Centre for Health Economics and Policy Analysis

CHEPA WORKING PAPER SERIES

Paper 10-02

Willingness-to-Pay for Parallel Private Health Insurance: Evidence from a Laboratory Experiment

June 17, 2010

Neil J. Buckley [a],+, Katherine Cuff [b], Jeremiah Hurley [b,c,d], Logan McLeod [e], Robert Nuscheler [f], David Cameron [b,c]

[a]Department of Economics, York University
4700 Keele Street
Toronto, Ontario, Canada M3J 1P3
[b]Department of Economics, McMaster University
1280 Main Street West
Hamilton, Ontario, Canada L8S 4M4
[c]Centre for Health Economics and Policy Analysis, McMaster University
1200 Main Street West
Hamilton, Ontario, Canada L8N 3Z5
[d]Department of Clinical Epidemiology and Biostatistics, McMaster University
1200 Main Street West
Hamilton, Ontario, Canada L8N 3Z5
[e]Department of Public Health Sciences, University of Alberta
11350 - 83 Avenue
Edmonton, Alberta, Canada T6G 2G3
[f]Department of Economics, University of Augsburg
Universitätsstr. 16
D-86159 Augsburg, Germany

+ Corresponding author
Department of Economics, York University
1144 Vari Hall
4700 Keele Street
Toronto, Ontario, Canada, M3J1P3
Email: nbuckley@yorku.ca
Phone: (416)736-2100 x30106
Fax: (416)736-5987

McMaster University

CHEPA WORKING PAPER SERIES

The Centre for Health Economics and policy Analysis (CHEPA) Working Paper Series provides for the circulation on a pre-publication basis of research conducted by CHEPA faculty, staff and students. The Working Paper Series is intended to stimulate discussion on analytical, methodological, quantitative, and policy issues in health economics and health policy analysis. The views expressed in the papers are the views of the author(s) and do not necessarily reflect the views of the Centre or its sponsors. Readers of Working Papers are encouraged to contact the author(s) with comments, criticisms, and suggestions.

NOT FOR CITATION WITHOUT PERMISSION

Abstract

Debate over the effects of public versus private health care financing has been, and continues to be, active in both academic outlets and policy circles. Theoretical literature on parallel health care financing is often built on untested behavioural assumptions and the empirical evidence generally depends upon the institutional details of the specific health care systems under analysis. This paper contributes to the literature on parallel health care finance by developing and executing a revealed preference laboratory experiment based on the theoretical model of parallel health care finance in Cuff et al. (2008). The theoretical model involves individuals in the labour force with varying severities of illness who demand health care from a limited supply of health care resources. Health care resources are purchased by the public sector and rationed free of charge to individuals, or purchased by individuals through a private insurance market. The general theoretical model is converted into a discrete experimental representation of a large-scale economy where individuals are price takers, the probability of receiving public health care is exogenous and the willingness-to-pay (WTP) for private health insurance is elicited from subjects. The experimental design includes two within-subject factors based on the theoretical model: the public sector rationing rule (rationing based on need or severity versus rationing based on a random allocation) and the probability of being publicly treated (high versus low). The experimental design also includes two between-subjects treatments based on the frame of the experiment (neutral frame versus health frame) and on the distribution of private health insurance prices (high prices versus low prices). The results show the public system's allocation rule and the probability of receiving health care from the public system both significantly affect an individual's WTP for private health insurance in the predicted direction, although the WTP values tend to be above the actual theoretical predictions. When the public system allocates health care based on need, the average WTP is lower than under random allocation. A higher probability of receiving health care from the public system elicits a lower WTP regardless of how the public system allocates health care. It is also found that when the public system allocates health care based on need, the WTPs are significantly higher under a neutral frame than a health frame.

Keywords: Health Care Finance, Rationing, Framing Effects
JEL Classifications: I11, C91

ACKNOWLEDGEMENTS

We acknowledge helpful comments from Michel Grignon, Stuart Mestelman, Andrew Muller and other members of the New and Emerging Team project on experimental methods applied to health economics at McMaster University. In particular we would like to acknowledge Dale Marsden and Aleksandra Gajic for research support. We would also like to thank conference participants from the 2009 Canadian Economics Association meetings and seminar participants from University of Innsbruck, Social Science Research Centre Berlin and University of Regensburg for helpful comments. This research was funded by the Canadian Institutes of Health Research (Grant # 76670). The views expressed are those of the authors alone.

1.0 INTRODUCTION

The 2005 Supreme Court of Canada decision in Chaoulli vs. Quebec (Supreme Court of Canada 2005) brought Canada one step closer to a health care system with parallel private health insurance. Parallel private health insurance consists of private insurance for services covered by the public insurance system. It can be used to purchase privately provided services for which a person is eligible for public coverage. A number of Canadian provinces explicitly prohibit such insurance; others inhibit the development of parallel private insurance through payment, and other related, regulations (Boychuk 2006).[1] The Supreme Court ruled that in the presence of unreasonable wait times Quebec's prohibition against private insurance for publicly insured services violated the Quebec Charter of Rights and Freedoms. After the court decision, Quebec passed legislation permitting the purchase of private insurance for three publicly financed services with long wait times (hip replacement, cataract surgery and knee replacement). The province later expanded legislation to cover a set of fifty publicly financed services (Quebec Bill 33 and its modifications). Lawsuits similar to Chaoulli have been filed in other provinces, including Ontario and Alberta. And professional organizations such as the Canadian Medical Association have begun exploring the possibility of offering such private insurance.[2]

The Supreme Court's decision re-ignited a broader debate about the merits and de-merits of introducing parallel private insurance in Canada. Advocates claim the introduction of parallel private finance can reduce wait times and financial pressure within the public system and increase access to, and quality of, needed health care (Globerman and Vining 1998; Crowley 2003; Montreal Economic Institute 2005; Esmail 2006). Opponents argue that parallel private finance will draw health care resources away from the public system, leading to increased wait times in the public system, reduced access for low-income individuals and reduced quality of public health care services (Yalnizyan 2006; Canadian Health Coalition 2006). Both sides in the debate often presume that relaxing regulatory and legal constraints on parallel private insurance and parallel privately financed provision more generally will beget a large and thriving private insurance sector; both assume demand will be forthcoming and debate only the impacts of such insurance.

---

[1] Parallel private insurance is distinct from supplemental private insurance to cover services not insured by the public health care system (e.g., dental care, drugs for some individuals), which a majority of Canadians hold.
[2] Prior to the introduction of Medicare, physician-sponsored insurance constituted an important element in the private insurance industry.

However, even where permitted, the size of parallel private markets vary widely: In the UK, the proportion of people holding parallel private insurance has remained steady for two decades at about 11% of households, even in the presence of a subsidy in the early 1990s; in Australia the share of households with parallel private insurance fell steadily until the combination of a large premium subsidy and a regulated premium structure stabilized coverage beginning in 2000; in Scandinavia, the private insurance market has never exceeded 1-2% of expenditures. Studies of the demand for parallel private insurance in these and other countries consistently find that demand is price inelastic, strongly positively correlated with income, and depends in part on political attitudes (Mossialos and Thomson 2004). But such studies are unable to test explicitly the impact of what is perhaps the most important factor driving demand for parallel private insurance: people's perceived functioning of the public system. Such insurance is only worthwhile if one perceives that a necessary service won't be available in a timely manner or at all in a moment of need in the public system. A small number of studies have tried to identify the impact of public wait times on the demand for private insurance, but because such studies rely on observational data they suffer from possible reverse causation: higher demand for such insurance may cause longer wait times (Besley et al. 1999; Besley 2001; Propper and Maynard 1989; Propper 2000).

The purpose of this paper is two-fold. First, we take an alternative approach to studying the impact of parallel public and private finance: laboratory-based revealed choice experiments. Laboratory experiments allow one to study in a controlled manner how changes in an institutional environment affect the behaviour of participants. And although economic experiments are used primarily to test economic theory, they have also informed the development of many important public policies. Experimental studies have been influential, for instance, in designing markets for pollution trading permits (e.g., results from Cason (1995) and Cason and Plott (1996) influenced the Environmental Protection Agency's emission allowance auction). Stated-preference experiments are widely used in health care, especially to inform economic evaluations and health technology assessment (e.g., Drummond et al. 2005; Ryan et al. 2008; Tsuchiya and Dolan 2009). However, the application of revealed-choice laboratory experiments to health care systems research has been very limited, despite their growing use in many other areas of economics. By using laboratory experiments to investigate questions of parallel health care finance, we are able to address some of the possible limitations of observational studies including methodological concerns in the absence of good natural experiments, lack of necessary high-quality data, and concerns about generalizability (since

observational findings are often dependant upon the institutional details of the specific health care system(s) under analysis while laboratory environments involve the parsimonious use of institutional details that clearly map subject decisions to experimental outcomes).

