit) we do not include this in our primary specification so we can include these 15% of observations.

²¹See Section 4 above for a high-level discussion of identification.

the data are 14.4%, 12.5%, and 13.3%. For the personalized commitment treatments, the average additional commitment amount predicted by the model is 1 rupee (close to the minimum) while in the data this average is 4.55 rupees. The model fit overall is thus quite strong in replicating the levels of these three key sets of moments. It does not fit perfectly for the ordering of these moments across treatments, in part because the model imposes a structure on this ordering that may not always occur in practice due to sample size (or alternative factors not considered in the model, see the above discussion section).

Welfare and Distributional Implications. We now turn to the welfare implications of the different commitment contract and discount offerings in our setting. We focus on both the mean welfare implications and the distributional implications, since commitment contracts generate both winners, who use the commitment to help them go to the doctor, and losers who make a commitment but don't follow through on it.

We compute consumer welfare from a long run, or $t = 0$ perspective. This means that consumer welfare in each treatment equals their discounted benefit of going to the doctor, minus the costs of going (both pecuniary and non-pecuniary), and minus any pledged commitments amount lost due to not following through:

$$CW_{i,t} = 3 * (\delta b_i - c_i - f_i + d_i \mathbf{1}[D_{i,t} = 1] - m_{i,t} \mathbf{1}[C_{i,t} = 1, A_{i,t} = 0]) \quad (8)$$

Here, $\mathbf{1}[D_{i,t} = 1]$ is an indicator taking on value one if a consumer receives a per-visit discount in treatment t and $\mathbf{1}[C_{i,t} = 1, A_{i,t} = 0]$ is an indicator taking on value one if a consumer accepts a commitment contract but loses their pledged amount due to not following through on that commitment.²² We multiply per-visit welfare by 3 since all decisions (including commitment) are made with respect to a bundle of three visits in the model, a useful rescaling to match the empirical setting. The difference in welfare between two treatments T and T' is:

$$\Delta CW(T, T') = \sum_{i=1}^N CW_{i,T} - \sum_{i=1}^N CW_{i,T'}$$

It is also useful to consider consumer welfare net of the discount. If a planner providing discounts cares about the health behavior, but considers the discount to be a net transfer, e.g. coming from some other transfer program, then this is the relevant measure of consumer welfare. We define this consumer welfare net of discounts as:

$$\begin{aligned} CW_{i,T}^{ND} &= 3 * (\delta b_i - c_i - f_i - m_{i,T} \mathbf{1}[C_{i,T} = 1, A_{i,T} = 0]) \\ \Delta CW^{ND}(T, T') &= \sum_{i=1}^N CW_{i,T}^{ND} - \sum_{i=1}^N CW_{i,T'}^{ND} \end{aligned}$$

²²For the treatments where a discount is bundled with take-up of a commitment contract, $D_{i,t} = 1$ only for consumers who accept a contract.

Finally, we also consider a social welfare criterion that incorporates firm costs and profits. If the planner counts firm welfare the same as consumer welfare, then a consumer paying money to the firm from a commitment contract, but not following through, is just an even transfer, rather than a cost. Thus, social welfare gives a sense of how much a treatment improves consumer medical benefits, net of firm costs and consumer attendance costs. A key aspect of social welfare is what a visit to the doctor costs the provider. In our setting, it is possible that providers have marginal costs from patients actually visiting the doctor, as well as fixed costs of setting up the clinic and being present, whether patients show up or not. Our primary approach assumes that costs are fully marginal, denoted c^* and equal to the cost of a visit in the baseline environment of 30 rupees. This assumption is a zero profit condition on the margin: in practice true firm costs could be higher or lower than this amount depending on our partner’s operations. Discounts provided in the treatments are considered to be transfers, and thus do not impact social welfare. Given this, we define social welfare as follows:

$$\begin{aligned} SW_{i,T} &= 3 * (\delta b_i - c_i - c^*) \\ \Delta SW(T, T') &= \sum_{i=1}^N SW_{i,T} - \sum_{i=1}^N SW_{i,T'} \end{aligned}$$

This social welfare criterion will yield strictly higher welfare than consumer welfare netting out discounts, CW^{ND} , since some of the lost consumer surplus is transferred to the firm but the firm doesn’t incur the cost of the consumers’ care.

All welfare criteria take a long-run perspective in the sense that they do not consider present bias β to be welfare-relevant, and only consider welfare from the perspective of the $t = 0$ decision-maker. Additionally, for our primary discussion we take the perspective that the benefits b_i estimated via revealed preference reflect the true benefits of preventive treatment. Shortly, we discuss some alternative exercises where consumers under-value such services. We also make the standard utilitarian assumption that welfare can be captured by a simple summation of individual utilities, with equal weighting across people.

Table 6 presents these welfare results. Relative to the control treatment, the discount only treatment increases CW by 5.47 rupees, averaged across all consumers, including those with no welfare change. There are no losers relative to control, by construction, and consumer welfare increases because more consumers go to the doctor and those already going pay less. The first panel in Figure 11 plots the entire distribution of consumer welfare impacts for the discount-only treatment relative to control, except for the zeros, which have considerable mass (and thus are excluded to make the rest of the figure easier to interpret). Even once discounts are netted out, the discount-only treatment increases consumer welfare (CW^{ND}) by 1.51 rupees per consumer, since attendance and thus health benefits increase.

In contrast, the undiscounted fixed commitment contract treatment causes consumer welfare

losses (under both criteria) of 5.08 rupees per consumer. 3% more consumers attend the doctor under the fixed contract, which leads to a benefit, but over half (59%) of the consumers taking up this contract do not follow through on their pledge and lose money as a result. On net, the availability of this fixed commitment contract is worse for consumers than no intervention at all. The second panel of Figure 11 plots the entire non-zero distribution of these welfare impacts. The undiscounted personalized commitment contract treatment similarly reduces consumer welfare by an average of 4.12 rupees. It is slightly better than the fixed contract, since consumers choose lower commitment amounts in this treatment, and lose less money when they fail to follow through at similar rates.

When discounts are bundled with commitment, the flexible commitment contract becomes worse than the fixed contract. For example, using CW^{ND} as a criterion, the personalized contract with discount is 2.89 rupees worse on average than control, compared to 2.42 for fixed without discount. This occurs because, with personalized contracts, bundled discounts bring many naive and partially naive consumers into the commitment contract, but they choose too low a commitment amount to ensure follow-through. 12.51% of consumers lose money under personalized contracts with discount, the highest percentage of any treatment, while fewer consumers go to the doctor (14.25%) than under the fixed contract with discount (18.00%).

The fixed commitment contract with discount is instead calculated to make consumers slightly better off than the control (or the discount-only treatment) under baseline consumer welfare CW , due to the stronger commitment. The third panel of Figure 11 plots the entire non-zero distribution of the baseline consumer welfare impacts for this treatment comparison. 8.71 pp of consumers lose money relative to the discount only treatment, but 9.6 pp more consumers are predicted to go to the doctor. We urge caution in interpreting this calculation, since the model over-predicts follow-through relative to the data in this particular treatment. The last two panels of Figure 11 illustrates the non-zero consumer welfare impacts (CW) of moving from personalized to fixed commitment contracts, for the cases with and without bundled discounts.

These results have a number of implications. First, the option to enter into a commitment contract can reduce consumer welfare on average, and generate more losers than winners. This occurs in our environment where consumers exhibit (i) a high degree of present-bias and (ii) are relatively unsophisticated about that present-bias.

Second, whether fixed or personalized commitment contracts are a better option depends on the tradeoff between the money partially naive consumers might lose if they don't follow through versus the increased probability they will follow through with a higher commitment amount. With bundled discounts, fixed contracts are better than personalized ones in our setting since the impact of the fixed contract on follow through outweighs the incremental losses from naive consumers who don't follow through. Without bundled discounts, the reverse is true, and the losses from those not following through outweighs the benefit from increasing the probability of follow-through.

Third, the discounts to consumers in our environment have a meaningful and positive impact under both consumer welfare metrics, relative to corresponding treatments without those discounts. This need not necessarily be the case if the discount selects in consumers with high naivete about present-bias relative to those who leverage these discounts to attend the doctor.²³ Fourth, though the treatments with discounts are positive for consumer welfare on average, they do generate more losers than the treatments without discounts: with fixed contracts the bundled discounts generate 4.5% more losers relative to fixed contracts with no discounts, with an analogous increase in losers of 7.6% for personalized contracts.

Finally, though the best treatment for average baseline consumer welfare CW in our setting is the fixed contract with bundled discounts treatment, once discounts are netted out in CW^{ND} the discount-only treatment is the best. Once discounts are netted out of consumer welfare, all commitment contract treatments are welfare negative, due to limited consumer follow-through on these commitments.

The preceding discussion relates to the consumer welfare impacts of different treatments. Table 6 also presents the social welfare results. Social welfare is positive for all treatments relative to the baseline control, and especially for the commitment contracts with bundled discounts. Overall, more consumers go to the doctor under those treatments, and the benefits to the consumer of going outweigh the social costs (consumer costs of attendance and firm costs). For example, for the fixed CC treatment with discount, the change in social welfare relative to control is 27.13, while the consumer welfare change, net of discounts, is -2.42. The analogous statistics for the personalized CC treatment with discount are -2.89 and 39.31. Thus, in these treatments, while social welfare increases, on average all of the surplus is given to the firm while consumers are on average, net losers (though some gain and some lose). The table shows similar, but smaller impacts for the CC treatments not bundled with discounts. These results illustrate both that (i) firms can benefit from offering commitment contracts, at the consumers' expense and (ii) that it is important for such policies whether a planner wants to consider consumer or social welfare as their benchmark.

One important assumption behind the welfare analysis in this section is that the perceived benefits from doctor visits that we estimate correspond to the actual welfare-relevant benefits. Recent research suggests consumers do not correctly perceive the benefits of preventive medical treatments – something that may be especially true in our setting where consumers have limited experience with formal scientific bio-medicine [see Baicker et al. (2015)]. Though we do not integrate exact estimates on the value of the hypertension treatments we study for welfare, which is a difficult challenge not well addressed in the medical literature, Table 6 does present consumer welfare results, with CW^{ND} , for each treatment, assuming the actual benefit consumers get from going to the doctor is larger than their perceived benefits. We present three cases where consumer actual benefits equal 500, 5,000 or 50,000 rupees (approximately 10, 100, and 1,000 USD per visit,

²³In addition, discounts could bring in consumers who value the treatment at less than f , which would lead to a consumer welfare loss in our setting when netting out discounts.

respectively).

As the benefits from treatment get larger, the number of losers (i.e. those who purchase contracts but do not go to the doctor) and their welfare losses remain the same, but the benefits to winners are now larger. This means that now, even if commitment contracts only encourage a small percentage of consumers to follow through, they might be worth implementing, since those consumers might benefit by a much larger amount than the losers lose. This can be seen in the bottom rows of Table 6 where the discount commitment contracts treatments start to deliver massive benefits relative to the treatments without discounts or without commitment contracts, since they encourage a meaningful portion of additional consumers to go to the doctor. These results highlight that if consumer health benefits are larger than what they perceive them to be, commitment treatments are more attractive even if many consumers are partially naive and don't follow through on their commitments. Note that, in practice, we find that follow-through on commitment (conditional on sign-up) is only weakly correlated with baseline systolic BP (p-value = 0.128), and uncorrelated with baseline diastolic BP (p-value = 0.451) and BMI (p-value = 0.594). That is, individuals with worse health status, who might have higher benefits from attending, are only weakly more likely to follow through on their commitments.

Of course, these conclusions rely on a welfare function that weights all people equally regardless of their baseline health status, or level of utility. A more sophisticated social planner might place more weight on preventing losses than achieving gains, or might place different relative value on the health gains. For all of these reasons, these welfare calculations remain somewhat illustrative.

Counterfactual Contracts. In addition to studying the welfare impacts of our observed treatments, we use our structural estimates to study the welfare impacts of unobserved contract configurations. While out-of-sample prediction leans more heavily on the modeling assumptions, it allows us to illustrate some of the broader patterns of how different discount-commitment bundles impact consumer welfare, with some potentially important implications for contract design in other settings.

