

ECONOMIC INCENTIVES FOR A HEALTHY DIET: A COMPARISON OF POLICIES IN A CANADIAN CONTEXT

Peter Kennedy

Melanie Toner

Department of Economics

University of Victoria

27 April 2010

ABSTRACT

This paper examines the potential impact of policies to promote dietary health in the context of a tax-financed universal health care (TFUHC) system, as found in Canada and many other countries, in which a universal level of treatment is set by government policy. We construct a model in which a low quality diet raises the risk of disease, and disease in turn causes a loss of labor productivity. Low quality diets are less expensive than high quality diets, so dietary choices are worse for lower income households. The effect of disease can be partly offset by medical treatment. We show that dietary choices under TFUHC are distorted by a moral hazard problem which leads to diets that are lower in quality than is first-best. We then examine three different policy interventions to address the dietary distortion: a reduced level of treatment; risk-based premiums for health care; and a quality-based tax on food. We calibrate the model with Canadian data on type 2 diabetes, and derive the net benefit and distributional impact for each policy.

1. INTRODUCTION

Diets high in calories but low in quality are the source of some serious health problems in wealthy countries. In particular, diets rich in saturated fat, refined sugar and salt are closely associated with a variety of serious chronic diseases including, cardiovascular disease, hypertension, type 2 diabetes, arthritis and some types of cancer. The social cost of these preventable diseases is substantial.

A variety of factors lie behind poor quality diets. In this paper we focus on the link between dietary quality, the relative price of low and high quality calorie sources, and the moral hazard associated with health insurance. It is well documented that foods high in saturated fat and sugar are cheaper – on a per calorie basis – than healthier foods.¹ There is also good evidence that low income households obtain a lower fraction of their calories from high quality calorie sources, such as fruits and vegetables, than do high income households [Kirkpatrick and Tarasuk (2003), Heart and Stroke Foundation of Canada (2006)]. Moreover, excess body fat – closely associated with low quality diets – is negatively correlated with household income [Le Petit and Berthelot (2005)]. These empirical relationships point to a strong overall link between relative food prices and dietary behavior across income groups. We construct a model founded directly on that link. Households who choose low quality diets in our model do so because low quality food is less expensive on a per calorie basis. This yields an outcome in which lower income households have poorer quality diets and higher rates of diet-related disease than do higher income households.

We use this model of dietary choice to examine how diet outcomes and disease incidence might be affected by various policy options for the promotion of healthier diets. Our base case – against which we assess the policy alternatives – is a stylized representation of the Canadian health care system. This tax-financed universal health care (TFUHC) system provides a universal level of medical treatment – determined by government policy – to all agents who become ill. Treatment users do not pay a price for that treatment. Instead, treatment costs are funded out of income tax revenue. Many countries around the world have similar systems. The separation of treatment use and

¹ For example, see Drewnowski and Specter (2004), Critser (2003) and World Health Organization (2002). This is especially true of “fast food”, when the cost of time is taken into account; see Chou *et. al.* (2004).

treatment funding under TFUHC creates a moral hazard problem: agents do not take an efficient level of preventative care, reflected in dietary choices, to reduce the risk of illness. Consequently, there is too much low quality food consumed and too much diet-related disease.

An additional distortion arises if income tax revenue is also used to fund the provision of public goods and other government services. A low quality diet raises the likelihood of illness and a consequent reduction in productivity. Lower productivity in turn translates into a smaller contributions to tax revenue and public goods provision. This fiscal externality is not a consequence of the TFUHC system *per se* but it does complicate the policy remedies needed to address the TFUHC-induced moral hazard problem.

We use the basic model to compare the impact of three different policy options to improve incentives for dietary health within the TFUHC system: a restricted level of treatment – relative to the first-best level – for all agents with diet-related disease; risk-based premiums for health care; and a quality-based tax on food (commonly called a “junk food tax”). All of these policies are the subject of debate in Canada and elsewhere. Proponents of reform cite the need to better match private costs with high-risk dietary behavior, but opponents raise concerns about the potentially regressive impact of introducing sharper incentives. In this paper we attempt to address both sides of this debate.

Our analysis begins with a representative agent model. This simple model allows the illumination of some key incentive issues and the characterization of the various policy options in terms of their impact on dietary choices. We then extend that model to one with agents who differ according to income, and calibrate the model with Canadian data on type 2 diabetes. The calibrated model allows an examination of how the candidate policies might compare in practice, in terms of both their incentive effects and their distributional impacts.