Second, we test specific predictions regarding people's willingness-to-pay (WTP) for parallel private insurance derived from a theoretical model of parallel private finance (Cuff et al. 2008). The theoretical model involves individuals with varying severities of illness who demand health care from a limited supply of health care resources. Health care resources are purchased by the public sector and rationed free of charge to individuals, or purchased by individuals through a private insurance market. The model makes specific predictions regarding individuals' WTP for parallel private insurance and, in particular, how WTP varies with changes in the way in which the public sector rations health care and with changes in the likelihood that someone seeking care in the public system will receive treatment. The general theoretical model is converted into a discrete experimental representation of a large-scale economy where individuals are price takers, the probability of receiving public health care is exogenous and the WTP for private health insurance is elicited from subjects.[3] The experimental design includes two within-subject treatments: (1) the public sector's allocation rule determining which individuals will receive health care (allocation according to need versus random allocation) and (2) the probability of receiving health care through the public sector (high versus low). The design also includes a between-subjects treatment with respect to the exogenous distribution of private health insurance prices (high prices versus low prices). Finally, because this is the first health-related experiment of this type of which we are aware, and because preferences with respect to health may differ from preferences more generally, we use a second between-subjects treatment to test for differences between the health scenario and a formally identical "neutral" scenario framed purely in terms of a monetary gains and losses.

We find that individuals' willingness to pay for private insurance varies with the public sector's allocation rule and the probability of receiving health care in the public system as predicted, although point estimates of WTP tend to be above the theoretical predictions. Despite subjects knowing that the price of private insurance is drawn independently each period, the price draws from previous periods significantly affect WTP for private health insurance. Finally, WTP differs across the health and neutral frames only when public sector

[3] While eliciting WTP for private health insurance is crucial for testing the predictions made by Cuff et al. (2008), estimation of WTP for private insurance is also important to the parallel financing debate when determining whether demand is sufficient to sustain an optional private health care system.

care is rationed based on need. In particular we find that, under a needs-based rationing rule, individuals are willing to pay less for private insurance when the decision is framed in terms of health care than when it is framed using neutral terminology.

The paper is organized as follows. Section 2 provides a summary of the Cuff et al. (2008) model of parallel public and private health care finance. Section 3 discusses our laboratory implementation of the theoretical model. We present our experimental results in Section 4 and offer a discussion of these results in Section 5.

2.0 THEORETICAL MODEL

To model the interaction between parallel public and private health care financing and the allocation rule of the public health care system, Cuff et al. (2008) assume a continuum of individuals that differ along two dimensions: income (Y) and severity of illness (s). An individual knows her income but is unaware of her severity of illness. In the base model, income and severity are independently distributed.[4] The support of the income distribution is the interval $[0,\overline{Y}]$ and for the severity distribution the unit interval $[0,1]$. Both the distributions and support of income and severity are known by everyone. All individuals are sick in the population but differ in the degree or severity of their illness.

An individual's illness can be cured immediately and fully (regardless of severity level) with the receipt of one unit of health care (e.g., a combination of health care providers or equipment required to deliver health care). If an individual fails to obtain treatment she incurs a monetary loss from illness given by sY. Notionally, the length of her illness is proportional to her severity and, because she is unable to work and generate income while ill, she loses income sY. If an individual receives health care she is cured immediately and suffers no loss of income. An individual can receive health care from one of two sources: from the public system that offers free care but cannot guarantee access to care; or from a private insurer that can guarantee access to care for a positive price. Conditional on receiving health care, there is no difference in public and private health care: health outcomes are identical and both public and private health care services are produced using one unit of a health care resource.

[4] Cuff et al. (2008) also consider the case of negative correlation between income and severity.

The health care resources (H) used to produce health care services are fixed in supply; the fixed supply is insufficient to treat all individuals in society.[5] The market for health care resources is competitive and both public and private providers bid for contracts with suppliers of health care. The public sector bids for health care resources according to its ability-to-pay, as determined by the public health care budget (B). Private insurers bid for health care resources according to their willingness-to-pay, which is based on individuals' WTP for private insurance that guarantees access to care regardless of severity level. Thus, the two sectors compete directly for the limited health care resources, resulting in a single market-clearing price (P) for the health care resource such that all resources are allocated across the two sectors.[6]

Because the supply of health care resources is limited, the public sector must ration access to publicly financed health care. The analysis considers two different public-sector rationing rules: needs-based allocation and random allocation. Needs-based allocation assumes the public sector allocates scarce health care resources to beneficiaries in accordance with need (defined by severity of illness). Need-based allocation is the stated objective of many publicly financed health care systems (van Doorslaer et al. 1993), including Canada's (Marchildon 2005). The second allocation rule, random allocation, assumes the public sector allocates its scarce health care resources to beneficiaries independent of severity.[7] No public insurance plan deliberately allocates its health care services randomly, but every system contains elements of random allocation because of informational problems.[8] One could think of this allocation rule as a 'lower bound' of other forms of allocation mechanisms that use some (possibly imperfect) information about individuals such as their need for health care (Wijkander 1988).

[5] Cuff et al. (2008) discuss the relaxation of the fixed supply assumption.

[6] Private insurers only charge a single price for health insurance equal to the market-clearing price. In this sense, the private health insurance contracts are actuarially fair since the private insurers earn zero profit although the price is not equal to each person's expected loss.

[7] This form of rationing could also be called uniform since all individuals regardless of income or severity face the same likelihood of being treated. Equivalently, it could be called rationing with no prioritization. Both allocation rules are independent of income. This would no longer be true with needs-based allocation if income and severity were correlated.

[8] A recent survey of MRI clinics in Canada, for instance, found that: "Few sites have documented criteria to guide triaging decisions. No site had a method of quality assurance to determine whether or not the prioritization was being performed consistently. Thus, it is entirely possible that patients with the same medical indication for an MRI examination, at the same centre, could be placed in different prioritization categories, with very different wait times . . . This inconsistency in defining prioritization categories and the considerable variation in the number of categories leads to significant inconsistencies in access to MRI from site to site even within a given province." (Emery et al. 2009, p. 82).

Cuff et al. (2008) determine the equilibrium in the market for health care resources by first calculating each individual's WTP for private insurance. At the time an individual decides whether to purchase private insurance, she does not know her severity of illness or the probability of receiving health care in the public system. Assume individuals are expected-utility maximizers. A privately insured individual pays a premium P and in return receives guaranteed health care regardless of her severity. The individual's resulting payoff is $Y-P$. In contrast, an individual in the public health care system has the following expected payoff $\pi^e Y+(1-\pi^e)[1-E(s\,|\,\text{no care})]Y$ where π^e is the individual's expected probability of receiving health care in the public system and $[1-E(s\,|\,\text{no care})]Y$ is the individual's expected income if not treated in the public system. An individual's maximum WTP for private insurance is the amount of money that equates her expected payoff without private insurance to her (guaranteed) payoff with private insurance:[9]

(1) $$WTP=(1-\pi^e)E(s\,|\,\text{no care})Y\,.$$

This expression makes two things clear: first, the lower is an individual's expected probability of treatment in the public system, the higher is her maximum WTP; second, the higher the expected loss from not receiving public health care, the higher the individual's maximum WTP.

Both the probability of receiving health care in the public system and the expected severity conditional on not receiving public health care depends on how the public system allocates its scarce health care resources. When the public sector uses random allocation, the probability of receiving public health care is the same for all individuals and does not depend on individual severities. The expected severity of those who do not receive public care under this allocation rule is simply the unconditional expected severity level. Assuming a uniform severity distribution, this can be written as, $E_R(s\,|\,\text{no care})=E(s)=1/2$.[10] Inserting the unconditional expected severity level into equation (1), we obtain the WTP for private insurance under random allocation:

(2) $$WTP_R=(1-\pi_R^e)Y/2.$$

[9] The individual's income minus their maximum WTP is simply the individual's certainty equivalent, i.e., the amount of income that leaves the individual indifferent between facing the gamble of remaining in the public health system where there is some risk of not being treated. Given risk neutrality, the certainty equivalent is equal to the expected payoff of the gamble.

[10] We use indexes R and N to denote variables under random and needs-based allocation, respectively.

When the public sector uses needs-based allocation, it allocates its limited health care resources to treat only those individuals with the highest severities; consequently, only individuals with severity above some threshold receive public health care. Thus, the expected probability of receiving public health care translates into an expected severity threshold that, assuming a uniform severity distribution, can be written as $s^e = 1 - \pi^e$. Those with severities equal to or above s^e are expected to receive public health care, while those with severities below s^e are not expected to receive public health care. The expected severity conditional on not receiving public care is $E_N(s \mid \text{no care})] = s^e/2 = (1-\pi_N^e)/2$ and the WTP for private insurance is:

$$WTP_N = (1-\pi_R^e)^2 Y/2. \tag{3}$$

Starting from equations (2) and (3), Cuff et al. (2008) determine the equilibrium under each allocation rule by solving a system of three equations.[11] The first equilibrium condition is that the health care resource market clears. The second equilibrium condition is that the health care resources acquired by the public sector must be consistent with rationed demand for public health care and the third equilibrium condition is that expectations are confirmed in equilibrium. From this system of three equations, Cuff et al. (2008) show the equilibrium price for health care resources, which is also equal to the price of private health insurance, is lower, and the probability of receiving health care in the public system is higher under needs-based allocation than under random allocation.

The intuition for this result is straightforward. From equations (2) and (3), we know that for equal (expected) allocation probabilities, an individual's WTP for private insurance is lower when the public sector allocates according to need rather than randomly. This follows because under needs-based allocation individuals receive public health care if their severity is above the severity threshold, thus avoiding the largest losses if not treated publically; but under random allocation individuals risk not receiving care even if they have a high severity, so large losses are not necessarily avoided if not treated publically. This higher WTP under random allocation causes private insurers to bid more for health care resources, generating greater competition for health care resources and a correspondingly higher equilibrium price under random allocation than under needs-based allocation. Because the public sector's nominal budget is fixed, the public sector can purchase fewer health care resources, causing stricter rationing of health care. Hence, the probability of receiving health care services from the public sector is lower

[11] See Appendix A for details.

under random allocation. The scope for a private health insurance market in parallel to the public system is smaller under needs-based allocation than under random allocation.