We study a grid of fixed commitment contracts for the discount d and fixed per-visit commitment amount m . We focus on fixed contracts since personalized contracts will always have consumers choosing the same amount m^* if they take up the contract (which in our environment is typically close to 0). In practice, the fixed contracts with $m = 0$ mimic the personalized contracts in our setting closely. We study five values of d (0, 7.5, 15, 22.5 and 30) and four values of m (0, 15, 30 and 45) such that we have 20 commitment contract in our simulations. We also present the discount only results (no bundled commitment contract) for comparison.

Table 7 presents four statistics for the case where the entire population we study is offered each potential contract. It presents (i) the % of consumers attending the doctor under each offered contract, (ii) the mean consumer welfare impact of each contract netting out discounts (CW^{ND}),

(iii) the mean baseline consumer welfare impact (CW) of each contract relative to our control treatment, and (iv) the % losers (from a CW perspective) relative to the control case. The simulations assume that only one contract is offered to consumers in each case, and thus does not consider things like menu design with multiple commitment contract options.

A number of results emerge. First, from a positive standpoint, doctor attendance is always increasing with the discount d conditional on a given commitment amount m . For example, for $m = 15$, the percentage of consumers going to the doctor increases from 7.36 to 21.61 as the discount increases from 0 to 30 (which is basically giving the service away for free, except for the commitment amount). Second, for fixed d , doctor attendance is always increasing in m in our simulations. In our model, weakly more consumers will take up contracts with higher commitment amounts, while those larger commitment amounts will increase participation at $t = 1$. If consumers are apprehensive about higher commitment amounts because of uncertainty about follow-through ability, something that we do not consider, or due to liquidity constraints, this result might not hold. Overall, the % of losers is decreasing with m as m increases, though the amount lost per loser increases.

Figure 12 displays the consumer welfare results, netting out discounts (CW^{ND}), from Table 7. A first key insight is that welfare is not monotonic in the commitment level m , given a discount level d . For example, when $d = 0$, availability of a commitment contract where $m = 15$ decreases welfare by an average of 5.08 rupees per person, but when m is raised to 45 the availability of this contract now increases welfare by an average of 8.38 rupees per person. The top panel of the figure reveals that this is a consistent pattern across the different d levels we consider: relative to the corresponding discount only treatments, commitment contracts with level $m = 0$ and $m = 15$ are welfare decreasing, but those with $m = 30$ and $m = 45$ are welfare increasing. This reveals that if partial naivete about present-bias is the main driver for lack of follow through to commitment, either (i) removing the option to commit or (ii) making the commitment stronger would be welfare improving in our setting, where $m = 15$ is the fixed commitment contract commitment level. When the ability to commit is removed, consumers don't lose money through commitments they don't follow through on. Conversely, when the commitment amount is strengthened, even most partially naive consumers will follow through anyway because the commitment is so meaningful.

Our results also have interesting implications for discounts, and how they are bundled with different commitment levels. The bottom panel of Figure 12 plots the welfare impacts of different discount amounts as a function of the possible commitment amounts m . Whether discounts are complementary to commitment amounts m depends on the level of m . For high levels of m discounts are welfare decreasing. Given that almost all consumers signing up for contracts follow through at high m , welfare is decreasing in d because such discounts induce consumers with benefits lower than the social costs to utilize preventive care. This can also be seen in the fact that, for the discount-only treatments, welfare is increasing in d for $d \leq 15$ but decreasing for $d > 15$. For low

commitment levels m the impacts of discounts are non-monotonic: at low d welfare is increasing in d , given low m , but at high d welfare is decreasing in d , given low m

Taken together, our results suggest that commitment contracts and bundled discounts induce subtle trade-offs whose welfare implications depend on the specifics of the empirical context. Changes to up front fixed commitment amounts induce a trade-off between the amount that losers lose, the number of losers, and the number of consumers going to the doctor. In our setting, the amount lost by losers generally outweighs the gains from increased attendance by winners at low m , but the reverse is true at high m . Introducing bundled discounts changes the nature of this trade-off, with higher discounts being better at low m , but not so at higher m . Importantly, bundled discounts may also lead to a traditional source of inefficiency whereby consumers who value a product at less than its social cost buy it because of the discount. Finally, if welfare were concave, a planner would want to consider the number of losers (which is increasing in d , given m) and the amount of those losses (increasing in m).

These results are most relevant to settings where consumers have limited uncertainty about the product they are purchasing and their stochastic costs of following through on commitment. In our setting, if m gets very large, this must be welfare increasing, since consumers follow through all the time. With budget constraints or stochastic costs that can be large, this need not be true. The cases we study of $m = 30$ and $m = 45$ use our model to extrapolate out of our environment, and in our environment $m = 45$ is enough to get essentially all consumers signing up for commitment contracts to follow through. That said, the trade-offs explored in this analysis will be present in essentially any environment where consumers have the ability to choose commitment contracts, though the impacts of these tradeoffs will vary as some of these assumptions are relaxed.

6 Conclusion

We conducted a field experiment in rural India to evaluate whether commitment contracts and price discounts increase preventive health visits by hypertensive patients. The results are generally disappointing from the point of view of harnessing commitment contracts in the health sector: take-up of the contracts was modest, few of those who purchased the contracts ended up utilizing health services, and objective health outcomes (blood pressure, weight) do not change in the treatment groups. Under plausible model assumptions, offering individuals commitment contracts reduces social welfare in the context we study.

A methodological contribution of the project is to design a natural field experiment with a tight link to a theoretical model, generating both clear reduced-form facts, and structural estimates of the key utility and belief parameters. This approach remains relatively rare in the development economics literature, although it builds on a recent stream of work in structural behavioral economics (DellaVigna et al. 2012; Augenblick et al. 2015). An advantage of this approach is that the identification of structural parameters here relies on exogenous variation due to randomization.

Recovering the structural parameters allows us to conduct welfare and counterfactual analysis, which ultimately may be tested through additional experimentation.

A substantive contribution is to provide one of the first estimates of the distributions of sophistication about present bias. The theoretical literature has highlighted the importance of this parameter (Laibson 1997; O’Donoghue and Rabin 1999). Only sufficiently-naive agents are predicted to engage in procrastination; conversely, only sufficiently-sophisticated agents will demand commitment and use it to achieve the first-best. Understanding the distribution of sophistication is thus crucial to understanding the nature of self-control challenges, and to identify the appropriate policy responses. Yet, to date, we have very few empirical estimates, with Augenblick and Rabin (2015) being a notable exception. We provide evidence of partial naivete on average, with substantial variation. We complement the existing literature by providing evidence from a field setting, and by relying not on incentivized predictions of future behavior, but instead by examining the amount of costly commitment that agents choose when designing their own commitment contracts.

Our findings and the underlying theory both suggest the need for caution in the design of commitment contracts. If partial naivete is common, as our results indicate, then agents are likely to demand costly commitment, but then systematically accept commitments which are not strong enough to succeed. One policy response would be to restrict the set of available commitments to *strong* commitments. But this poses a tradeoff with consumers’ demand for flexibility in an uncertain world (Laibson 2015; Amador et al. 2006). Another implication of widespread partial naivete is that the many agents who do *not* demand a particular commitment device might simply be over-optimistic about their self-control problems, as opposed to not having a self control problem to begin with. That is, even offering strong commitments may leave unhelped a number of consumers with self-control problems.

A potential solution to these problems emerging from our approach would involve the design of optimal contracts, utilizing the estimated distributions of beta and beta-hat to design a menu of commitment options. Indeed, with enough individual choice data, one could in principle design optimal contracts for each individual, as in Andreoni et al. (2016). While this remains an intellectually valuable exercise, it faces obvious shortcomings as a scalable public health tool to target consumers. A practical solution could instead be to allow consumers to experience different levels of incentives before they are offered commitment, providing them an opportunity to learn how strong an incentive they require to overcome their self-control problem in a given context.

Despite these open questions, it seems clear that commitment contracts are not a panacea for improving preventive health behaviors. What, then, is the best approach for boosting preventive care in low income countries? This is a critical public policy issue in many developing countries, in Asia, Africa and Latin America, as chronic lifestyle diseases, metabolic disorders and other chronic conditions related to diet and obesity plague increasingly prosperous - and aging - populations.

One, perhaps unsurprising, lesson that does emerge from our data is that simple subsidies are

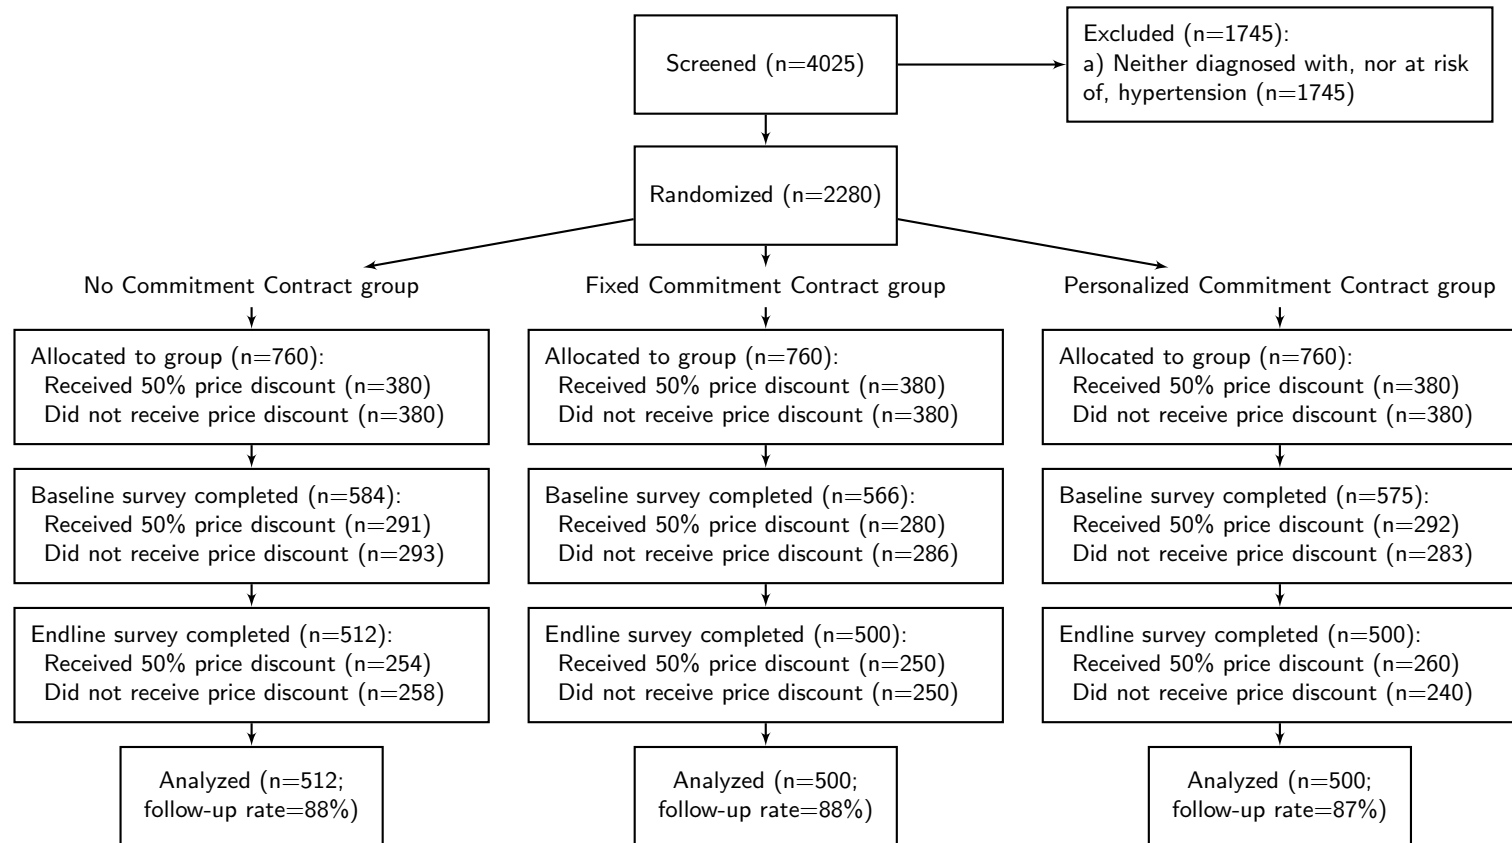
effective tools for driving higher demand for these services. Yet, even with a forceful intervention targeting an at-risk population, deploying subsidies, commitment, information provision, a high-quality provider and home-visit reminders, overall utilization remains very low in our setting. Reducing hassle costs and better understanding the dimensions of provider quality consumers care about is likely to be a more promising approach, as is providing direct incentives to utilize preventive care.