Our paper contributes to an extensive literature on obesity and diet-related health issues generally; see Cutler *et. al.* (2003) and Philipson and Posner (2008) for recent overviews. A number of papers in that literature have examined the impact or predicted impact of specific policies to promote dietary health but we are not aware of any work

which attempts the sort of comparative policy assessment and distributional analysis undertaken here. Variyam and Cawley (2006) examine the impact of mandated additions to the information content of nutrition labels in the US and find a significant negative impact on body weight among consumers who read the labels, and that the associated benefits far outweigh the costs of the program. Kuchler *et. al.* (2005) estimate the price-elasticity of certain salty snack foods in the US and argue that a small tax on these foods would have only a minimal impact on the consumption of those foods. In contrast, Brownwell and Hogen (2004) cite research for the US which indicates that “junk food” is in fact highly price elastic, and that a tax on this type of food could cause a large reduction in consumption.

On the other side of the calorie ledger, Cawley *et. al.* (2005) estimate the effect of a mandated increase in the time children spend in physical activity in US schools and find no impact on body fat measures for those children. Laniv *et. al.* (2009) examine food consumption and exercise in the context of a single model. They show that a junk food tax could actually increase the incidence of obesity if individuals spend less time in physical exercise in order to prepare more time-intensive healthy meals.

Bhattacharya and Sood (2006) examine the link between private health insurance contracts and dietary behavior. They argue that if overweight agents do not face risk-based premiums then these agents – who consume more treatment resources than those with healthier diets and exercise habits – impose a negative externality on other agents in the insurance pool. This is akin to the moral hazard problem under TFUHC.²

The rest of our paper is organized as follows. Section 2 presents the theoretical model. Section 3 derives the first-best solution in the representative-agent model and assesses the dietary choices under the existing TFUHC system relative to the first-best benchmark. Section 4 then examines each of the three candidate policy interventions in the context of the TFUHC system. Section 5 calibrates the model with Canadian data, and Section 6 examines the three policy options using the calibrated model. Section 7 contains some concluding remarks.

2. THE THEORETICAL MODEL

Let c_i denote the number of calories from food i consumed by a representative agent.

The total calories consumed by that agent are

$$(1) \quad C = \sum_{i \in F} c_i$$

where F is the set of all foods. Within all foods there are two types of calorie source: high quality and low quality. This dichotomous distinction is clearly an over-simplification but it captures in a tractable way the well-recognized fact that some calorie sources are healthier than others. For example, it is generally agreed within the medical profession that saturated fat and sucrose are less healthy calorie sources than protein and the complex carbohydrates found in grains, fruits and vegetables. Let ϕ_i denote the fraction of high quality calories in food type i . Then the overall fraction of high quality calories in the diet of the representative agent – her overall “diet quality” – is

$$(2) \quad q = \frac{\sum_{i \in F} \phi_i c_i}{C}$$

Let p_i denote the *price per calorie* for food type i . Then total expenditure on food by this agent is

$$(3) \quad e = \sum_{i \in F} p_i c_i$$

As noted in the introduction, foods high in fat and sugar tend to be relatively less expensive – on a per-calorie basis – than healthier foods. Here we assume the following specific relationship between price per calorie and food quality:

$$(4) \quad p_i = \rho + \theta \phi_i$$

where ρ is the price of the lowest quality food, and $\theta > 0$ is a slope parameter. Thus, foods with a higher fraction of high quality calories have a higher price per calorie. This simple linear specification for the price of calories is restrictive but it allows expenditure on food to be expressed in terms of q alone, without regard to the specific consumption

² There is also an extensive literature on moral hazard in health care more generally. See for example, Zweifel and Manning (2000), Kenkel (2000), and the earlier seminal work on moral hazard by Ehrlich and Becker (1972).

profile that underlies a given dietary quality. In particular, substitution of equation (4) into equation (3) yields

$$(5) \quad e = [\rho + q\theta]C$$

The simplicity of this reduced form linking food expenditure to diet quality for any given calorie intake allows the model to be calibrated without the need for data on consumption profiles across individual food types. Critically, note from (5) that expenditure on food is increasing in dietary quality. This is the pivotal relationship in the model.