To summarize, compared to random allocation, needs-based allocation creates two effects that reduce WTP for private insurance that translate into a lower equilibrium price for private insurance. The first is a direct effect: because the highest losses are avoided under needs-based allocation, the individual's WTP for the private insurance is lower. The second is an indirect effect. The lower WTP for private insurance under needs-based allocation shrinks the private insurance market, freeing resources for the public system. The public system now has more health care resources, increasing the probability of receiving care in the public sector, reducing the expected loss even further and, with it, the WTP for private insurance.

3.0 LABORATROY IMPLEMENTATION

One of the main purposes of the experiment is to test the WTP predictions in the model of Cuff et. al. (2008). In this regard, we test two effects. First, we test the above-noted direct effect of a change in the public sector's allocation rule. To do this we control the indirect effect by fixing the probability of receiving health care in the public system, π, exogenously12 and then vary the experimental treatment of allocation rule to be either random allocation or needs-based allocation. We do this using a within-subject design that exposes all subjects to both rationing rules. Second, we test the effect of a change in the probability of treatment in the public system on an individual's WTP. Independent of the rationing rule, the model predicts a negative association between the probability of treatment (π) and WTP. To do this we introduce the treatment variable 'public sector probability of receiving care' that can assume two values, high $(\pi = 0.7)$ and low $(\pi = 0.3)$. Again, we test this using a within-subject design.

In the experiment, an individual must submit her WTP for private insurance before she learns her severity of illness. Since severity is a random draw, subjects face a situation of decision-making under uncertainty. To elicit individuals' true WTP we use an incentive-compatible mechanism developed by Becker, DeGroot and Marshak (BDM, 1964).13 The BDM

[12] As the experiment is only concerned with eliciting the individual's WTP for insurance rather than determining the market equilibrium, there is no concern here about whether the public insurer's budget is balanced or possible strategic interactions between subjects. Further, income heterogeneity is not needed and to simplify the experiment and to reduce possible noise all individuals are assigned the same income.

[13] Since we focus on the use of revealed-preference instruments in this study, an alternative to the BDM mechanism would have been to use the procedure suggested by Holt and Laury (2002). The Holt-Laury mechanism was originally

mechanism has been widely used in economics and is designed to induce a subject to truthfully reveal the certainty equivalent to a lottery (Davis & Holt (1993), Kahneman et al. (1990), Irwin et al. (1998), Rutstrom (1998), and Shogren et al. (2001)).[14] The original BDM mechanism calls for subjects to first be given a unit of some commodity. Subjects are then asked to state their minimum price for selling the good. Third, subjects are notified that a "buy price" is randomly drawn from some distribution. If the subject's stated sell price is less than or equal to the randomly drawn buy price, the subject receives the buy price and gives up the good. If the stated sell price is above the buy price, then the subject keeps the good.

We modified the mechanism as follows to suit our context. Given some probability of public care π and an allocation rule, the expected utility of relying on the public system is given by $\pi u(Y)+(1-\pi)u([1-E(s\,|\,\text{no care})]Y)$, where $u()$ is the individual's utility from income. In contrast, net income with private insurance of price P is $Y-P$ and the corresponding utility is $u(Y-P)$. The market price for insurance is drawn randomly from some non-empty interval. It seems natural to consider the interval $[0,Y]$ as the interval, where Y is the subject's income and the maximum a subject can pay for private health insurance. The distribution function of market price is denoted by $F(P)$ with corresponding density $f(P)>0$. Denote a subject's stated WTP, or bid, as $\tilde{p}$. The subject purchases insurance and receives payoff $u(Y-P)$ whenever the randomly drawn price P is less than or equal to her WTP for insurance, $\tilde{p}$. The subject purchases the insurance at price P, so she pays her actual bid $\tilde{p}$ only if it equals P. With probability $1-F(\tilde{p})$ the price draw yields a market price that exceeds the subject's stated WTP so she does not purchase private insurance and remains in the public system. The expected utility of a subject who faces the BDM mechanism and submits a bid $\tilde{p}$ is therefore:

(4) $$Eu=(1-F(\tilde{p}))(\pi u(Y)+(1-\pi)u([1-E(s|\text{no care})]Y))+\int_0^{\tilde{p}} u(Y-P)dF \quad .$$

The BDM mechanism can be shown to be incentive compatible as follows. Consider a subject who overstates his or her true WTP. Doing this lowers her expected utility because it will sometimes cause her to purchase private health insurance when the market price is above her true valuation. Analogously, if the subject understates his or her WTP she will sometimes be

designed to elicit risk preferences, e.g., coefficients of risk aversion (see Anderson and Mellor (2008) for a recent application to health economics and see Harrison and Rutström (2008) for a detailed comparison of these two risk elicitation techniques). Possible concerns about the use of the BDM mechanism in our environment are discussed below but the BDM mechanism was used due to its similarity to the market mechanism described by the theoretical model reported in Cuff et al. (2008).

[14] The certainty in our model is given by $Y-P$ so that the mechanism is designed to give the right incentives to truthfully reveal the WTP for a switch from uncertainty (public sector) to certainty (private sector).

unable to purchase private health insurance when her true valuation exceeds the market price. To demonstrate incentive compatibility analytically, differentiate (4) with respect to $\tilde{p}$ to obtain the following first order condition:

(5) $$f(\tilde{p})[u(Y-\tilde{p})-(\pi u(Y)+(1-\pi)u((1-E(s|\text{ no care}))Y))]=0.$$

The expression in square parentheses is exactly the situation described in the theory with a linear utility function. The second-order condition is satisfied, so that bidding the true WTP is, indeed, incentive compatible: the individual maximizes her expected utility by bidding her true WTP. Incentive compatibility holds as long as the price distribution has full support; the actual distribution used is – in theory – irrelevant.

Although theoretically appealing, three concerns about the incentive compatibility of the BDM mechanism in practice have been raised: (i) subject bids are sensitive to the boundaries of the distribution for market price; (ii) subjects' preferences may not conform to expected utility theory; and (iii) subjects may not be able to compute their optimal bid. We address these concerns in turn.

First, bids elicited using the BDM mechanism are sensitive to the choice of endpoints of the distribution over which the market price is drawn (Bohm et al. 1997). If the upper bound is set higher than a subject would be expected to pay, the stated WTPs are inflated. In our work, the upper bound is set to a subject's assigned income (500 lab dollars), which is the maximum price she can bid (and therefore would be expected to bid). Because the support is constant across treatments we are confident that even if such bias is present it would affect treatments equally so that WTP comparisons across regimes are still valid.

Second, the BDM assumes that individuals are expected-utility maximizers, so the BDM mechanism may not be incentive compatible if a subject's preferences do not conform to expected-utility theory. This is particularly a problem if the value of the good is uncertain as this can lead to preference reversals (Karni and Safra 1987). But more recently a similar concern has been raised even when the value of the good is certain (Horowitz 2006). In order to minimize potential problems from this feature of the BDM mechanism, we focus on differences in WTP across treatments, noting that the absolute magnitudes of stated WTP may be suspect.

Third, given the complex nature of the lottery presented, a subject may not be able to compute her optimal bid. Harrison (1992) raises concerns about the role of payoff dominance in experimental settings. If there is payoff dominance, it may be strategic for a subject to search for his or her optimal bid by choosing a WTP starting value and to modify it slowly in the direction which leads to increasing the subject's payoff each period. Irwin et al. (1998)

investigate the interaction between payoff dominance (reward saliency) and cognitive effort in a decision task. They find it is not necessary to have a steep payoff schedule to induce optimizing behaviour when a subject is able to deduce their optimal strategy from the initial information provided. However, if a subject must search for the optimal strategy, the steepness of the payoff schedule can make a difference. We set a favourable environment for subjects to be able to deduce their optimal strategy by using exogenously set probabilities of receiving care in the public sector and by using simple uniform distributions for severities.

The experiment was conducted as follows: Subjects were told to imagine that they were workers who get sick. All subjects were told their initial income (500 lab dollars ($L)) for each period and that they would lose a random and initially unknown proportion, s, of their income if they did not receive health care. Subjects were told that their severity of illness, s, is a random number drawn from a uniform distribution on the unit interval. A subject had two health care options: private health insurance (with guaranteed access to health care but positive price) and public health care (with positive probability of not receiving health care but is free of charge), where public and private health care are, in all other respects, identical and always cure perfectly so that receiving care means going back to work and not losing income. The BDM mechanism was explained, including an illustration of why individuals should bid their true valuation.