References

- Acland, Dan and Matthew R. Levy**, “Naivete, Projection Bias, and Habit Formation in Gym Attendance,” *Management Science*, 2015, *61* (1), 146–160.
- Amador, Manuel, Iván Werning, and George-Marios Angeletos**, “Commitment vs. flexibility,” *Econometrica*, 2006, *74* (2), 365–396.
- Andreoni, James, Karrar Callen, Michael, Jaffar, Muhammad Yasir Khan, and Charles Sprenger**, “Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan,” 2016.
- Ashraf, Nava, Dean Karlan, and Wesley Yin**, “Tying Odysseus to the Mast: evidence from a Commitment Savings Product in the Philippines,” *Quarterly Journal of Economics*, 2006, *121* (2), 635–672.
- Association of Physicians of India**, “Epidemiology of Hypertension,” *Journal of the Association of Physicians of India*, 2013, *61* (6).
- , “Management of Hypertension,” *Journal of the Association of Physicians of India*, 2013, *61* (6).
- Augenblick, Ned and Matthew Rabin**, “An Experiment on Time Preference and Misprediction in Unpleasant Tasks,” *mimeo*, 2015.
- , **Muriel Niederle, and Charles Sprenger**, “Working Over Time: Inconsistency in Real Effort Tasks,” *Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein**, “Behavioral hazard in health insurance,” *The Quarterly Journal of Economics*, 2015, *130* (4), 1623–1667.
- Banerjee, Abhijit, Angus Deaton, and Esther Duflo**, “Health, health care, and economic development: Wealth, health, and health services in rural Rajasthan,” *The American economic review*, 2004, *94* (2), 326.
- Bryan, Gharad, Dean Karlan, and Scott Nelson**, “Commitment Devices,” *Annual Review of Economics*, 2010, *2*, 671–698.
- DellaVigna, Stefano and Ulrike Malmendier**, “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics*, 2004, *119* (2), 353–402.
- and – , “Paying not to go to the gym,” *The American Economic Review*, 2006, *96* (3), 694–719.
- , **John A List, and Ulrike Malmendier**, “Testing for altruism and social pressure in charitable giving,” *The Quarterly Journal of Economics*, 2012, p. qjr050.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson**, “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya,” *American Economic Review*, 2011, *101* (6).
- Dupas, Pascaline**, “Health behavior in developing countries,” *Annu. Rev. Econ.*, 2011, *3* (1), 425–449.

- **and Edward Miguel**, “Impacts and determinants of health levels in low-income countries,” *Handbook of Economic Field Experiments*, 2017, 2, 3–93.
- Gine, Xavier, Dean Karlan, and Jonathan Zinman**, “Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation,” *American Economic Journal: Applied Economics*, 2010, 2, 213–235.
- Handel, Benjamin R.**, “Adverse selection and inertia in health insurance markets: When nudging hurts,” *The American Economic Review*, 2013, 103 (7), 2643–2682.
- Heidhues, Paul and Boton Koszegi**, “Exploiting naivete about self-control in the credit market,” *The American Economic Review*, 2010, 100 (5), 2279–2303.
- John, Anett**, “When Commitment Fails – Evidence from a Regular Saver Product in the Philippines,” *mimeo*, 2015.
- Karlan, Dean and Leigh L Linden**, “Loose knots: strong versus weak commitments to save for education in Uganda,” 2017.
- Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan**, “Self-Control at Work,” *Journal of Political Economy*, 2015, *Forthcomin*.
- Kenkel, Donald S.**, “The demand for preventive medical care,” *Applied Economics*, 1994, 26 (4), 313–325.
- Kremer, Michael and Rachel Glennerster**, “Improving Health in Developing Countries,” *Handbook of Health Economics*, 2011, 2, 201–315.
- Laibson, David**, “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 1997, 112 (2), 443–478.
- , “Why don’t present-biased agents make commitments?,” *The American Economic Review*, 2015, 105 (5), 267–272.
- Mohan, Viswanathan, Yackoob Seedat, and Rajendra Pradeepa**, “The rising burden of diabetes and hypertension in southeast asian and african regions: need for effective strategies for prevention and control in primary health care settings,” *International journal of hypertension*, 2013, 2013.
- O’Donoghue, Ted and Matthew Rabin**, “Doing It Now or Later,” *American Economic Review*, 1999, 89 (1), 103–124.
- Royer, Heather, Mark Stehr, and Justin Sydnor**, “Incentives, commitments, and habit formation in exercise: evidence from a field experiment with workers at a fortune-500 company,” *American Economic Journal: Applied Economics*, 2015, 7 (3), 51–84.
- Schilbach, Frank**, “Alcohol and Self-Control : A Field Experiment in India,” *mimeo*, 2017.
- Train, Kenneth E.**, *Discrete choice methods with simulation*, Cambridge university press, 2009.
- World Health Organization**, “A global brief on hypertension: silent killer, global public health crisis: World Health Day 2013,” 2013.


Figure 1: Flowchart of participants' progress through phases of the trial



37

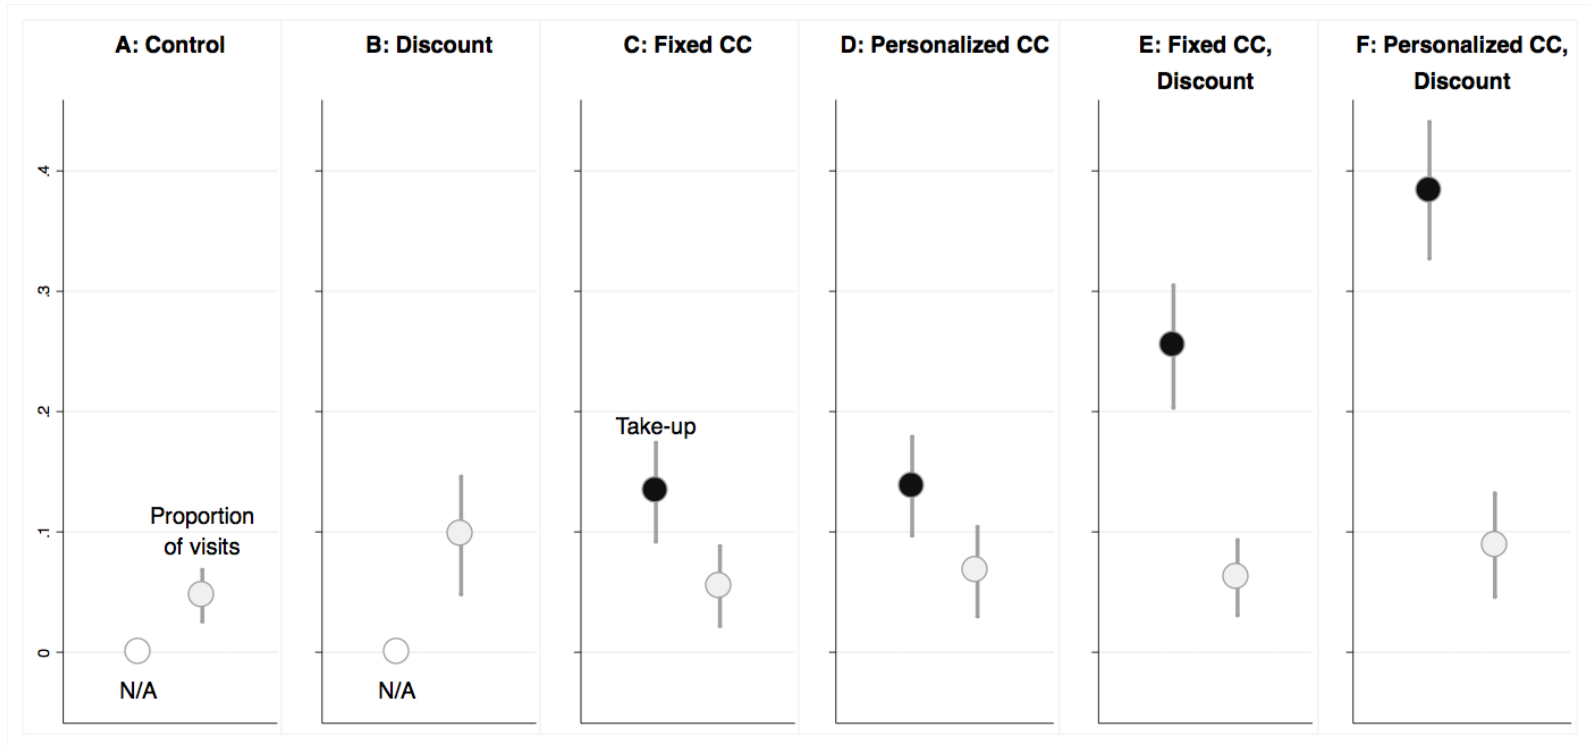
Notes: The initial round of screening was carried out during May-June 2012. The baseline survey and commitment contract offers were done during October 2012 - January 2013. The weekly hypertension health camps were held during August 2012 - June 2013. The endline surveys were completed during May-July 2013.

Figure 2: Experimental Interventions



Notes: Control group participants can visit the health camps on a pay-as-you-go basis (30 rupees per consultation), while commitment contract group participants were offered the opportunity to pay for multiple visits upfront. This is combined with either a fixed (15 rupees) or personalized (self-chosen) commitment amount, which they then receive back during future visits. Finally, the above three treatments are cross-cut with a simple discount treatment, where participants were charged 15 rupees per consultation.


Figure 3: Commitment Contract Take-up and Fraction of Clinic Visits by Treatment Group



39

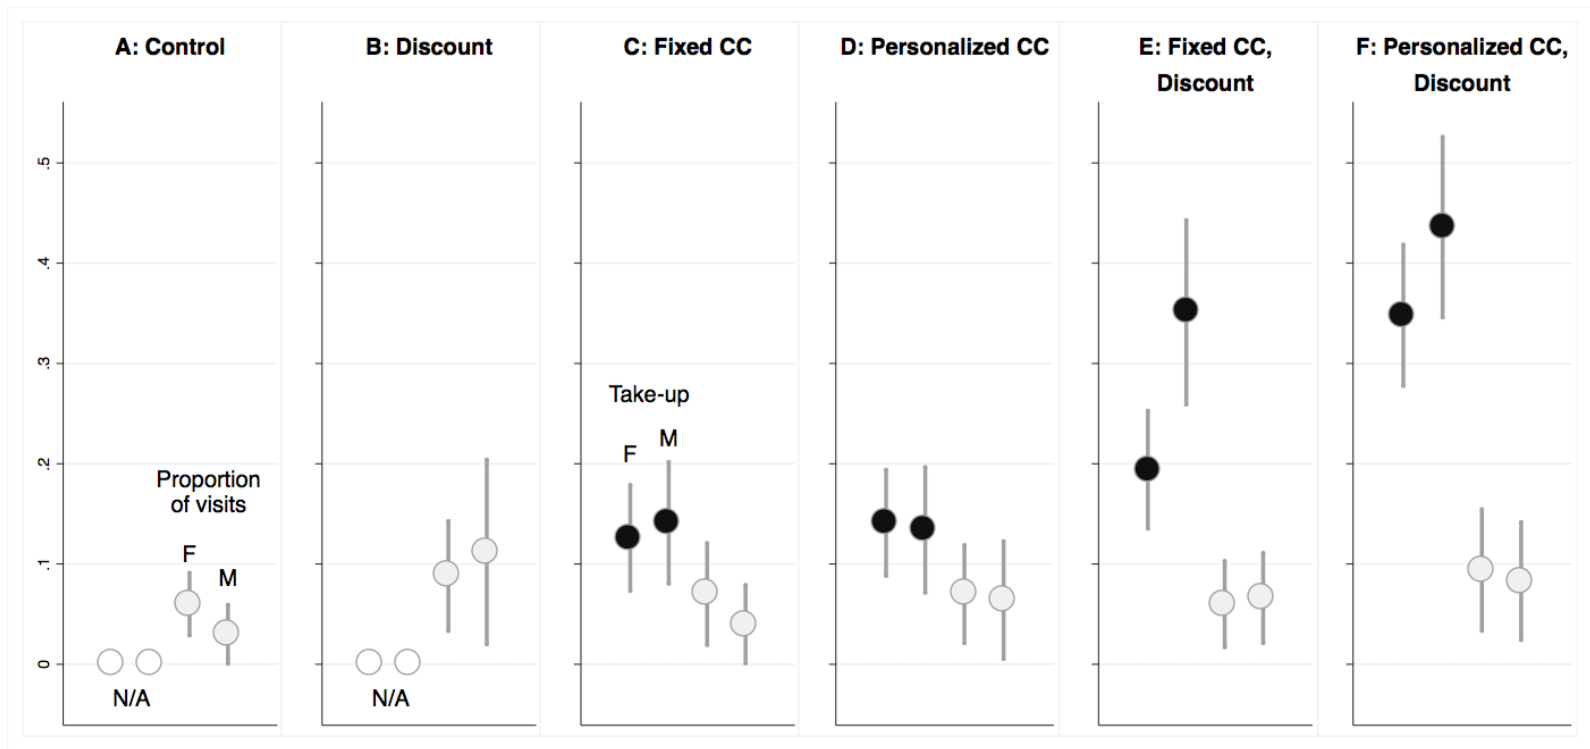
Notes: Take-up refers to the fraction of study participants in each treatment group that signed up for the commitment contract on offer. It is therefore not applicable to the control and discount groups (panels A and B respectively), as participants in these groups were not offered commitment contracts. Proportion of visits is the average for participants in each treatment group during the six-months intervention period, out of the recommended three visits. Black circles denote average take-up rates of commitment contracts in each treatment group, while gray circles denote average proportions of health camp visits in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure 4: Personalized Commitment Contract Contributions