Next consider the impact of diet on the risk of illness. In reality, there are two aspects to this: total calorie intake; and diet quality. The total number of calories is important because caloric intake relative to physical exercise levels is a key determinant of body weight. However, modeling the relationship between total calorie intake, exercise and income is beyond the scope of this paper. While there appears to be some empirical link between low income and a high calorie intake relative to exercise levels, the relationship is complicated by a host of other social factors [Philipson & Posner (2008)]. Moreover, there is no clear link between diet quality and overall calorie intake. In this model, we allow disease risk to depend on both total calorie intake and diet quality, but we assume that calorie intake is entirely independent of income and diet quality, and hence fixed for the purposes of the analysis. We focus exclusively on the link between disease and diet quality.

Let $\bar{\pi}(C)$ denote the exogenous risk of disease, linked to total calorie intake – fixed in this model – and other factors unrelated to diet quality, and let \bar{q} denote the diet quality above which there is no quality-related reduction in disease risk. That is, \bar{q} constitutes a “perfect quality” diet from a health perspective. We assume that the probability with which the agent becomes diseased is

$$(6) \quad \pi = \begin{cases} \bar{\pi}(C) + \delta(\bar{q} - q)^2 & \text{for } q \leq \bar{q} \\ \bar{\pi}(C) & \text{for } q > \bar{q} \end{cases}$$

where $\delta > 0$ is a parameter relating poor diet quality to disease risk. We have assumed a quadratic relationship in (6) to reflect the epidemiological evidence that the risk of disease seems to rise more than proportionately with body fat and related indicators of a poor quality diet [Douketis *et. al.* (2005)]. Note that we can henceforth ignore values of

$q > \bar{q}$ since quality is costly, and quality above \bar{q} yields no incremental benefits. Thus, quality above \bar{q} will never be chosen.

The effect of disease on an individual is to reduce her labor productivity and income. However, productivity can be at least partially restored through medical treatment. We model the level of medical treatment in terms of the extent to which the productivity of a diseased agent is restored to its disease-free level. In particular, let y denote the income earned by the agent if she is healthy. Then her income if she becomes ill is ym , where $m \in [0,1)$ is the degree of medical treatment she receives. This specification embodies a strong assumption that earnings fall to zero in the event of untreated illness. Depending on the illness, we might instead expect earnings to fall, but to some positive level. Our chosen specification is driven purely by the requirements of the calibration exercise in section 5, where our data on productivity losses from disease do not allow a slope and intercept parameter on the treatment-productivity relationship to be identified separately.

The cost of providing treatment level m to a diseased individual is

$$(7) \quad h(m) = \alpha \left(\frac{m}{1-m} \right)$$

where $\alpha > 0$ is a cost parameter. This specification implies a strictly convex cost of treatment, and that fully restorative treatment (where $m = 1$) is unachievable.

The government in this economy provides two types of service, funded by an income tax. The *exogenous* income tax is levied on all agents – at a flat rate t – and the revenue is used to finance publicly provided medical treatment, at a level specified by government policy, as well as the provision of public goods and other government services. Let g denote per-capita government spending on these non-treatment goods and services, which are accessed and valued by all agents equally. We will henceforth refer to g simply as per-capita spending on “the” public good. Spending on the public good is determined as a residual after treatment costs are financed from tax revenue.

The specification of utility is kept as simple as possible. In particular, utility in state k for an agent with full-health income y is

$$(8) \quad u_k = y(1-t)[k + (1-k)m] + g - e - \beta q$$

where $k = 1$ if the agent is healthy and $k = 0$ otherwise, and e denotes expenditure on food. Note that calories *per se* provide no direct utility; calories are merely a necessity needed to allow the enjoyment of other goods and services purchased with after-tax income. However, individuals may prefer certain types of food over others. In particular, foods that are high in fat and refined sugar may simply taste better than healthier foods to some people. For such people, $\beta > 0$ in (8). In reality, this parameter is likely to differ widely across individuals, and is not necessarily correlated in any way with income. For some people, β may even be negative. This heterogeneity of tastes will pose problems for the calibration of the model in section 5, and at that point we will simply set $\beta = 0$. However, it is instructive to allow for $\beta \neq 0$ in the analytical results.