The experiment ran for 20 decision periods. At the beginning of each period, subjects were given an income of $L500, the probability of receiving public health care, and the public sector's allocation rule. Subjects were then asked to indicate their WTP for private health insurance. After submitting their WTP, subjects were told the randomly drawn market price for private health insurance, and whether their WTP was high enough to purchase private health insurance. If their WTP was greater than or equal to the market price for private insurance they purchased private insurance at the market price and they were told their random severity draw and their final payout for the period ($L500 less the market price for insurance). If their WTP was less than the market price, they did not purchase insurance and remained in the public system. They were then told their randomly drawn severity level as well as whether they were selected to receive free public health care, based on the public sector allocation rule. If they were selected to receive care, they retained their full income for the period ($L500). If they were not selected to receive care, they lost a proportion of their income equal to their severity level $(\$L500 \cdot s)$, and their final payout was $\$L500 \cdot (1-s)$. Subjects were instructed that their decisions affect their real earnings; at the end cumulative payouts in lab dollars were converted

to Canadian dollars using the same conversion rate for all subjects.[15]

All subjects faced 10 periods each of needs-based allocation and random allocation by the public sector with 5 periods each of high and low probability of receiving public care under each allocation rule, for a total of 20 decision periods. Any change in the allocation rule or the probability of receiving public care was explicitly announced to subjects. Respective WTP decisions were grouped into four sequential sets of five decision periods each. The order a subject experienced each set of 5 decision periods was systematically varied between subjects to control for order effects. Random draws for severity levels and random health care allocations where made independently across all subjects and decision periods. Hence submitted WTP data are independent across individuals in the experiment but each subject's WTP decisions are likely correlated across decision-making periods.[16]

Table 1 presents the unique theoretical predictions of WTP for private health insurance derived from equations (2) and (3) for the four within-subject treatments (income is set to 500 lab dollars for each subject). Each subject faced all four possible within-subject treatments, in varying order, as shown in Table 2.

We faced one additional issue in implementing the BDM mechanism. The theoretically predicted WTP values (Table 1) were sufficiently low that drawing random market prices from a uniform distribution between zero and the subject's income implied that subjects, acting as predicted, would rarely purchase private insurance. This could cause them to become frustrated and disengage from the experiment. To inhibit this 'frustration factor', we instead drew market price randomly from a skewed beta distribution (Figure 1 illustrates the beta distributions used), where the parameters of the beta distribution were set to increase the likelihood of a subject being able to purchase private insurance.[17] Furthermore, we implemented a second between-subjects treatment using two mean levels of the randomly drawn BDM market price. The low-frustration treatment set the parameters of the beta distribution so that on average 50% of the random draws would be below the theoretically predicted WTPs. The high-frustration treatment set the parameters of the beta distribution so that on average 20% of the random draws would be below the theoretically predicted WTPs. Frustration level was held constant for each subject for all 20 decision periods..

[15] The average payout for the experimental session (lasting on average 70 minutes) was CDN$17.60 per subject, which is above the minimum wage.

[16] In our empirical analysis we take this (potential) correlation into account by using a random effects panel series regression model.

[17] It is worth noting that although the BDM price distribution was skewed, its support was still the full range of each subject's income and thus the theoretical incentive compatibility property of the BDM mechanism remains unchanged. A test of whether this skewed price distribution causes effects similar to those found by Bohm et al. (1997) when changing the endpoints of the price distribution will be left for the results section.

Finally, to test whether the health frame of our scenarios had salience, we also implemented a neutral frame scenario using a between-subject design. The neutral frame was structurally identical to the health frame but was presented using neutral terms. In the neutral frame, subjects were given an "endowment" (not an "income") and asked to indicate their WTP for a "guaranteed loss avoidance" (not "insurance") to prevent losing part of their endowment because of an "uncertain adverse event" (not a "random severity of illness"). The price of the guaranteed loss avoidance was random and if their stated WTP was greater than or equal to the subsequently revealed price, subjects purchased the guarantee of loss avoidance out of their endowment at the revealed price and avoided any loss due to the uncertain adverse event. Subjects who did not purchase the guaranteed loss avoidance were informed of their random "loss level" and received free loss avoidance depending on the described rule: random or highest losses avoided. The chance of avoiding any loss (i.e., 30% chance or 70% chance) and the free loss avoidance rule were revealed prior to subjects submitting their WTP. After submitting their WTP, subjects were informed of their individual loss level and whether they were selected for free loss avoidance, thus keeping their full endowment for the period. Nowhere in the neutral frame were the terms "health", "health care", "insurance", "sickness", or "severity" used. The predicted WTP are identical to those associated with the health insurance scenario (Table 1). Table 3 presents information on these two between-subjects treatments and provides a breakdown of the sample sizes in each treatment.

Subjects were recruited using classroom notices, campus posters, university websites and word of mouth. Thirty-seven sessions were administered between February 5 and July 10, 2008 in the McMaster Experimental Economics Laboratory on McMaster University campus. In total, two hundred and sixty-three subjects completed the experiment. Sessions lasted between 45 to 90 minutes with an average of around 70 minutes, including approximately 15 minutes of on-screen computer-narrated instructions. After the experimental component of the session ended, subjects completed a brief on-screen demographic survey. Subjects were then individually paid their cash earnings in private. The experiment used the University of Zurich's Z-tree software program (Fischbacher 1999). The experiment was approved by the McMaster University Research Ethics Board.

4.0 DISCUSSION

We begin by comparing the observed mean WTP values against the theoretically predicted values for two periods only, periods 1 and 20 (see Table 4). Period 1 WTP values are of

interest because they are independent and free from the effects of previous severity draws and previous insurance price draws from the BDM mechanism, while period 20 WTP values reflect experience gained under all four combinations of allocation rule and probability of receiving public care. For each period the mean WTP values vary among the different experimental treatments as predicted by the theory: the highest WTP occurs under random allocation with a low probability of public-sector treatment; the next highest occurs under needs-based allocation with a low probability of public-sector treatment; third is random allocation with a high probability of public-sector treatment; and the lowest WTP occurs under needs-based allocation with a high probability of public-sector treatment. The differences between the treatment combinations within each period are statistically significant (see the Appendix B for more details). In all cases the observed values exceed the predicted values, a point we will return to below. Finally, for both random and needs-based allocation, when the probability of public-sector treatment is high the observed WTP values are stable between the first and twentieth periods (i.e., $96 vs. $90 for random, and $73 vs. $77 for needs-based), displaying little evidence of a learning effect. In contrast, under both allocation rules when the probability of public-sector treatment is low the WTP values differ notably across periods 1 and 20, albeit in different directions (i.e., $166 vs. $190 for random, and $144 vs. $103 for needs-based). It appears that when the insurance stakes were higher, because of the low probability of public-treatment, subjects at first did not appreciate the implications of the differing allocation rules, bidding similar amounts under each ($144 and $166 in period 1), but experience provided insight regarding the greater value of insurance under random allocation than under needs-based allocation that reduced the chances of large losses associated with high-severity draws, causing the WTPs to diverge over time as predicted ($190 vs. $103 in period 20).

Table 5 presents, for each combination of allocation rule, public sector care probability, “frustration level” and frame, the mean WTP for private insurance and the proportion of subjects that purchased private insurance.18 Again, the order in the mean WTP values among the different experimental treatments (i.e., following down any given column in Table 5) conforms to the theoretical predictions but the mean WTP values in general exceed the predicted values (i.e., the observed values in any given row of the table). The only case for which the differences between the observed and predicted values are not statistically significant is random allocation with a low probability of public-sector treatment.19

[18] These summary values exclude data from the first two of five replications for each allocation rule-treatment probability combination due to the learning effects evidenced in tables 4 and 6.
[19] For this case however, the predicted WTP is $175, the largest prediction of all the experimental treatments.

Figures 2 through 5 illustrate mean per period WTP values by allocation rule and probability of receiving care in the public sector, supporting the statistics presented in Table 5. Each of these figures presents the theoretical prediction and the mean WTP values under both frames and frustration levels. It is clear that WTP levels are above the predicted values across the board, especially in the needs-based allocation treatment with a high probability of treatment in the public sector.

The proportion of subjects who successfully purchase private insurance is predicted to be 0.5 in the low-frustration treatments and 0.2 in the high-frustration treatments as determined by the specific price distributions used for the BDM mechanism. The proportion of subjects who purchased private insurance was close to the 0.5 predicted level for the low-frustration treatments in all cases but one — needs-based rationing with a high probability of public-sector treatment, for which the proportions were 0.61 and 0.71 respectively for the health and neutral frames. In contrast, the proportions of insured subjects were well above the 0.2 predicted level for the high-frustration treatments, with again the highest values (0.58 and 0.61) occurring with needs-based allocation and a high probability public-sector treatment.20

Table 6 presents the results of a random-effects panel regression model for WTP for private insurance. The model includes variables representing the public sector allocation rule, the probability of public sector treatment, the frame, the frustration-level for the BDM mechanism and associated control variables and interaction terms. The assumption of random-effects at the individual level fit our experimental context well, for which any subject-specific effects (individual heterogeneity) are independent of changes in the experimental treatments (exogenous regressors). Errors between subjects are independent because there was no interaction among subjects in the experiment. Two specifications are presented. Specification 1 is based on 19 periods of data (all 20 periods less the first period because of the lagged variables included in the model). Because the subjects faced the four experimental treatments (2 allocation rules and 2 probabilities of public-sector treatment) in varying orders, we tested for order effects (see variables “Periods 1-5”, “Periods 6-10”, etc.). We also included variables to control for the order of the five replicated decisions within each treatment (“Replication 1”, etc.). Model specification 1 suggests a significant order effect for Periods 1-5 and Periods 11-15 and significant replication effects for the first replication. This difference is not surprising and is consistent with learning effects early within each allocation rule (which

[20] Outlier behaviour in the treatment with needs-based rationing and a probability of public treatment of 0.7 is likely due to the fact that the optimal WTP value was the lowest of all treatments, equal to approximately 5% of each subject’s income. Evidence that the theoretical model is poor at describing subject choices in the presence of low probability losses.

began in periods 1-5 and periods 11-15) and early within each treatment group of five replicated decisions (more specifically the first replication). Model specification 2 was run based on WTP decision from replications 2 to 5 from Periods 6-10 and Periods 16-20 only to control for these learning/order effects. The results are similar across both specifications, evidence that while the order effects were significant they were likely small in magnitude. We focus below on the results from specification 2.