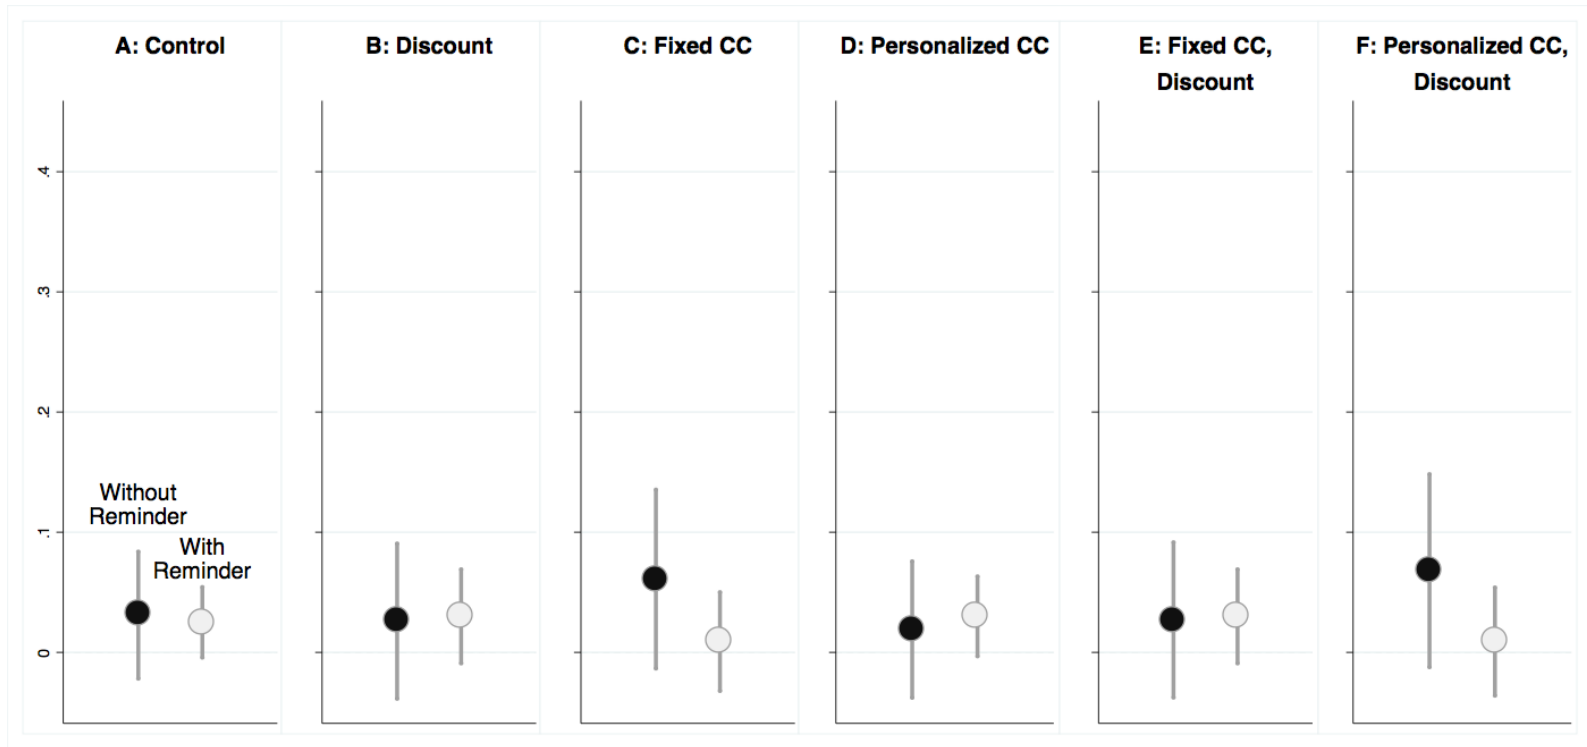
Notes: For participants who signed up for a personalized commitment contract, they could choose the amount m to be received at each future health camp visit. In the fixed commitment contract group, it was Rs. 15. In addition, the discount group also received 50% off the price of consultations (normally Rs. 30). The sample used here is comprised of the 114 study participants that signed up for a personalized commitment contract. Of these, 25 were in the personalized CC group, while 89 were in the personalized CC + discount group. There was only one case where a participant chose a commitment amount (Rs. 40 per visit) larger than that in the fixed CC.

Figure 5: Commitment Contract Take-up and Fraction of Clinic Visits by Treatment Group and Gender



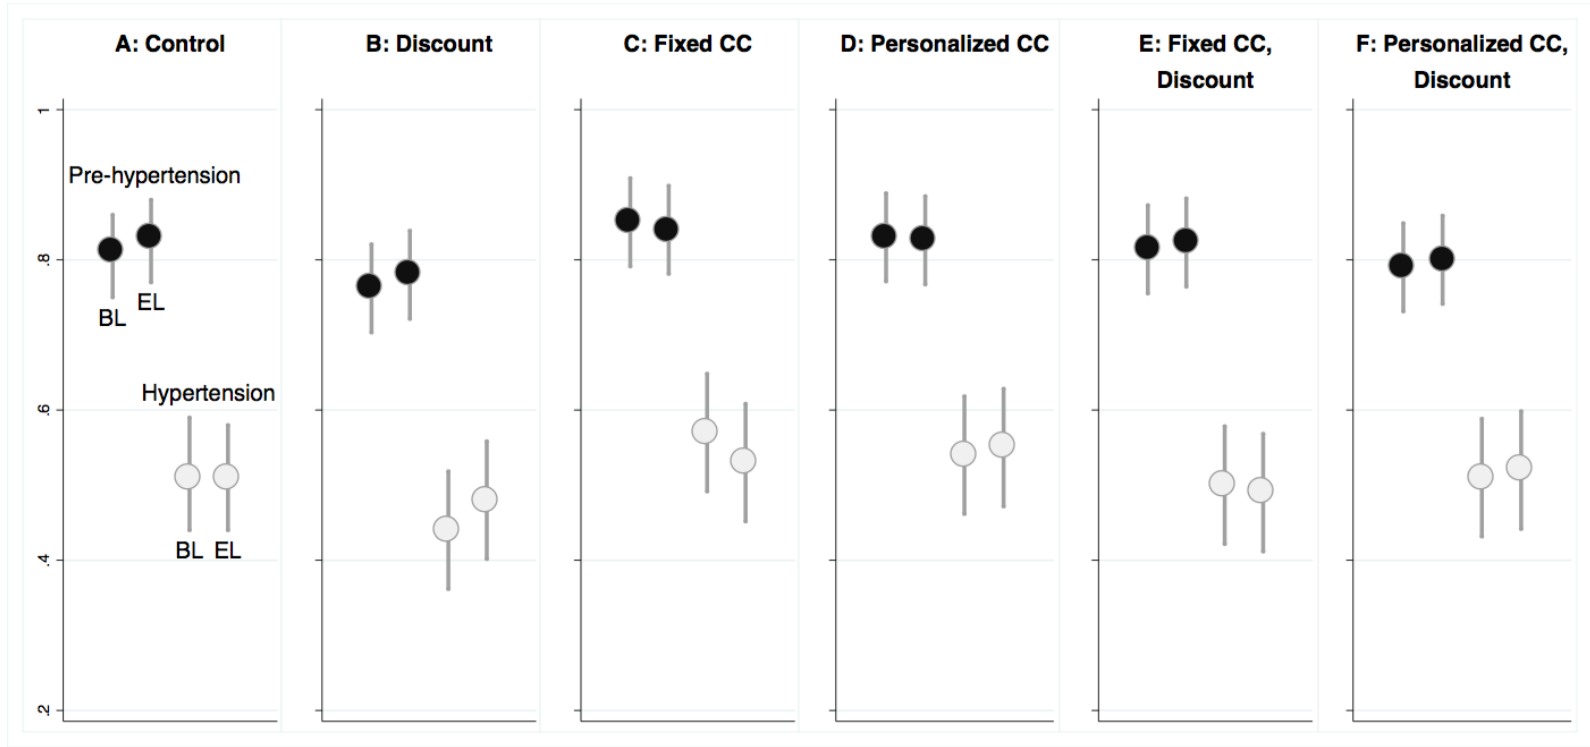
Notes: Take-up refers to the fraction of study participants in each treatment group that signed up for the commitment contract on offer. It is therefore not applicable to the control and discount groups (panels A and B respectively), as participants in these groups were not offered commitment contracts. Proportion of visits is the average for participants in each treatment group during the six-months intervention period, out of the recommended three visits. Black circles denote average take-up rates of commitment contracts in each treatment group, while gray circles denote average proportions of health camp visits in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure 6: Post-Reminder Clinic Visits by Treatment Group