Note that the level of the public good has no impact on dietary choice (due to the additive separability assumption). This simplifies the distributional analysis in section 6 significantly. If instead there is some interaction between g and dietary choice then general equilibrium feedback effects mean that an iterative process is needed to derive equilibrium outcomes in the calibrated model. This renders the determination of optimal policy much more complicated but in the end adds little of real substance to the analysis.³

Our analysis proceeds in two stages. In sections 3 and 4 we derive analytical results in the context of the representative agent model. In sections 5 and 6 we calibrate the model with Canadian data, and use the calibrated model to conduct a numerical assessment of the policy options.

3. DIETARY CHOICES UNDER TFUHC

In this section we characterize the dietary choices that arise under the existing TFUHC system, in the absence of any incentive-based policies. This provides the base case against which the policy options are compared. We begin by deriving the first-best solution as a benchmark.

³ It is worth noting that the linear manner in which g enters utility is less restrictive than it might appear. In particular, if g is instead concave in expenditure but is chosen optimally, then the analytics remain unchanged. We are grateful to an anonymous reviewer for pointing this out.

3.1 First-Best Choices

The “first-best” solution is defined here to be that which maximizes aggregate social surplus, but taking the income tax rate – and residual expenditure on the public good – as given. In an environment with homogeneous agents, this simply amounts to maximizing the utility of a representative agent. In particular, the first-best solution solves a planning problem in which diet quality and treatment level are chosen for a representative agent as follows:

$$(9) \quad \max_{q,m} \pi[y(1-t)m] + (1-\pi)y(1-t) - (\rho + q\theta)C + g - \beta q$$

$$\text{subject to } \pi = \bar{\pi} + \delta(\bar{q} - q)^2 \text{ and } g = \pi y m t + (1-\pi) y t - \pi h(m)$$

The first constraint relates dietary choice to the risk of disease. (Note that we have dropped the functional dependence of $\bar{\pi}$ on C , since C is fixed). The second constraint is a balanced-budget condition (in per-capita terms): expenditure on the public good must be equal to expected income tax revenue net of treatment costs. Note that tax revenue depends on health outcomes, via the effect of health on productivity and income. Hence, there arises the potential for a *fiscal externality* associated with private dietary choices. This will be important with respect to the implementation of first-best choices via policy intervention under the TFUHC system.

For the purposes of comparing the first-best solution with outcomes under TFUHC, it is useful to solve the planning problem in two stages. In the first stage q is chosen for a given value of m ; in the second stage m is chosen. The interior solution to the first stage is

$$(10) \quad q^*(m) = \bar{q} - \frac{(1-m)(\theta C + \beta)}{2\delta[y(1-m)^2 + \alpha m]}$$

The interior solution to the second stage is

$$(11) \quad m^* = 1 - (\alpha/y)^{1/2}$$

This solution for m^* is interior ($0 < m^* < 1$) if $\alpha < y$, meaning that the marginal cost of treatment evaluated at no treatment is less than income. This condition seems reasonable and is henceforth assumed. (Moreover, the condition is satisfied in the context of the calibrated model; see section 5.4).

Note from (11) that the optimal treatment level is an increasing function of income. In an economy with heterogeneous agents – as modeled in section 5 below – this means that high income agents should receive a higher level of treatment than low income agents, because the value of restoring the productive capacity of high income agents is higher. This notion is clearly at odds with the central tenet of a TFUHC system in which all individuals have equal access to treatment regardless of income. We examine the welfare implications of this defining feature of universal health care in section 6 below.

Substitution of (11) back into (10) yields the solution for diet quality:

$$(12) \quad q^* = \bar{q} - \frac{(\theta C + \beta)(\alpha / y)^{1/2}}{2\alpha\delta[2 - (\alpha / y)^{1/2}]}$$

This first-best dietary choice has three noteworthy properties. First, $q^* < \bar{q}$ for any finite y when $\beta \geq 0$. Some degree of unhealthy food intake – and associated elevated risk of disease – is generally optimal because unhealthy food is less costly than healthy food, and if $\beta > 0$, because unhealthy food tastes better. Second, q^* is increasing in α . This reflects the fact that a high cost of restorative treatment calls for more preventative care, in the form of a higher quality diet. Finally, q^* is increasing in y ; the cost of disease is increasing in income (because treatment does not fully restore productivity), and hence the optimal degree of preventative care via diet quality is also increasing in income.