The public sector allocation rule has a significant effect on stated WTP for private insurance (p-value of 0.00). The results for specification 2 in Table 6, for example, indicate that in the health frame with low-frustration BDM price draws and a low probability of public treatment subjects bid on average $37 less for private insurance when the public system allocates health care according to need than when it allocates randomly. The significant interaction term between needs allocation and neutral frame signifies that this allocation rule affect is cut in half (a change of approximately $18) under a neutral frame rather than the health frame. A second significant interaction term between needs allocation and a high frustration price level indicates that subject WTP values are only approximately $24 ($36.98-$13.35) less when the public system allocates according to need than when it allocates randomly when insurance prices tend to be high.

As predicted, the probability of being treated in the public system also has a significant effect on the stated WTP for private insurance. Subjects, on average, were willing to pay $74 less for private insurance when the probability of treatment in the public sector was high rather than low regardless of allocation rule, frame or frustration level (and interactions with all these factors are not statistically significant). This finding proves that the probability of treatment in the public system is a very strong and salient component of WTP for private insurance. Not only is the effect of a change in public treatment probability statistically significant and large in magnitude, it is not affected by other aspects of the decision environment such as the allocation rule, the decision-making frame and the level of prices.

Both the frame effect and the BDM frustration effect are statistically significant only under needs-based rationing (p-value=0.022). Under needs-based allocation WTP values were significantly lower in the health frame than the neutral frame; they were also significantly lower under the low-frustration mechanism than the high-frustration mechanism. A priori, our hypothesis was that neither frame nor frustration level would affect WTP for private health insurance since neither has any effect on the financial incentives, so this is an unexpected finding. The higher WTP values found in the high frustration treatment compared to the low frustration treatment could mean that a change in the price density (i.e. making it a more

asymmetric distribution by increasing the probability of high prices) has a similar effect on elicited WTPs to that which Bohm et al. (1997) found when the endpoints of the price distribution were changed. Finally, while last period's severity level does not affect WTP for insurance (p-value=0.345), last period's random market price has a significant affect on WTP (p-value=0.000). Again, this is an unexpected finding given that subjects were told that price draws are independent across periods.

5.0 DISCUSSION

These findings have important implications both for the potential role of laboratory experiments in health-sector research and for health policy. We consider each in turn.

As noted already, this is, to our knowledge, the first attempt to create an experimental environment to test alternative systems of health care finance. The results indicate that laboratory experiments have considerable potential to provide insight into a range of institutional design issues in the health sector. The subjects appear to have understood what in many respects was a complex environment foreign to Canadians (who are most familiar with the public insurance system only); learning effects are present but bids stabilize after a couple of periods in each parameterization of the environment; and the mean bids for private insurance varied as expected.

Although no framing effects were expected, differences between the health and neutral frames arose where they are most plausible: needs-based allocation. Subjects appear to have treated the random allocation as akin to a financial gamble, regardless of how we framed the setting. Needs-based allocation is both more familiar and expected in a health frame. In the health frame, needs-based allocations cause no cognitive dissonance and subjects responded more to the incentive effects of needs-based rationing in the health frame than in the neutral frame. Framing all issues of relevance to health settings is not possible in laboratory experiments, especially the full cost (including pain and suffering) to an individual of being ill or injured. In this work we focused only on monetary costs. Future experimental work will have to explore alternative ways to represent the full array of costs and benefits in health settings.

The experimental design allowed us to investigate the role of two factors that would be almost impossible to study using observational data. No data document either the extent to which individuals perceive that the health care system allocates resources according to need, or their perception of the probability that the system would be unable to provide needed services. Yet, these unmeasured perceptions are potentially major drivers of demand for an alternative to the publicly financed health care system.

Our results also provide further caution regarding the BDM elicitation mechanism. The frustration level imposed by the distribution of randomly drawn market prices for private insurance, while not predicted to have an effect by the theoretical model being tested (Cuff et. al. (2008)), did affect subjects stated WTP. This may be an artifact of the BDM mechanism. As noted above, Bohm et al. (1997) find elicited WTP values are sensitive to the choice of endpoints of the distribution over which the market price is drawn. Moving an endpoint of a uniform distribution is equivalent in some respects to increasing the probability mass at the upper or lower tail of the distribution. It appears that changing the shape of the market price distribution by making it skewed toward higher values to induce “high-frustration” appears to have artificially increased our subjects WTP. A similar argument could be made as to why we observe WTPs higher than the theoretical predictions even under low frustration. This sensitivity of WTP bids to aspects of the elicitation mechanism that should not influence bids is a cause for concern with the BDM mechanism. Because our analysis focused on differences between bids in alternative treatment environments, as long as any such bias was fixed it was removed by taking differences. But in settings for which the absolute level of WTP is of primary importance, such bias may compromise findings. Testing hypotheses regarding the absolute level of WTP for private insurance in the context of the Cuff et al. (2008) model might best be achieved using an experimental setting in which the market price for health care resources is endogenously determined along with the probability of treatment in the public sector. We leave such interests for future work.

From a policy perspective, the results highlight some critical, relatively neglected issues in the debate about parallel private health insurance. The policy discussion has focused on wait times per se as a driver of demand for privately financed care. Our results emphasize that also important is the process of rationing. If rationing is needs-based, people are less likely to pursue private financing for two reasons. First, they are insured against the highest losses associated with a severe condition. But second, anecdotal evidence suggests that for reasons of both efficiency and equity, people are more tolerant of a wait if they know that the people ahead in the queue are in greater need. But if they perceive that resources are not in fact allocated based on need, both self-interest and sense of inequity encourage use of a private alternative. A critical policy implication is that if system managers can demonstrate to the public that the system does allocate according to need, the viability of privately financed services will be reduced: few things will inhibit the growth of private insurance more than the good performance of the public system.

References

Anderson, L.R., and J.M. Mellor (2008). Predicting Health Behaviors with an Experimental Measure of Risk Preference. *Journal of Health Economics* 27: 1260-1274.

Besley T, Hall J, Preston I. 1999. The Demand for Private Health Insurance: Do Waiting Lists Matter? *Journal of Public Economics* 72:155-81.

Besley T. 2001. The Demand for Health Care and Health Insurance. *Oxford Review of Economic Policy* 5(1):21-33.

Becker, G. M., M. H. DeGroot, and J. Marschak (1964). Measuring Utility By a Single-response Sequential Method." *Behavioral Science*, 9:226-32.

Bohm, P., J. Linden and J. Sonnegard (1997). Eliciting Reservation Prices: Becker-Degroot-Marschak Mechanism vs. Markets. *The Economic Journal.* 107:1079-1089.

Boychuk G. 2006. Provincial Approaches to Funding Health Services in the Post-Chaoulli Era. *Unpublished Document.*

Canadian Health Coalition Factsheet: Private Health Insurance: Busting the Myths. Canadian Health Coalition: Ottawa; 2006.

Cason, T.N. (1995). An experimental investigation of the seller incentives in the EPA's emission trading auction, *American Economic Review* 85, 905-22.

Cason, T.N. and C.R. Plott. (1996). EPA's New Emissions Trading Mechanism: A Laboratory Evaluation, *Journal of Environmental Economics and Management* 30, 133-160.

Crowley B. Top Ten Things People Believe about Canadian Health Care but Shouldn't. Atlantic Institute for Market Studies: Halifax, NS; 2003.

Cuff, K., J. Hurley, S. Mestelman, A. Muller and R. Nuscheler (2008). Public and Private Health Care Financing with Alternate Public Rationing Rules. *McMaster University Centre for Health Economics and Policy Analysis Working Paper No. 07-07*, Hamilton, ON

Davis, D.D. and C. A. Holt (1993). *Experimental Economics*. Princeton, NJ: Princeton University Press.

Drummond M, Sculpher M, Torrance G, O'Brien B, Stoddart G. 2005. *Methods for the Economic Evaluation of Health Care Programmes*. Oxford: Oxford University Press.

Emery D, Forster A, Shojania K, Magnan S. 2009. Management of MRI Wait Lists in Canada. *HealthCare Policy* 4(3):76-86.

Esmail N (2006). Setting the Record Straight on Private Health Care. Fraser Institute: Vancouver, BC.

Fischbacher, U. (1999). z-Tree – Zurich toolbox for readymade economic experiments. Working Paper 21, Institute for Empirical Research in Economics, University of Zurich.

Globerman S and A. Vining (1998). A Policy Perspective on Mixed Health Care Financial Systems. *Journal of Risk and Insurance*; 65(1): 57-80.

Harrison, G.W. and E. Rutström (2008). "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (eds.), Risk Aversion in Experiments (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Harrison, G.W. (1992). Theory and Misbehavior of First-Price Auctions: Reply. *American Economic Review*, 82(5): 1426-1443.

Holt, C.A., and S.K. Laury (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92(5): 1644-1655.

Horowitz, J.K. (2006). The Becker-Degroot-Marschak Mechanism is not Necessarily Incentive Compatible, even for non-Ransom Goods. *Economic Letters*, 93: 6-11.