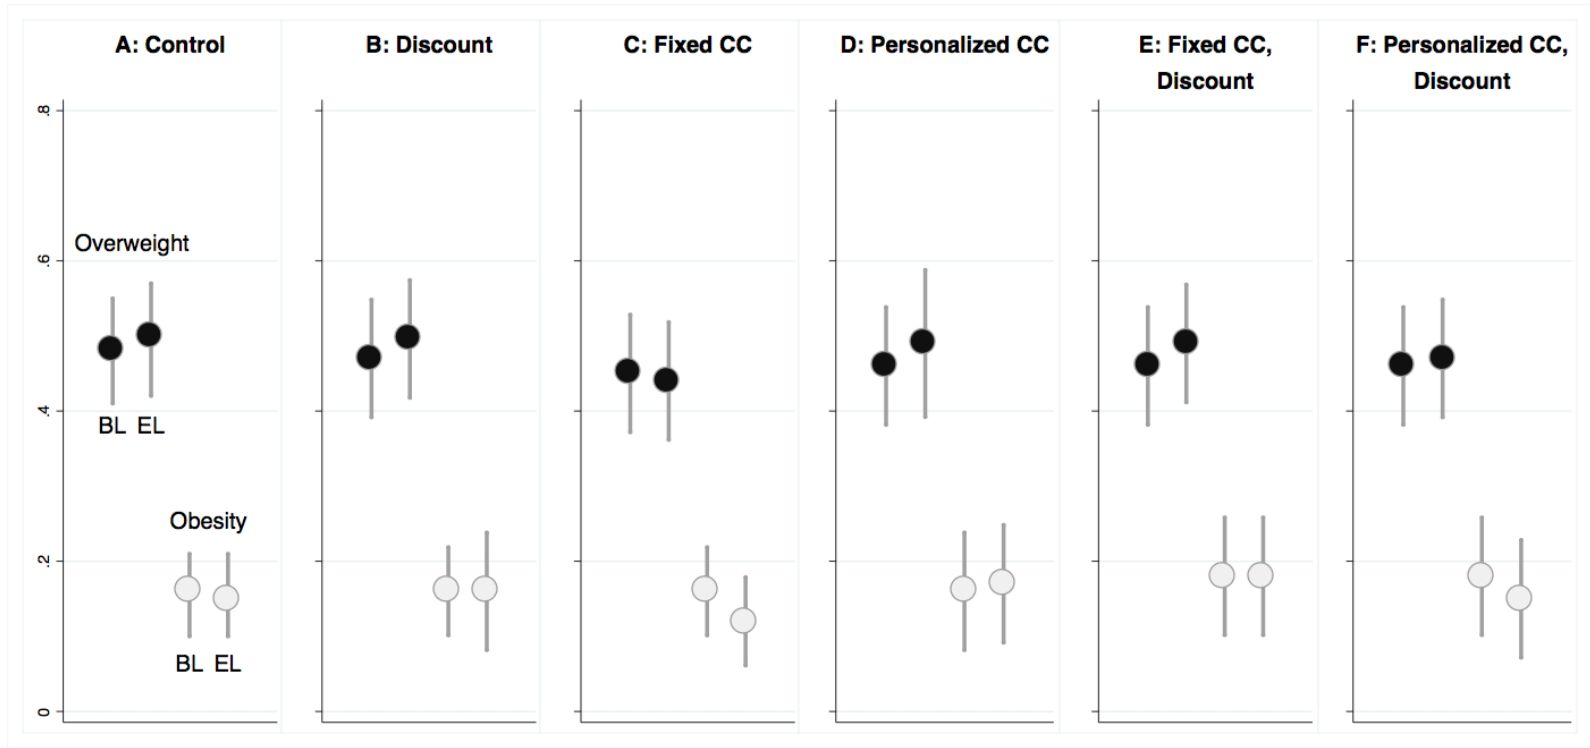
Notes: This figure shows the proportion of participants who visited the doctor after receiving randomized reminders. These were delivered in person, by enumerators at least two weeks before the discount coupons and commitment contracts expired. Hence, participants had at least two more opportunities to visit the doctor. Black circles denote average proportions of health camp visits in each treatment group for those that did not receive a reminder, while gray circles denote average proportions of post-reminder visits in each treatment group for those that did receive a reminder. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure 7: Hypertension Status by Treatment Group




Notes: A normal systolic (diastolic) blood pressure is below 120(80) mmHg. A systolic (diastolic) blood pressure number between 120(80) and 139(89) mmHg is considered to be pre-hypertension. A systolic (diastolic) blood pressure number of 140(90) mmHg or higher is considered to be hypertension. Black circles denote the average baseline (BL) and endline (EL) incidence of pre-hypertension and hypertension in each treatment group, while gray circles denote the average baseline (BL) and endline (EL) incidence of hypertension in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure 8: Weight Status by Treatment Group



Notes: Overweight is defined as a body mass index (BMI) of 25 or above. Obesity is defined as a BMI of 30 or above. Weight measurements were carried out by enumerators during baseline and endline surveys. Black circles denote the average baseline (BL) and endline (EL) incidence of overweight in each treatment group, while gray circles denote the average baseline (BL) and endline (EL) incidence of obesity in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

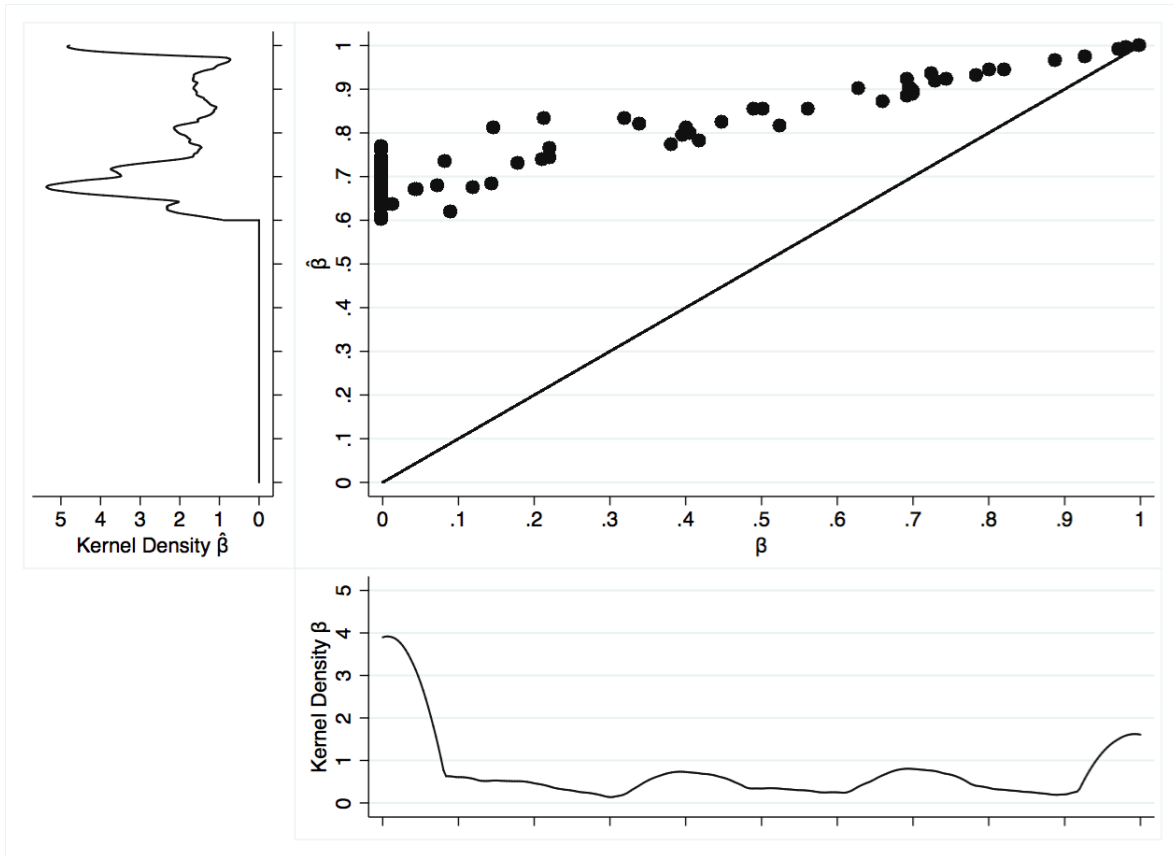
Figure 9: Identification of $\hat{\beta}$ and β in Theoretical Model



1. Predicts will not visit doctor even with CC, and does not visit even with CC.
2. Predicts will visit doctor only with CC, but does not visit even with CC.
3. Predicts will visit doctor only with CC, and visits only with CC.
4. Predicts will visit with and without CC, but does not visit even with CC.
5. Predicts will visit with and without CC, but visits only with CC.
6. Predicts will visit with and without CC, and visits with and without CC.

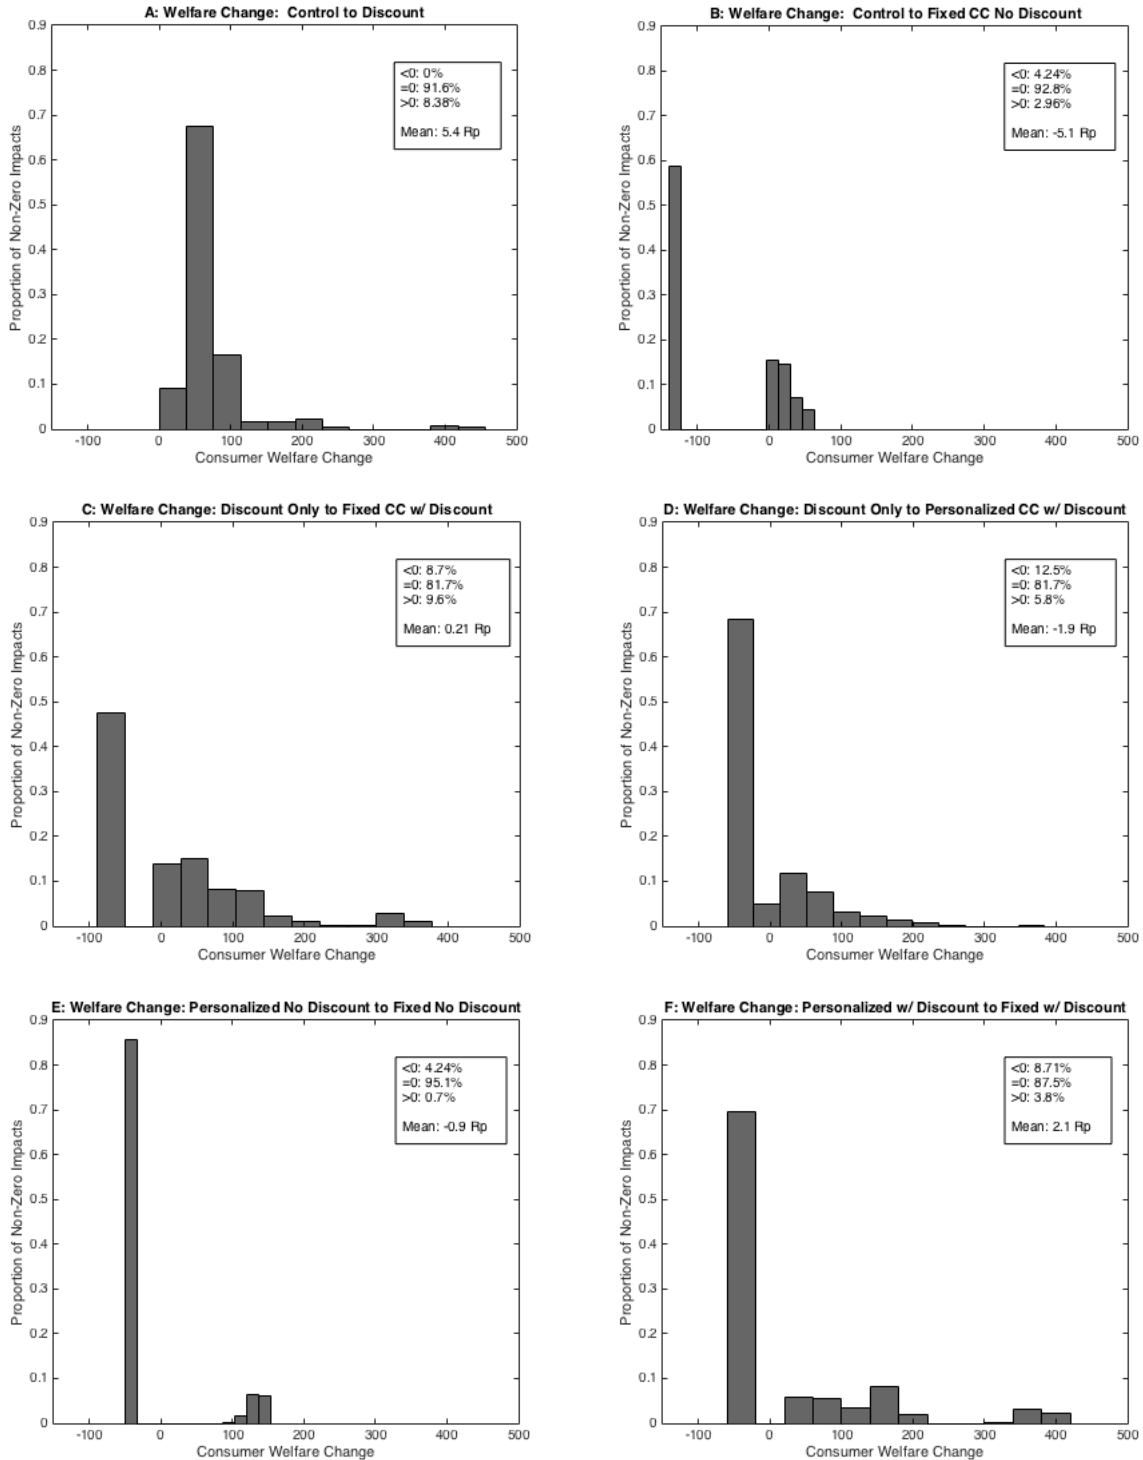
Notes: This figure illustrates how we can identify the parameters $\hat{\beta}$ and β from commitment contract choices and subsequent doctor visits. The white region correspond to individuals who predict they will not visit the doctor even with a commitment contract. The light gray region correspond to individuals who predict they will visit the doctor only with a commitment contract. The dark gray region correspond to individuals who predict they will visit the doctor both with and without a commitment contract. In the theoretical model of Section 4, b is the benefit from visiting the doctor, f is the standard per visit fee, c is the non-financial cost of going to the doctor, m is the commitment amount and δ is the exponential discount factor.