3.2 The Base Case: Dietary Choices Under TFUHC

The TFUHC system supplies a fixed, universal level of treatment, provided free of charge, to all diseased agents, financed out of income tax revenue. The government chooses the level of treatment provided, and individual agents are prohibited from purchasing additional private insurance. This approximates the prevailing policy regime in Canada, and it is the base case to which we calibrate the model in section 5 below.

The decision problem for a representative agent under a TFUHC system that provides treatment level m is

$$(13) \quad \max_q \quad \pi m y(1-t) + (1-\pi)y(1-t) - (\rho + \theta q)C + g - \beta q$$

$$\text{subject to } \pi = \bar{\pi} + \delta(\bar{q} - q)^2$$

and the solution is

$$(14) \quad q_0(m) = \bar{q} - \frac{(\theta C + \beta)}{2\delta y(1-t)(1-m)}$$

where the “0” subscript denotes the base case (Policy 0). The essential property of this dietary choice is that diet quality is below first-best; that is, $q_0(m) < q^*(m)$.

This dietary choice under TFUHC exhibits two distortions relative to the first-best solution. The first distortion is due to tax-financing of the public good. The impact of disease on tax receipts – and hence, the provision of the public good – is external to each individual agent. Each agent therefore takes too little care to reduce the risk of disease in terms of dietary choices. The second distortion is a moral hazard problem stemming directly from the TFUHC structure: the provision of medical treatment at no marginal cost to the user erodes her incentive to make precautionary dietary choices.⁴

4. POLICY OPTIONS UNDER TFUHC

We now examine three policy options to improve dietary incentives under the TFUHC regime. The first policy option – provision of second-best treatment – takes the distorted dietary behavior as given, and attempts to modify incentives by providing less-than-first-best treatment. The other two policy options – risk-based health care premiums and a quality-based tax on food – target individual behavior directly while still providing first-best treatment. It will become clear that these latter two policy options are equivalent in the context of the representative agent model but they differ (and lose their first-best qualities) when implemented in a world with heterogeneous agents.

4.1 Policy 1: Second-Best Treatment

The dietary distortion under TFUHC arises because the individual does not face the true cost of her high-risk diet. The simplest – but imperfect – policy response to that problem is to shift more risk onto the individual by providing less than first-best treatment. Under this policy, government takes the dietary behaviour in (14) above as given, and then chooses m – the treatment level – to maximize social surplus.

To solve this problem we first need to derive expenditure on the public good as a residual when treatment level m is funded out of income tax revenue. Expected per-capita income tax revenue is

$$(15) \quad T_1(m) = \pi_1(m)myt + [1 - \pi_1(m)]yt$$

where $\pi_1(m) = \bar{\pi} + \delta[\bar{q} - q_0(m)]^2$, and where the “1” subscript denotes Policy 1. Note that the dietary behaviour from (14) is taken as given under this policy, so $\pi_1(m)$ depends on $q_0(m)$. Expected per-capita expenditure on medical treatment is

$\pi_1(m)\alpha m/(1-m)$. Thus, per-capita expenditure on the public good is

$$(16) \quad g_1(m) = \pi_1(m)[ymt - \alpha m/(1-m)] + [1 - \pi_1(m)]yt$$

The treatment level is then chosen to maximize the utility of the representative agent:

$$(17) \quad \max_m \pi_1(m)ym(1-t) + [1 - \pi_1(m)]y(1-t) - [\rho + \theta\bar{q}(m)]C + g_1(m) - \beta q_0(m)$$

Let m_1 denote the solution to (17). This solution is too cumbersome to report here but its key property is straightforward to derive (see Appendix A): $m_1 < m^*$; that is, the optimal treatment level is lower than first-best.

In the absence of any policy to influence dietary behaviour directly, the under-provision of treatment – relative to first-best – is an optimal response to the moral hazard problem and the fiscal externality. These distortions lead to lower-than-first-best quality diets, which in turn cause an excessive incidence of illness, and an associated over-utilization of the health care system. This over-utilization drives up the cost of providing treatment at any given level of treatment, and the optimal response by government is to set a treatment level lower than first-best.

While this policy does create dietary incentives in the right direction – it is clear from (14) that diet quality rises if treatment level is reduced – it comes at a significant cost, in the form of lower health restoration and lower productivity. The next two policy options studied target the dietary distortion more directly.

⁴ It is worth noting that a third type of distortion could arise here if the treatment level is not rigidly controlled. In particular, an *ex post* moral hazard problem may exist if the agent can choose how intensively to utilize the treatment level provided.