Irwin, J.R., G.H. McClelland, M. McKee, W.D. Schulze, and N.E. Norden (2007). Payoff Dominance vs. Cognitive Transparency. *Economic Inquiry*, 36: 272-285

Karni, E. and Z. Safra (1987). "Preference Reversal" and the Observability of Preferences by Experimental Methods. *Econometrica*, 55(3): 675-685.

Kahneman, D., J.L. Knetsch, and R.H. Thaler (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6): 1325-1348.

Marchildon G. 2005. *Health System in Transition: Canada*. Copenhagen: WHO on behalf of the European Observatory on Health Systems and Policies.

Montreal Economic Institute (2005). Using Private Insurance to Finance Health Care. Montreal Economic Institute: Montreal, QC.

Mossialos E, Thomson S. 2004. *Voluntary Health Insurance in the European Union*. World Health Organization: Geneva.

Propper C. 2000. The Demand for Private Health Care in the UK. *Journal of Health Economics* 19(6).

Propper C, Maynard A. 1989. *The Demand for Private Health Care and the Demand for Private Health Insurance in Britain*. Centre for Health Economics Discussion Paper 53, York.

Rutstrom, E.E. (1998). "Home-Grown Values and the Design of Incentive Compatible Auctions." *International Journal of Game Theory*. 27(3): 427–441.

Ryan M, Gerard K, Amaya-Amaya M. 2008. *Using Discrete Choice Experiments to Value Health and Health Care*. Dordrecht, Netherlands: Springer.

Shogren, J.F., S. Chob, C. Koo, J. List, C. Park, P. Polo, and R. Wilhelmi (2001). Auction mechanisms and the measurement of WTP and WTA. *Resource and Energy Economics*. 23: 97–109.

Supreme Court of Canada. 2005. *Chaoulli vs. Quebec*. Ottawa.

Tsuchiya A and P. Dolan. 2009. Equality of what in health? Distinguishing between outcome egalitarianism and gain egalitarianism. *Health Economics,* 18(2): 147-159.

Van Doorslaer E, Wagstaff A and F Rutten (1993). Equity in the Finance and Delivery of Health Care: An International Perspective. Oxford University Press: Oxford.

Wijkander, H. (1988). "Equity and Efficiency in Public Sector Pricing: A Case for Stochastic Rationing." *Econometrica* 56(6): 1455–1465.

Yalnizyan A (2006). Getting Better Health Care: Lessons From (and for) Canada. Canadian Centre for Policy Alternatives: Ottawa.

Table 1: Theoretical Predictions of Levels and Changes to WTP by Within-Subject Treatments (for an Income of 500 Lab Dollars per Period)

Allocation Rule	Random		Needs-based	
Probability of Public Treatment	0.3	0.7	0.3	0.7
Theoretical Prediction for WTP for Private Insurance	$175	$75	$122.5	$22.5
Predicted change in WTP with change in Allocation Rule	$52.5		-$52.5	
Predicted change in WTP with change in Probability of Public Treatment	$100	-$100	$100	-$100

Table 2. Experimental Design: Order of Within-Subject Treatments

Treatment #	1	2	3	4	5	6	7	8
Allocation Rule for Periods 1-10	Random				Need			
Probability of Receiving Care from the Public System, Periods 1-5	0.3	0.7	0.7	0.3	0.3	0.7	0.7	0.3
Probability of Receiving Care from the Public System, Periods 6-10	0.7	0.3	0.3	0.7	0.7	0.3	0.3	0.7
Allocation Rule for Periods 11-20	Need				Random			
Probability of Receiving Care from the Public System, Periods 11-15	0.3	0.3	0.7	0.7	0.3	0.3	0.7	0.7
Probability of Receiving Care from the Public System, Periods 16-20	0.7	0.7	0.3	0.3	0.7	0.7	0.3	0.3

Table 3. Experimental Design: Between-Subjects Treatments

Frame	Health		Neutral	
Frustration Level	High	Low	High	Low
Number of Subjects	56	56	75	76

Note: Due to an error in the instructions, data from periods 11-20 for 47 and 48 of the subjects from the neutral framed high and low frustration treatments, respectively, were eliminated from the analysis presented in this paper. The error did not contaminate data from periods 1-10 for these subjects.

Table 4: Observed and Predicted Mean WTP, Periods 1 and 20

Allocation Rule	Random		Needs-based	
Probability of Receiving Public Health Care	Low (0.3)	High (0.7)	Low (0.3)	High (0.7)
Theoretical Prediction	$175	$75	$122.50	$22.50
Mean WTP (Period 1)	$166	$96	$144	$73
Mean WTP (Period 20)	$190	$90	$103	$77

Table 5: Mean WTP, by Experimental Treatment

				Frustration Level			
				Low		High	
Allocation Rule	**Probability of Public Treatment**		**WTP Prediction**	**Health**	**Neutral**	**Health**	**Neutral**
Random	Low (0.3)	*WTP*	175	182	173	208 +	191
		Insured		0.48	0.47	0.32 +	0.29 +
	High (0.7)	*WTP*	75	89 +	97 +	114 +	116 +
		Insured		0.51	0.46	0.33 +	0.33 +
Needs-Based	Low (0.3)	*WTP*	122.5	135 +	154 +	179 +	192 +
		Insured		0.51	0.56	0.35 +	0.44 +
	High (0.7)	*WTP*	22.5	49 +	61 +	77 +	112 +
		Insured		0.61 +	0.71 +	0.58 +	0.61 +
			Predicted Proportion Insured:	0.5	0.5	0.2	0.2

Note: The "WTP" is the mean willingness-to-pay for private health insurance and "Insured" is the proportion of subjects who purchased private health insurance.

+ : denotes an estimate that is significantly higher than the prediction at the 5% level (no symbol means no difference)

Table 6: Random Effects Regression Results of WTP for Private Health Insurance*

	All Periods		Periods 7-10 and 17-20 only	
Specification:	**1**		**2**	
Needs-based Allocation Rule	-36.45	***	-36.98	***
	(3.67)		(5.84)	
Neutral Frame	-6.81		-5.05	
	(10.43)		(13.18)	
High Probability of Public Treatment	-75.12	***	-74.01	***
	(3.68)		(7.35)	
High Frustration	16.04		9.05	
	(12.24)		(16.02)	
Lag Price	0.13	***	0.11	***
	(0.02)		(0.03)	
Lag Severity	11.12		6.81	
	(7.02)		(7.22)	
Intercept	145.41	***	144.16	***
	(9.81)		(12.31)	
Interaction Terms				
High Frustration*Lag Price	0.00		0.04	
	(0.02)		(0.03)	
Neutral*High Probability	13.72	***	3.76	
	(3.55)		(7.56)	
Neutral*Need	13.65	***	18.11	***
	(4.14)		(5.70)	
Neutral*High Frustration	-1.36		-3.24	
	(14.26)		(17.48)	
Need*High Probability	3.71		0.43	
	(3.54)		(7.32)	
Need*High Frustration	18.03	***	13.35	**
	(4.05)		(5.83)	
High Probability *High Frustration	-6.79	*	1.69	
	(3.84)		(8.17)	
Periods 1-5	17.42	**		
	(6.77)			
Periods 6-10	8.53		6.42	
	(6.62)		(6.93)	
Periods 11-15	13.34	*		
	(7.33)			

*Note: Table continues on the next page.

Table 6 (continued): Random Effects Regression Results of WTP for Private Health Insurance

	All Periods	**Periods 7-10 and 17-20 only**
Specification:	**1**	**2**
Replication 1	-15.25 ***	
	(2.80)	
Replication 2	-4.45	0.29
	(2.79)	(3.44)
Replication 3	-1.06	4.21
	(2.79)	(3.44)
Replication 4	-0.43	4.30
	(2.79)	(3.44)
Order interactions with Lag Price	Yes	Yes
Order Interactions with Lag Severity	Yes	Yes
Order interactions with High Frustration	Yes	Yes
N	4309	1724
R-squared (within)	0.370	0.339
R-squared (between)	0.071	0.169
R-squared (overall)	0.254	0.246

Note: Standard errors are in parentheses and significance levels are identified: * p<0.10, ** p<0.05, *** p<0.01.