Figure 10: Joint distribution of β and $\hat{\beta}$ estimates from primary specification



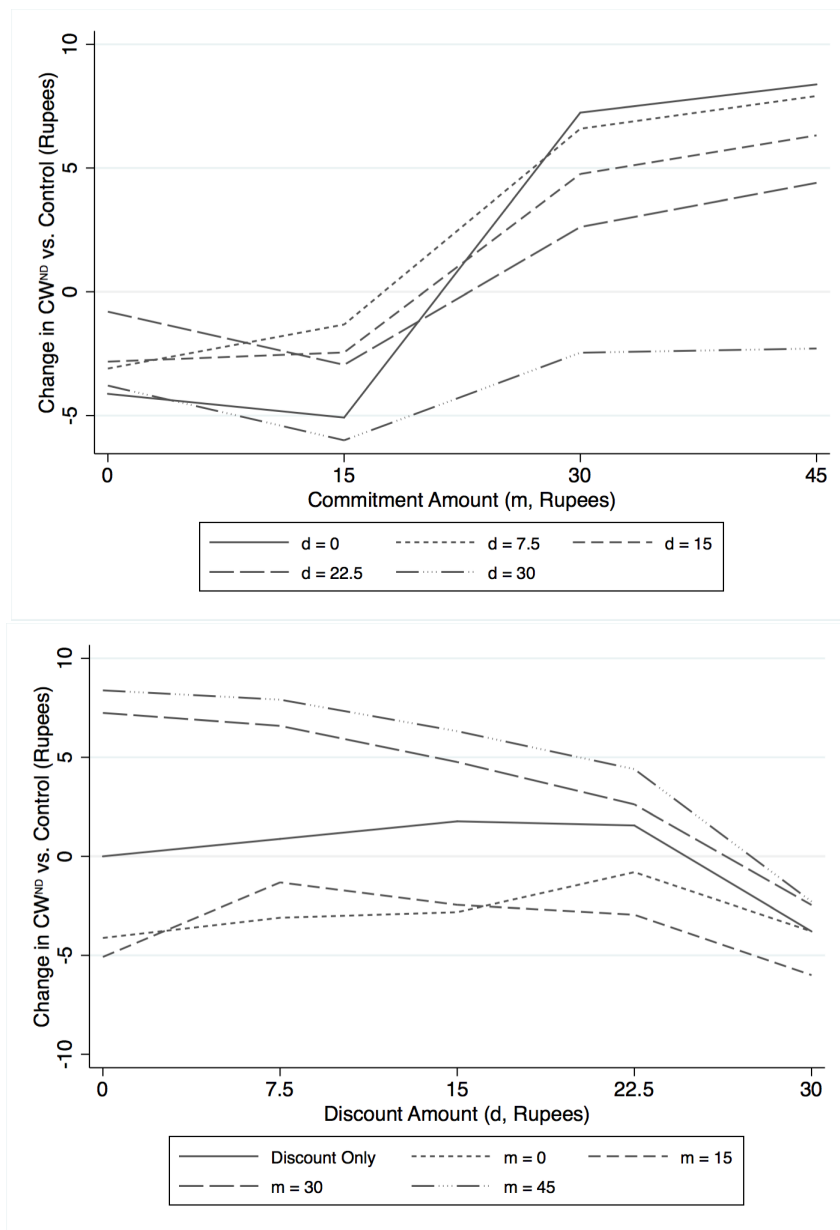
Notes: This figure plots the joint distribution of β and $\hat{\beta}$ estimates from our primary specification. Such a specification incorporates observable heterogeneity across a number of important dimensions. These include (i) our hypertension severity index (ii) our general health index (iii) gender (iv) whether the consumer was already taking hypertension medication prior to the study and (v) whether they are literate. The estimated mean of β is 0.365, with a population standard deviation of 0.395. The estimated mean of $\hat{\beta}$ is 0.795 with a population standard deviation of 0.13.

Figure 11: Distribution of Consumer Welfare Impacts from Different Treatments



Notes: This figure plots the distribution of non-zero consumer welfare impacts (CW) from moving across different treatment arms. These are computed assuming that our entire population is enrolled in a given treatment, and applying our parameter estimates (see Table 5) to study decision-making and subsequent outcomes. Furthermore, they are computed from an ex ante, or $t = 0$ perspective. This means that consumer welfare in each treatment equals their discounted benefit of visiting the doctor, minus the costs (both pecuniary and non-pecuniary), and minus any pledged commitment amounts lost due to not following through.

Figure 12: Consumer welfare (CW^{ND}) impacts of counterfactual commitment contracts as a function of the commitment level m and discount level d



Notes: This figure illustrates results from our counterfactual simulations that study alternative bundled discount-fixed commitment contract designs. Specifically, we study five different values of d (0, 7.5, 15, 22.5 and 30) and four values of m (0, 15, 30 and 45). The top panel plots the consumer welfare impacts, netting out discounts as transfers, of different commitment amounts m for each possible discount d , while the bottom panel plots these welfare impacts for different discount amounts corresponding to each possible commitment amount m .

Table 1: Summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Panel A: Demographic Characteristics					
Household Size	1728	5.51	2.26	1	21
Gender (Female = 0, Male = 1)	1728	0.41	0.49	0	1
Age	1728	53.7	14.3	30	105
HH Head Literate (Yes = 1, No = 0)	933	0.45	0.50	0	1
Household Income (1000 Rs)	1539	101.9	161.5	0	2000
HH Head Self-Employed Agriculture (Yes = 1, No = 0)	908	0.37	0.48	0	1
Panel B: Baseline Health Indicators					
Blood Pressure (Systolic)	1718	142.4	24.5	87	264
Blood Pressure (Diastolic)	1718	85.8	14.1	37	182
Weight (kg)	1677	64.9	14.7	28.3	164.9
Pre-hypertension (Yes = 1, No = 0)	1718	0.86	0.34	0	1
Hypertension (Yes = 1, No = 0)	1718	0.58	0.49	0	1
Overweight (Yes = 1, No = 0)	1677	0.51	0.50	0	1
Obesity (Yes = 1, No = 0)	1677	0.20	0.40	0	1
Panel C: Take-up and Service Utilization					
Contract Take-up (Yes = 1, No = 0)	1725	0.15	0.36	0	1
Proportion of Doctor Visits	1725	0.07	0.27	0	4
Any Doctor Visit (Yes = 1, No = 0)	1725	0.11	0.32	0	1
Panel D: Endline Health Indicators					
Blood Pressure (Systolic)	1512	138.8	24.2	64	245
Blood Pressure (Diastolic)	1512	85.0	15.4	12	161
Weight (kg)	1512	65.9	14.6	29.7	119.4
Pre-hypertension (Yes = 1, No = 0)	1512	0.82	0.38	0	1
Hypertension (Yes = 1, No = 0)	1512	0.50	0.50	0	1
Overweight (Yes = 1, No = 0)	1512	0.53	0.50	0	1
Obesity (Yes = 1, No = 0)	1512	0.20	0.40	0	1

Table 2: Main Study Outcomes by Treatment Group: Full Sample

	(1)	(2)	(3)
Panel A: Take-up and Service Utilization	Take-up	Any visit	Proportion of visits
Discount		0.056 (0.027)**	0.050 (0.025)**
Fixed CC	0.133 (0.021)***	0.006 (0.024)	0.008 (0.017)
Personalized CC	0.138 (0.021)***	0.010 (0.024)	0.020 (0.019)
Fixed CC + Discount	0.254 (0.026)***	0.036 (0.026)	0.015 (0.016)
Personalized CC + Discount	0.384 (0.029)***	0.048 (0.026)*	0.042 (0.022)*
Observations	1,725	1,725	1,725
R-squared	0.14	0.00	0.00
Control Group Mean (SD)		0.089 (0.285)	0.046 (0.192)
Panel B: Endline Health Outcomes	BP (systolic)	BP (diastolic)	Weight (kg)
Discount	0.614 (1.591)	0.081 (1.154)	-0.162 (0.487)
Fixed CC	-0.286 (1.558)	0.031 (1.153)	-0.442 (0.312)
Personalized CC	0.600 (1.506)	1.933 (1.180)	-0.128 (0.322)
Fixed CC + Discount	-0.461 (1.556)	-0.082 (1.076)	0.173 (0.324)
Personalized CC + Discount	1.105 (1.610)	1.747 (1.246)	-0.182 (0.292)
Observations	1,481	1,481	1,481
R-squared	0.48	0.29	0.91
Control Group Mean (SD)	139.7 (25.0)	85.0 (15.6)	67.4 (14.1)

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Village indicators are included in all specifications. Baseline blood pressure and weight measures are included as controls in Panel B regressions. This is done to improve the precision of our estimates. Regressions without baseline controls, as specified in our pre-analysis plan, are reported in Table A2.

Table 3: Main Study Outcomes by Treatment Group: “Ideal” Sample

	(1)	(2)	(3)
Panel A: Take-up and Service Utilization	Take-up	Any visit	Proportion of visits
Discount		0.106 (0.052)**	0.119 (0.076)
Fixed CC	0.172 (0.047)***	0.107 (0.055)*	0.077 (0.052)
Personalized CC	0.192 (0.046)***	0.031 (0.045)	0.025 (0.045)
Fixed CC + Discount	0.319 (0.055)***	0.074 (0.050)	0.040 (0.041)
Personalized CC + Discount	0.494 (0.057)***	0.130 (0.054)**	0.121 (0.067)*
Observations	439	439	439
R-squared	0.20	0.02	0.01
Control Group Mean (SD)		0.065 (0.248)	0.047 (0.246)
Panel B: Endline Health Outcomes	BP (systolic)	BP (diastolic)	Weight (kg)
Discount	0.276 (3.207)	-0.988 (1.905)	-0.318 (0.443)
Fixed CC	-0.802 (2.906)	-0.638 (2.008)	-0.910 (0.521)*
Personalized CC	3.557 (2.955)	3.526 (1.996)*	-0.228 (0.468)
Fixed CC + Discount	0.978 (2.979)	1.380 (1.830)	-0.972 (0.624)
Personalized CC + Discount	3.967 (3.034)	4.724 (2.128)**	-0.553 (0.646)
Observations	395	395	395
R-squared	0.49	0.36	0.95
Control Group Mean (SD)	143.4 (28.7)	85.6 (15.2)	69.4 (16.0)

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Village indicators are included in all specifications. Baseline blood pressure and weight measures are included as controls in Panel B regressions. The “ideal” sample is defined as respondents who both believe “it is possible to be healthy with hypertension if blood pressure is frequently monitored” *and* who trust the service provider.