4.2 Policy 2: Risk-Based Premiums

The health care premiums we consider under this policy are not meant to mimic those that would apply under private insurance; in particular, they are not designed to be actuarially fair (where the premium paid by an agent would be set equal to the expected treatment cost for that agent). The health care system is still financed primarily through income tax revenue. The purpose of the risk-based premiums is to create better incentives for dietary health.⁵

Premiums are based on observable pre-disease physiological indicators of diet, such as an elevated cholesterol level or a high level of body fat. In reality, observable factors like obesity are imperfectly correlated with poor diet; genetic factors also play an important role. We abstract from this practical difficulty, and assume that diet quality can be inferred from observable characteristics of the individual. Thus, the health care premium can be based directly on q . In particular, the premium levied on an agent with inferred diet quality q is given by

$$(18) \quad P_2 = \gamma(\bar{q} - q)^2$$

where the “2” subscript denotes Policy 2, and γ is a parameter to be chosen as part of the optimal policy. The decision problem for an agent facing this premium rule is

$$(19) \quad \max_q \quad \pi y m(1-t) + (1-\pi)y(1-t) - (\rho + \theta q)C - P_2 + g - \beta q$$

$$\text{subject to } \pi = \bar{\pi} + \delta(\bar{q} - q)^2 \text{ and } P_2 = \gamma(\bar{q} - q)^2$$

and the solution is

$$(20) \quad q_2(m) = \bar{q} - \frac{\theta C + \beta}{2[\delta y(1-t)(1-m) + \gamma]}$$

Note that $q_2(m)$ is increasing in γ . That is, the risk-based premium creates an incentive for the agent to adopt a more precautionary diet.

Now consider the policy-maker’s choice of γ and m . Expected per-capita revenue – from income taxation and health care premiums – is

⁵ The incentive effects of this risk-based health care premium could also be achieved with a risk-based user fee for treatment. These *ex ante* and *ex post* instruments are equivalent in our model because the linear utility function exhibits risk-neutrality, and because the treatment level is fixed.

$$(21) \quad R_2(m) = \pi_2(m)ymt + [1 - \pi_2(m)]yt + P_2(m)$$

where $\pi_2(m) = \bar{\pi} + \delta[\bar{q} - q_2(m)]^2$ and $P_2(m) = \gamma[\bar{q} - q_2(m)]^2$. Expected per-capita expenditure on treatment is $\pi_2(m)\alpha m/(1-m)$. Thus, per-capita expenditure on the public good (determined as a residual) is

$$(22) \quad g_2(m) = \pi_2(m)[ymt - \alpha m/(1-m)] + [1 - \pi_2(m)]yt + P_2(m)$$

The policy problem for government is to choose m and γ to maximize the utility of the representative agent, based on her response to the policy from (20) :

$$(23) \quad \max_{m,\gamma} \pi_2(m)ym(1-t) + [1 - \pi_2(m)]y(1-t) - [\rho + \theta q_2(m)]C \\ - P_2(m) + g_2(m) - \beta q_2(m)$$

First consider the solution for m :

$$(24) \quad m_2 = 1 - (\alpha/y)^{1/2}$$

This is equal to the first-best treatment level. That is, the introduction of risk-based health care premiums under TFUHC allows the level of treatment to be set at its first-best level. This is because the over-utilization of the health care system associated with poor dietary choices under TFUHC is corrected when the agent faces a risk-based premium. Note that this premium must be set to address both the moral hazard problem and the fiscal externality. To see this, consider the solution to (23) for γ :

$$(25) \quad \gamma_2 = \delta[m_2 + t](\alpha y)^{1/2}$$

There are two parts to this expression, corresponding to the two additive terms. The first part targets the moral hazard problem. This part is proportional to the level of treatment provided (m_2) because the incentive to adopt a precautionary diet falls as the level of restorative treatment rises. Thus, the health care premium required to correct the weakened precautionary incentive must rise with the treatment level. The second part of the health care premium targets the externality associated with tax-financing of public good provision. This fiscal externality is proportional to the income tax rate, and so too therefore is the required corrective premium.

Evaluating $q_2(m)$ from (20) at $\gamma = \gamma_2$ and $m = m_2$ yields $q_2 = q^*$. That is, the outcome for diet quality under the optimal policy is first-best. The risk-based health care

premium induces the agent to take full account of the risk associated with her dietary choices.