Figure 1. Beta Distributions for High and Low Price Frustration Levels

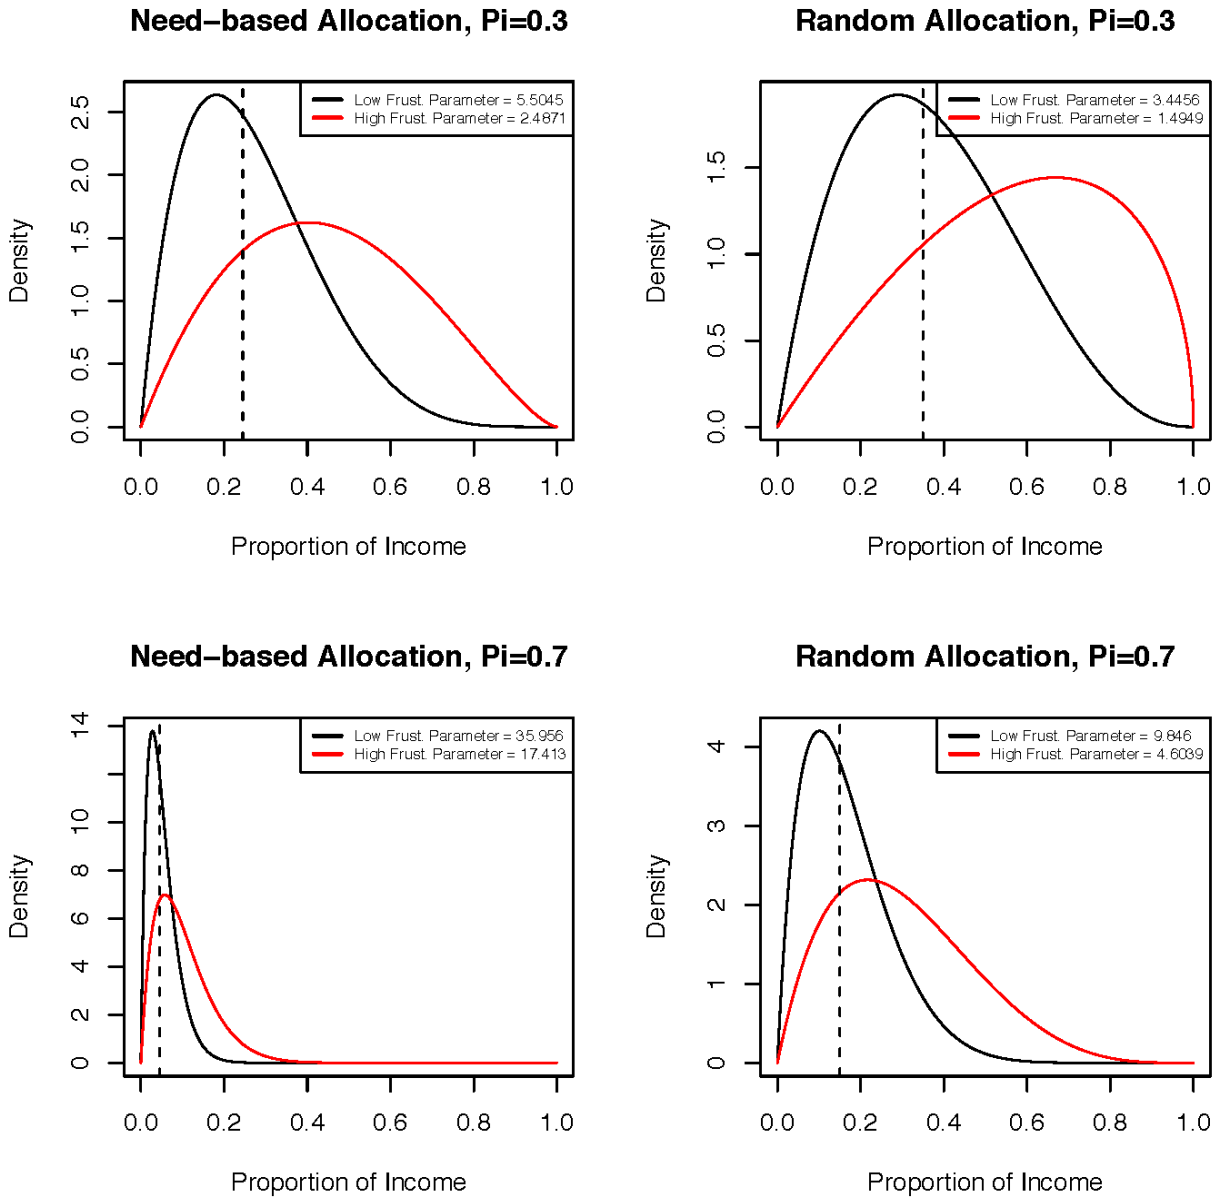

Figure 2. Mean WTP for Insurance under Random Allocation and Probability of Public Treatment of 0.3

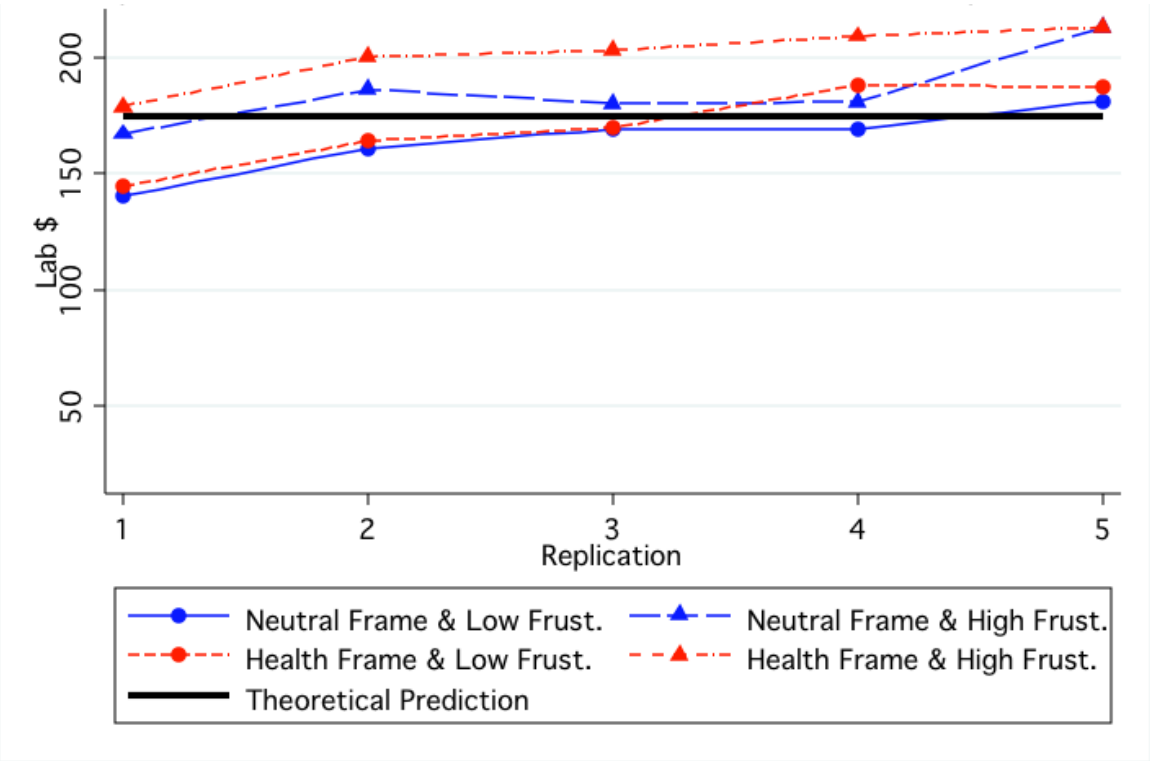

Figure 3. Mean WTP for Insurance under Random Allocation and Probability of Public Treatment of 0.7

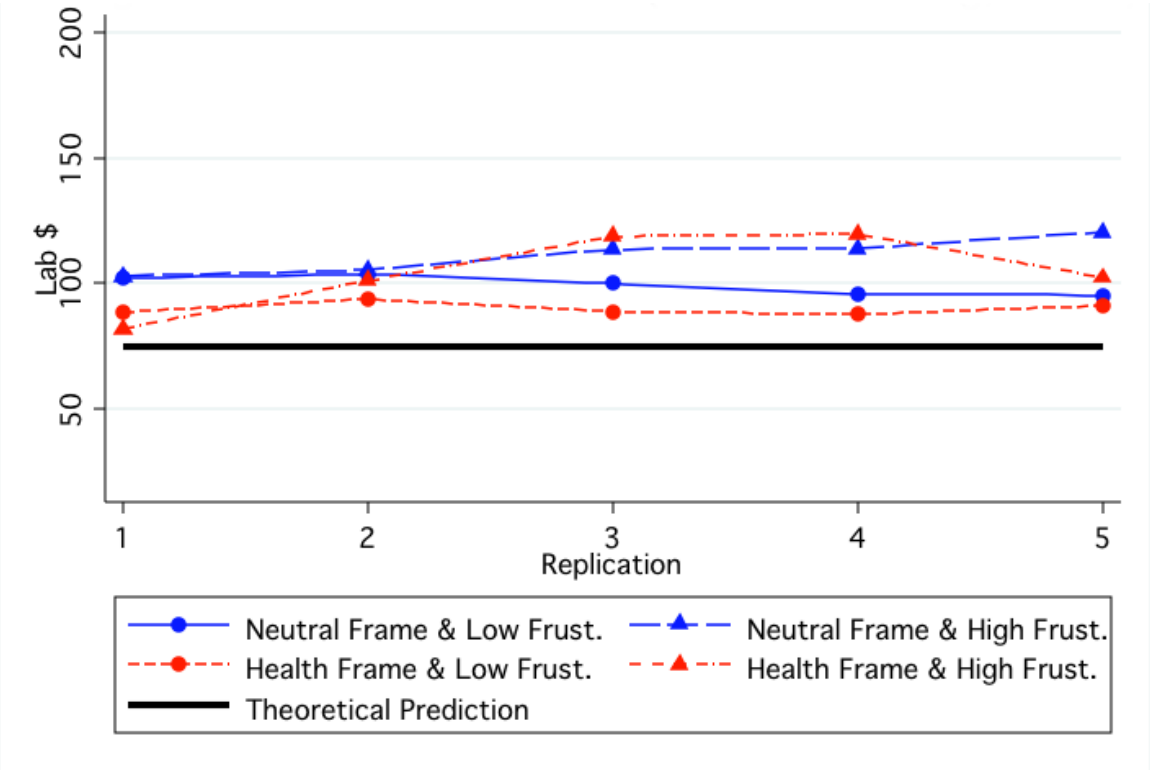

Figure 4. Mean WTP for Insurance under Needs-based Allocation and Probability of Public Treatment of 0.3