Table 4: Endline Health Outcomes: Incidence of Hypertension and Obesity

	(1)	(2)	(3)	(4)
Panel A: Full Sample	Pre-hypertension	Hypertension	Overweight	Obesity
Discount	-0.027 (0.031)	-0.001 (0.036)	0.020 (0.034)	0.032 (0.030)
Fixed CC	0.022 (0.030)	0.018 (0.038)	-0.012 (0.035)	0.010 (0.030)
Personalized CC	0.008 (0.030)	0.051 (0.038)	0.030 (0.035)	0.055 (0.031)*
Fixed CC + Discount	0.012 (0.030)	0.006 (0.037)	0.004 (0.034)	0.041 (0.030)
Personalized CC + Discount	-0.008 (0.030)	0.045 (0.037)	0.026 (0.033)	0.038 (0.029)
Observations	1,481	1,481	1,481	1,481
R-squared	0.18	0.30	0.41	0.34
Control Group Mean (SD)	0.83(0.37)	0.49(0.50)	0.56(0.50)	0.20(0.40)
Panel B: "Ideal" Sample	Pre-hypertension	Hypertension	Overweight	Obesity
Discount	0.027 (0.055)	-0.046 (0.068)	-0.012 (0.062)	0.032 (0.061)
Fixed CC	0.042 (0.061)	0.047 (0.073)	-0.023 (0.067)	-0.039 (0.061)
Personalized CC	0.092 (0.055)*	0.182 (0.075)**	0.015 (0.069)	0.156 (0.064)**
Fixed CC + Discount	0.056 (0.060)	0.068 (0.071)	0.013 (0.072)	-0.073 (0.055)
Personalized CC + Discount	0.067 (0.059)	0.158 (0.070)**	0.059 (0.065)	0.004 (0.056)
Observations	395	395	395	395
R-squared	0.19	0.33	0.41	0.33
Control Group Mean (SD)	0.84(0.37)	0.51(0.50)	0.59(0.49)	0.20(0.41)

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Pre-hypertension is equal to one if BP(systolic) > 120 or BP(diastolic) > 80, and to zero otherwise. Hypertension is equal to one if BP(systolic) > 140 or BP(diastolic) > 90, and to zero otherwise. Overweight is equal to one if BMI > 25, and to zero otherwise. Obesity is equal to one if BMI > 30, and to zero otherwise. Village indicators are included in all specifications. Baseline blood pressure and weight measures are included as controls in all regressions. The "ideal" sample is defined as respondents who both believe "it is possible to be healthy with hypertension if blood pressure is frequently monitored" and who trust the service provider.

Table 5: Structural Estimates of Theoretical Model

	(1)	(2)	(3)
	Baseline	Baseline w/ Correlations	Primary
β , mean	0.190	0.103	0.365
β , std. dev.	0.268	0.132	0.395
β , % = 1	0%	0%	15.1%
β , % = 0	47.8%	34.0%	39.8%
$\hat{\beta}$, mean	0.804	0.543	0.795
$\hat{\beta}$, std. dev.	0.04	0.12	0.13
mean $\frac{\hat{\beta}-\beta}{1-\beta}$, $\beta < 0.5$	0.04	0.47	0.66
mean $\frac{\hat{\beta}-\beta}{1-\beta}$, $\beta \geq 0.5$	0.04	0.55	0.66
κ_b , HBP Index	–	–	40.43
κ_b , Sickness Index	–	–	-13.70
κ_b , Male	–	–	45.37
κ_b , Prior Meds	–	–	-13.40
κ_b , Literate	–	–	36.56
κ_c , Employment	–	–	4.00
κ_c , Distance	1.36	0.31	–
κ_{β} , Male	–	–	-0.48
$\kappa_{\hat{\beta}}$, Male	–	–	0.047
ϵ , std. dev. (Rs)	69.2	190.0	197.7
b , mean (Rs)	30	30	61.3
c , mean (Rs)	32.1	30.5	26.5
δ , mean	0.780	0.687	0.234
$\delta * b - c$, 25th percentile	-49.96	-110.76	-55.50
$\delta * b - c$, mean	-3.6	-7.2	-19.3
$\delta * b - c$, 75th percentile	42.76	73.43	17.93
Log-Likelihood	-1399	-1416	-1327
N	1496	1496	1729

Notes: This table presents our structural estimates. Column 1 presents the estimates from a baseline model with limited observable heterogeneity. Column 2 is the same baseline model with additional correlations between unobserved heterogeneity in β and $\hat{\beta}$. Column 3 is our primary model, which incorporates observable heterogeneity on a range of potentially important dimensions (e.g. health and demographics). b and c are perceived benefits and costs of visiting the clinic. ϵ is the unobserved component of perceived benefits. δ is the exponential discount factor. $\delta * b - c$ is the perceived net benefit of visiting the clinic.

Table 6: Consumer Welfare and Social Welfare Impacts

	(1)	(2)	(3)	(4)	(5)	(6)
	Control	Discount	Fixed CC	Personalized CC	Fixed CC, Discount	Personalized CC, Discount
Panel A: Model Fit						
% Purchasing CC	n/a	n/a	7.2	7.2	26.7	26.7
% Commitment Amount > 0	n/a	n/a	100	1	100	1
% Visiting Doctor	4.37	8.37	7.36	7.63	18.00	14.25
Panel B: Baseline, Control						
Mean ΔCW^{ND}	0	1.51	-5.08	-4.12	-2.42	-2.89
Mean ΔCW	0	5.44	-5.08	-4.12	5.68	3.53
Mean ΔSW	0	1.51	8.27	9.53	27.13	39.31
Panel C: Baseline, No Doctor						
% Zero Impact	95.63	91.60	88.42	88.43	73.21	73.29
% Losers	0	0	4.24	4.95	8.71	12.51
Mean ΔCW	6.38	11.82	1.30	2.26	12.06	9.91
Mean if $\neq 0$	145.85	141.08	11.20	19.51	45.06	37.11
Median	0	0	0	0	0	0
Max	421.12	466.50	421.32	421.32	466.5	466.5
Min	0	0	-135	-90	-90	-51.2
10th Percentile	0	0	0	0	0	-45
90th Percentile	0	0	0	0	77.38	53.21
Panel D: Mean ΔCW^{ND}						
b = 500	8.73	19.09	8.73	8.79	33.16	26.69
b = 5000	146.03	285.64	241.09	218.66	602.49	476.67
b = 50000	1534.5	2936.60	2564.02	2317.56	6295.42	4966.20

Notes: This table reports the welfare impacts of each treatment under study. These are computed assuming that our entire population is enrolled in that treatment, and applying our parameter estimates (see Table 5) to study decision-making and subsequent outcomes. We compute welfare from a long run, or $t = 0$ perspective. This means that baseline consumer welfare, CW , in each treatment equals their discounted benefit of visiting the doctor, minus the costs (both pecuniary and non-pecuniary), and minus any pledged commitment amounts lost due to not following through. CW^{ND} nets out discounts given to consumers, such that discount d_i is a transfer and not a consumer benefit. Social welfare SW gives medical benefits net of consumer attendance costs and assumed firm costs.

Table 7: Counterfactual Simulations (Fixed Contracts)

	(1)	(2)	(3)	(4)	(5)
	d = 0	d=7.5	d=15	d=22.5	d=30
Panel A: % Visiting the Clinic					
Discount Only	4.38	5.96	8.37	10.4	18.7
m = 0	7.10	12.4	14.2	15.0	18.7
m = 15	7.36	15.21	18.00	20.24	21.61
m = 30	17.46	21.29	25.66	29.31	33.01
m = 45	17.91	22.01	26.70	30.17	34.89
Panel B: Mean ΔCW^{ND} , Baseline Control					
Discount Only	0	0.88	1.77	1.56	-3.79
m = 0	-4.12	-3.10	-2.82	-0.80	-3.79
m = 15	-5.08	-1.32	-2.45	-2.95	-6.00
m = 30	7.24	6.59	4.76	2.62	-2.46
m = 45	8.38	7.91	6.32	4.40	-2.29
Panel C: Mean ΔCW , Baseline Control					
Discount Only	0	2.22	5.54	8.57	12.92
m = 0	-4.12	-0.31	3.57	9.32	12.92
m = 15	-5.08	2.11	5.65	10.71	13.44
m = 30	7.24	11.38	16.30	22.39	26.24
m = 45	8.38	12.86	18.33	24.76	29.11
Panel D: % Losers, ΔCW , Baseline Control					
Discount Only	0	0	0	0	0
m = 0	4.95	9.6	12.4	13.8	0
m = 15	4.24	6.8	8.72	10.47	15.82
m = 30	0.51	0.73	1.12	1.40	4.42
m = 45	0	0	0	0	0

Notes: This table describes the results from our counterfactual simulations that study alternative bundled discount-fixed commitment contract designs. Specifically, we study five different values of d (0, 7.5, 15, 22.5 and 30) and four values of m (0, 15, 30 and 45). For each potential contract (i.e. d and m combination), we present four statistics where the entire population is offered the contract. These are: (i) the % of consumers visiting the clinic, (ii) the mean consumer welfare impact, netting out discounts, relative to our control treatment, (iii) the mean baseline consumer welfare impact relative to our control treatment, and (iv) the % of losers relative to the control treatment. The simulations assume that only one contract is offered in each case, and thus does not consider menu design with multiple commitment contract options.

A Additional Tables and Figures

Table A.1: Balance Check Across Treatment Groups

Variable	Control	Discount	Fixed CC	Personalized CC	Fixed CC, Discount	Personalized CC, Discount	F-test p-value
Panel A: Demographic Characteristics							
Household Size	5.50	5.29	5.59	5.77	5.33	5.59	0.098
Gender (Female = 0; Male = 1)	0.43	0.37	0.46	0.41	0.38	0.40	0.286
Age	53.4	53.6	54.0	54.0	53.8	53.2	0.978
Literate (Yes = 1, No = 0)	0.35	0.37	0.33	0.37	0.36	0.33	0.811
Household Income (1000 Rs)	120	94.2	92.3	106	99.4	99.2	0.582
Self-Employed Agriculture (Yes = 1, No = 0)	0.18	0.11	0.13	0.15	0.11	0.16	0.137
Panel B: Baseline Health Indicators							
Blood Pressure (Systolic)	143	141	145	143	142	141	0.401
Blood Pressure (Diastolic)	86	86	87	86	85	85	0.311
Weight (kg)	65.6	65.2	64.0	64.9	65.4	64.3	0.749
Panel C: Target Sample Characteristics							
Values Service (Yes = 1, No = 0)	0.66	0.64	0.62	0.71	0.67	0.66	0.317
Trusts Provider (Completely Trust =1, Completely Distrust = 5)	3.12	2.95	3.30	3.16	3.07	3.12	0.418

Notes: This table reports mean values of each variable for every treatment group. The final column reports the joint significance level of treatment indicators in a regression with village indicators. The respondent is classified as valuing the service if he/she believes “it is possible to be healthy with hypertension if blood pressure is frequently monitored.”

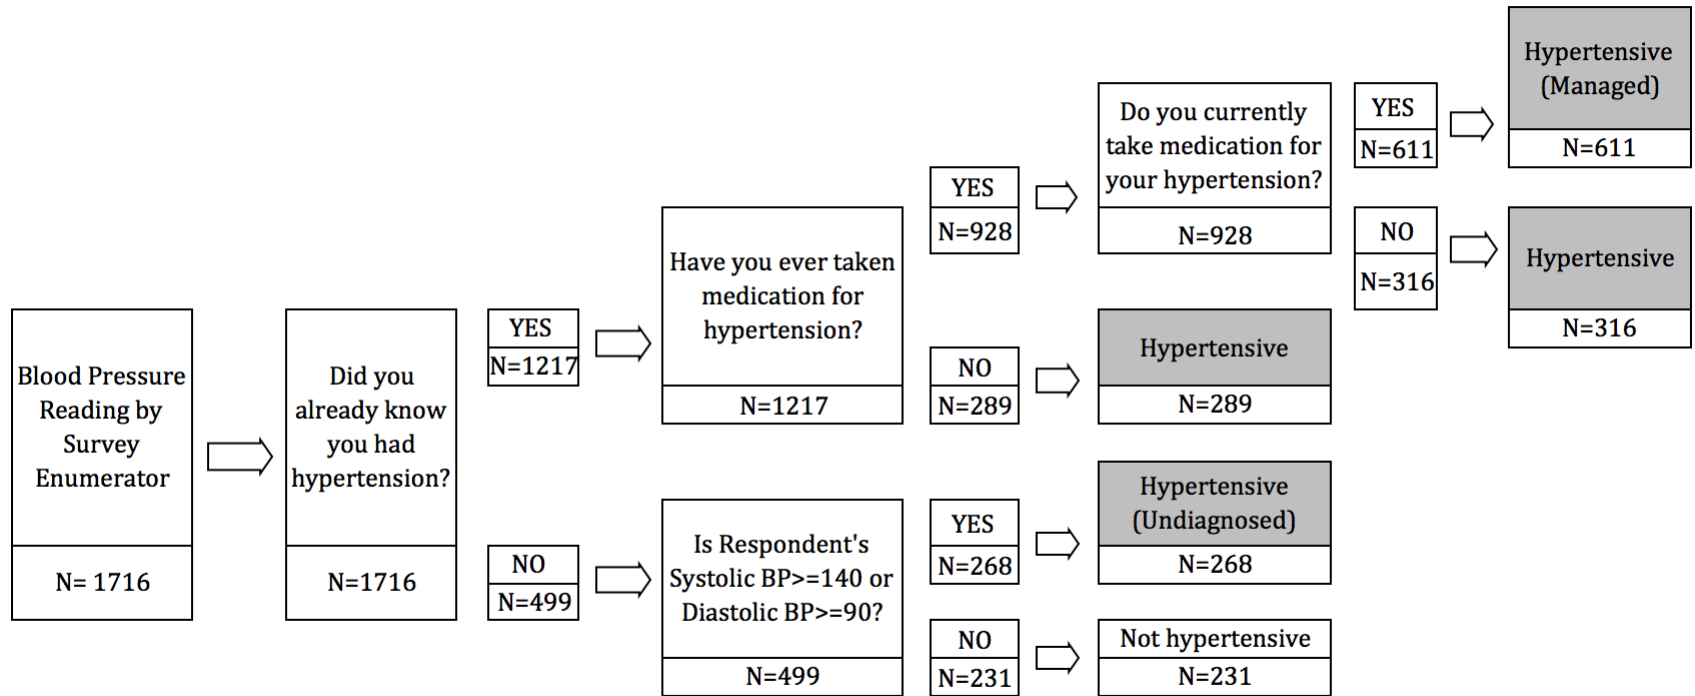
Table A.2: Endline Health Outcomes by Treatment Group: Without Baseline Controls

	(1)	(2)	(3)
Panel A: Full Sample	BP (systolic)	BP (diastolic)	Weight (kg)
Discount	-1.583 (2.197)	-0.845 (1.371)	-1.380 (1.281)
Fixed CC	-0.308 (2.175)	0.070 (1.392)	-2.753 (1.270)**
Personalized CC	-0.567 (2.141)	1.363 (1.407)	-2.090 (1.272)
Fixed CC + Discount	-2.515 (2.142)	-1.230 (1.256)	-0.705 (1.291)
Personalized CC + Discount	-1.041 (2.199)	0.454 (1.428)	-2.349 (1.265)*
Observations	1,512	1,512	1,512
R-squared	0.00	0.00	0.00
Control Group Mean (SD)	139.7 (25.0)	85.0 (15.6)	67.4 (14.1)
Panel B: "Ideal" Sample	BP (systolic)	BP (diastolic)	Weight (kg)
Discount	-4.336 (4.563)	-1.100 (2.499)	0.381 (2.706)
Fixed CC	-6.536 (4.479)	-2.969 (2.806)	-4.158 (2.907)
Personalized CC	-3.082 (4.341)	1.185 (2.527)	-2.525 (2.612)
Fixed CC + Discount	-7.405 (4.394)*	-2.253 (2.274)	-4.266 (2.478)*
Personalized CC + Discount	-6.547 (4.410)	0.167 (2.628)	-2.585 (2.669)
Observations	395	395	395
R-squared	0.01	0.01	0.01
Control Group Mean (SD)	143.4 (28.7)	85.6 (15.2)	69.4 (16.0)

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Village indicators are included in all specifications. The "ideal" sample is defined as respondents who both believe "it is possible to be healthy with hypertension if blood pressure is frequently monitored" *and* who trust the service provider.


Figure A.1: Managing Hypertension

58




Notes: Questions and sample sizes are taken from the baseline survey.

Figure A.2: Baseline and Endline Blood Pressure by Treatment Group




Notes: A normal systolic (diastolic) blood pressure is below 120(80) mmHg. A systolic (diastolic) blood pressure number between 120(80) and 139(89) mmHg is considered to be pre-hypertension. A systolic (diastolic) blood pressure number of 140(90) mmHg or higher is considered to be hypertension. Black circles denote average baseline (BL) and endline (EL) systolic blood pressure measures in each treatment group, while gray circles denote average baseline (BL) and endline (EL) diastolic blood pressure measures in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure A.3: Weight by Treatment Group



Notes: Weight measurements (in kilograms) were carried out by enumerators during baseline and endline surveys. Black circles denote average baseline body weights in each treatment group, while gray circles denote average endline body weights in each treatment group. The vertical gray lines correspond to 95% confidence intervals around coefficient estimates of treatment group indicators in a regression with village fixed effects.

Figure A.4: Preventive Visits Over Time



Notes: This figure shows the percentage of respondents who visited the clinic over time, by treatment group. The vertical line corresponds to the time when reminders were delivered to a randomly selected subset of study participants.

B Estimation Details

This appendix provides additional detail on the estimation methodology for our primary structural approach, described in Section 4 in the text. We estimate the model with a smoothed Accept-Reject simulated maximum likelihood methodology that, given the candidate parameters, matches the predicted decision paths for consumers in the population to their actual decision paths [see e.g. Train (2009) for an econometric discussion and Handel (2013) for another applied example]. Define the set of parameters to be estimated as Θ .²⁴ For a consumer with set of observables X in treatment T we match their sequence of decisions (which depends on T) to the predicted sequences of decisions for candidate parameters, and choose the parameters with the best match given choices across the control and all five treatments.

There are three types of decisions that could enter the decision path for a given individual. The first is choice of commitment contract (if offered). The second is what amount they commit

²⁴In our primary specification, these parameters include $\alpha_\beta, \kappa_\beta, \sigma_\zeta, \tau_1, \tau_2, \kappa_{\hat{\beta}}, \alpha_b, \kappa_b, \sigma_\epsilon, \alpha_C, \kappa_C,$ and δ .

specifically if the contract offered is a personalized contract. The third is whether the visit the doctor or not at $t = 1$. For the random coefficient parameters representing unobserved heterogeneity on each dimension, we take 50 simulated draws, above which the estimation results are stable.

We now describe the likelihood function construction. For doctor attendance, for a given simulated draw of the parameters, the probability someone visits the doctor if in the control group is:

$$P_s(d_{i,v} = 1|\Theta) = \mathbf{1}[C(X_i) + f \leq \beta\delta b_i - \min(\max(\frac{0.01}{\beta\delta b_i - C(X_i) - f}, -0.01), 0.01)]$$

Here, $d_{i,v}$ is an indicator variable that equals 1 if someone goes to the doctor and 0 otherwise. For a given simulated draw, without smoothing, this always equals 1 or 0. Accept-Reject smoothing helps the simulator function continuously by making this outcome probabilistic, but be equal to 1 or 0 in the limit as preferences become stronger for or against visiting the doctor. The smoothing term is the second fractional term, and helps optimization function but does not impact the ultimate results since it approximates the true binary values in the limit.

Given this value of P_s for each draw s , the overall probability of a doctor visit for candidate parameters Θ is:

$$P(d_{i,v} = 1|\Theta) = \frac{\sum_{s=1}^S P_s(d_{i,v} = 1|\Theta)}{S}$$

The construction of $P(d_{i,v} = 1|\Theta)$ for consumers in each of the five treatments is similar, but the per visit fees change accordingly.

For consumers who get to choose a commitment contract, we define $d_{i,c}$ equal to 1 if they accept the contract and 0 if they decline it. We introduce accept-reject smoothing on this decision as well, since it is a binary decision with simulated draws. For example, for a given simulated draw s , consumers who (i) are offered a fixed commitment contract with no discount and (ii) have preferences that satisfy equations 3 and 4:

$$P_s(d_{i,c} = 1|\Theta) = 1 - \max(\min(\frac{0.1}{C(X_i) + f - \hat{\beta}\delta b_i}, 0.01), 0)$$

Again, here the smoothing implemented brings this probably to 1 in the limit as the preference for commitment becomes larger and larger. We implement a similar condition for when a consumer declines the contract offer, in which case $P_s(d_{i,c} = 1|\Theta)$ limits to 0. For a given consumer, the likelihood of accepting a contract offer given Θ is:

$$P(d_{i,c} = 1|\Theta) = \frac{\sum_{s=1}^S P_s(d_{i,c} = 1|\Theta)}{S}$$

For consumers offered either a personalized commitment contract, or a commitment contract with a discount, the process for constructing the likelihood is similar, and follows the equations for accepting the contract laid out earlier in this section.

Finally, for consumers who accept a personalized commitment contract, they have a third decision: how much to commit. This is a continuous decision, except for people choosing the boundary commitment amount of $M_i = 0$, in which case they just pay the lump sum for all recommended doctor visits up front (with no additional commitment). For a given draw of parameters s , consumers who accept a commitment contract have a predicted commitment amount m_i that satisfies

equation 7. We use a uniform kernel likelihood for this decision such that:

$$P_s(m_{s,i,f} = m_{i,f}) = 1 \text{ iff } [m_{s,i,f}[\Theta] - 5 \leq m_{i,f} \leq m_{s,i,f}[\Theta] + 5]$$

Here, if the actual amount committed in the personalized contract is within 5 rupees of the predicted amount $m_{s,i,f}[\Theta]$ for a given draw, it counts as the same outcome, and otherwise counts as a 0, or different outcome.²⁵ The probability that someone commits their actual amount $m_{i,f}$ given Θ is:

$$P(m_{i,f}|\Theta) = \frac{\sum_{s=1}^S P_s(m_{s,i,f} = m_{i,f})|\Theta}{S}$$

Of these three possible decisions, all individuals always decide whether or not to attend the doctor, regardless of the treatment. The other two decisions depend on whether the individual is offered a commitment contract and whether they accept a personalized contract. Given this, we define the log-likelihood contribution for each individual based on their decision paths we need to consider. The log-likelihood contribution for individuals in the control and discount only treatments is:

$$SLL_i(\Theta) = [P(d_{i,v} = 1|\Theta)\mathbf{1}[d_{i,v} = 1] + [1 - P(d_{i,v} = 1|\Theta)]\mathbf{1}[(1 - d_{i,v}) = 1]]$$

For individuals who are offered a commitment contract, but don't accept a personalized contract, their log-likelihood contribution is:

$$\begin{aligned} SLL_i(\Theta) &= [P(d_{i,v} = 1|\Theta)\mathbf{1}[d_{i,v} = 1] + [1 - P(d_{i,v} = 1|\Theta)]\mathbf{1}[(1 - d_{i,v}) = 1]] * \\ &\quad [P(d_{i,c} = 1|\Theta)\mathbf{1}[d_{i,c} = 1] + [1 - P(d_{i,c} = 1|\Theta)]\mathbf{1}[(1 - d_{i,c}) = 1]] \end{aligned}$$

For individual who are offered a personalized contract, and accept that contract, the log-likelihood contribution is:

$$\begin{aligned} SLL_i(\Theta) &= [P(d_{i,v} = 1|\Theta)\mathbf{1}[d_{i,v} = 1] + [1 - P(d_{i,v} = 1|\Theta)]\mathbf{1}[(1 - d_{i,v}) = 1]] * \\ &\quad [P(d_{i,c} = 1|\Theta)\mathbf{1}[d_{i,c} = 1] + [1 - P(d_{i,c} = 1|\Theta)]\mathbf{1}[(1 - d_{i,c}) = 1]] * \\ &\quad P(m_{i,f}|\Theta) \end{aligned}$$

Given this, the simulated log-likelihood function that we maximize over Θ for the population is:

$$SLL(\Theta) = \sum_{i=1}^I SLL_i(\Theta)$$

In order to simplify estimation, we use an outer loop with a coarse grid for the intercepts for the costs and benefits of going to the doctor (α_b and α_c respectively). The final estimates are for the combination of these two parameters that yield estimates with the best likelihood function value.

²⁵In practice, varying this kernel threshold around 5 does not impact the results.