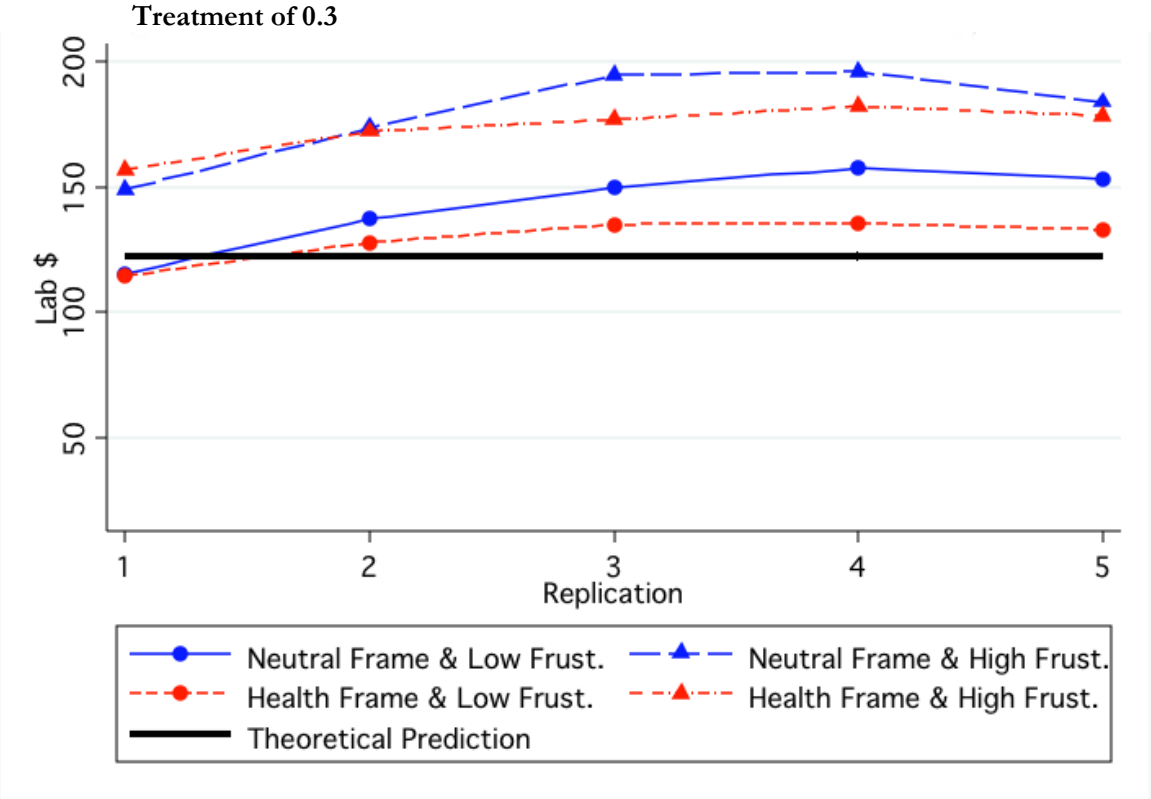

Figure 5. Mean WTP for Insurance under Needs-based Allocation and Probability of Public Treatment of 0.7

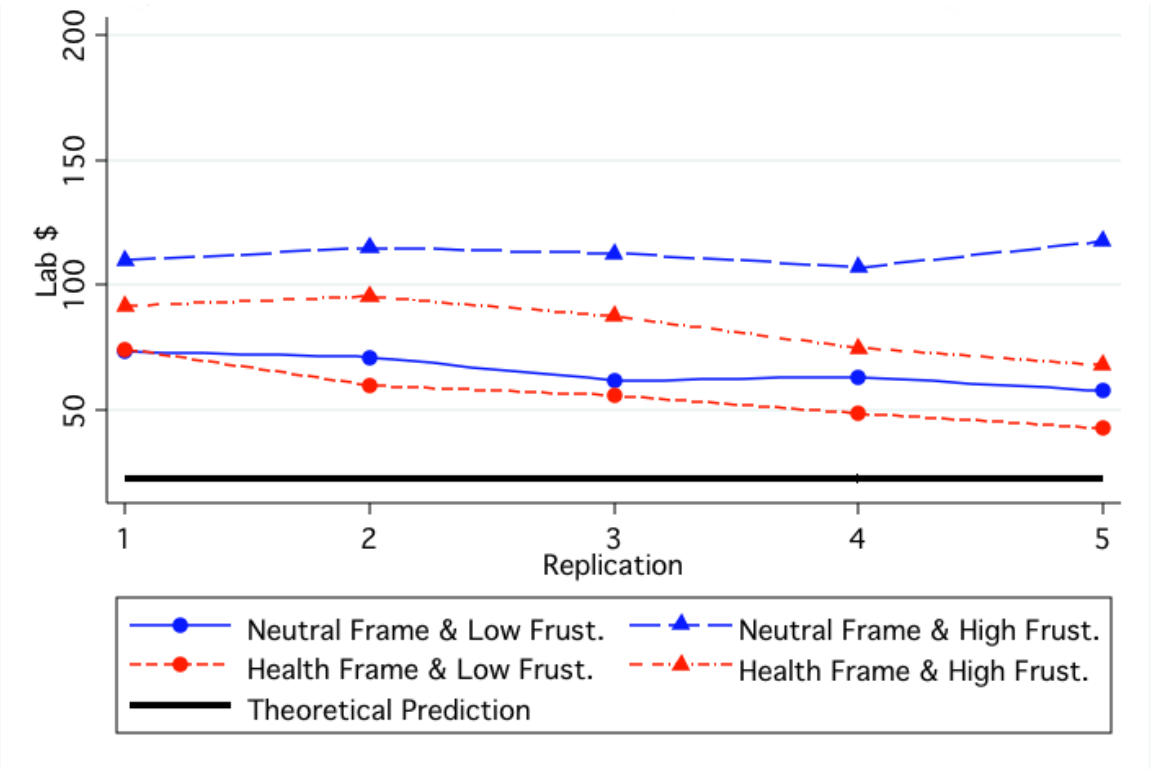

Appendix A

Suppose the price for private health insurance is P.[21] Equating P and WTP (equations (2) and (3)) and solving for income yields critical income levels $Y_i(P)$, $i = [C(public), E(private)]$ at which an individual is indifferent between both public and private insurers. Since the WTP for insurance is (strictly) increasing in income, the demand for private insurance at price P is given by the share of the population with an income of at least $Y_i(P)$. Assuming a uniform distribution of income, demand for private insurance is $1 - Y_i(P)/\bar{Y}$. The equilibrium conditions can now be written as:

(A1) $$\underbrace{1 - \frac{Y_i(P_i)}{\bar{Y}}}_{\text{private demand for resources}} + \underbrace{\frac{B}{P_i}}_{\text{public demand for resources}} = \underbrace{H}_{\text{total supply of resources}}, \qquad \text{[health care resources]}$$

(A2) $$\underbrace{\pi_i}_{\text{rationing probability}} \cdot \underbrace{\frac{Y_i(P_i)}{\bar{Y}}}_{\text{demand for public care (individuals w/o insurance)}} = \underbrace{\frac{B}{P_i}}_{\text{public suppy of health care}}. \qquad \text{[public sector]}$$

The solution of the two systems of equations

(A3) $$P_R = \frac{(1-H)\bar{Y}}{2} > \frac{(1-H)\bar{Y}}{2} - \frac{B}{1-H} = P_N, \qquad \text{[price]}$$

(A4) $$\pi_R = \frac{2B}{2B + (1-H)^2\bar{Y}} < \frac{2B}{(1-H)^2\bar{Y}} = \pi_N. \qquad \text{[public allocation prob.]}$$

[21] In the following we assume that expectations are confirmed.

Appendix B

Table B.1: Robust OLS Regression Results of WTP for Private Health Insurance

	Period 1		Period 18		Period 20	
Need (Allocation Rule)	-40.50	*	-58.13	**	-57.29	**
	(0.08)		(0.03)		(0.02)	
0.7 (Prob. of Public Treatment)	-85.18	***	-79.20	***	-102.36	***
	(0.00)		(0.00)		(0.00)	
Neutral (Frame)	-19.07		-25.99		2.68	
	(0.37)		(0.45)		(0.93)	
Neutral*Need	32.86		65.02		38.46	
	(0.25)		(0.18)		(0.38)	
Neutral*0.7	26.66		4.37		7.60	
	(0.33)		(0.92)		(0.86)	
Need*0.7	19.86		15.97		21.80	
	(0.51)		(0.61)		(0.44)	
Need*0.7*Neutral	-36.33		-59.96		-46.96	
	(0.33)		(0.29)		(0.38)	
Intercept	176.43	***	195.31	***	189.14	***
	(0.00)		(0.00)		(0.00)	
N	263		168		168	
F(7, 255)	10.3					
F(7, 160)			9.3		13.8	
Prob > F	0.00		0.00		0.00	
R-squared	0.212		0.236		0.327	
Root MSE	73.0		85.3		79.2	

* p<0.10, ** p<0.05, *** p<0.01

Note: p-values are reported in brackets below coefficients