An alternative to risk-base premiums is to influence dietary choices by changing the relative price of high and low quality food, via a quality-based tax. We consider that policy approach next.

4.3 Policy 3: A Quality-Based Tax on Food

The tax is levied as an excise tax per calorie, based on food quality. In particular, the excise tax per calorie on food type i is $\tau(1-\phi_i)$, where τ is the policy parameter, and ϕ_i is the fraction of high quality calories in food type i . Note that only “perfect quality” food – for which $\phi_i = 1$ – attracts a tax rate of zero. The tax on food means that expenditure on food for an agent with dietary quality q is now given by

$$(26) \quad e(\tau) = (\rho + q\theta + (1-q)\tau)C$$

Thus, the tax effectively flattens the relationship between food prices and food quality, thereby making low quality food less attractive. Under this policy, the choice problem for the representative agent is

$$(27) \quad \max_q \quad \pi y m(1-t) + (1-\pi)y(1-t) - (\rho + \theta q + (1-q)\tau)C + g - \beta q$$

subject to $\pi = \bar{\pi} + \delta(\bar{q} - q)^2$

and the solution is

$$(28) \quad q_3(m) = \bar{q} - \frac{(\theta - \tau)C + \beta}{2\delta y(1-t)(1-m)}$$

where the “3” subscript denotes Policy 3.

Now consider the policy-maker’s choice of τ and m . Expected per capita revenue – from income taxation and the tax on food – is

$$(29) \quad R_3(m) = \pi_3(m)y m t + [1 - \pi_3(m)]y t + \tau[1 - q_3(m)]C$$

where $\pi_3(m) = \bar{\pi} + \delta[\bar{q} - q_3(m)]^2$. Expected per capita expenditure on medical treatment is $\pi_3(m)\alpha m/(1-m)$. Thus, per capita expenditure on the public good (determined as a residual) is

$$(30) \quad g_3(m) = \pi_3(m)[y m t - \alpha m/(1-m)] + [1 - \pi_3(m)]y t + \tau[1 - q_3(m)]C$$

The policy problem for government is to choose m and τ to maximize the utility of the representative agent, based on her response to the policy from (28) :

$$(31) \quad \max_{m, \tau} \pi_3(m) y m (1-t) + [1 - \pi_3(m)] y (1-t) \\ - (\rho - (\theta - \tau) [1 - q_3(m)]) C + g_3(m) - \beta q_3(m)$$

First consider the solution to (31) for m :

$$(32) \quad m_3 = 1 - (\alpha / y)^{1/2}$$

This is equal to the first-best treatment level. Thus, the treatment level under the optimal food tax policy is the same as that provided under an optimal policy with risk-based health care premiums, and in both cases it is first-best. Now consider the solution to (31) for τ :

$$(33) \quad \tau_3 = \frac{[m_3 + t] \left(\theta + \frac{\beta}{C} \right)}{2 - (\alpha / y)^{1/2}}$$

This tax parameter has two parts, corresponding to the two terms in the square brackets of the numerator. The first part targets the moral hazard problem; it is proportional to the treatment level. The second part targets the fiscal externality; it is proportional to the income tax rate. The tax is able to correct both of these distortions, and is therefore able to implement the first-best dietary choice. In particular, evaluating $q_3(m)$ at $\tau = \tau_3$ and $m = m_3$ yields $q_3 = q^*$. Note that the food tax parameter is increasing in β ; a higher tax rate is needed to dissuade the agent from consuming low quality foods if she has a strong taste for such food.

4.4 Limitations of the Representative Agent Model

The results for the representative agent model help to illuminate some key incentive issues but the analysis to this point suffers from two major limitations in terms of guiding policy design in practice. First, the capacity to achieve first-best outcomes through corrective policy under TFUHC relies directly on the representative agent assumption. Recall from (11) that the first-best treatment level is increasing in income. In a setting where individuals differ according to income, a TFUHC system cannot achieve the first-best solution, by virtue of the fact that the treatment level is universal.

Note too that the optimal policy interventions must also be individual- specific. Consider for example the quality-based tax on food. The optimal tax rate under TFUHC – described in expression (33) above – is a function of income, and the taste parameter. This is not a problem in a representative agent economy, but it is a significant problem in an economy with heterogeneous agents. In such an economy, the tax on food would need to be tailored specifically to each individual agent. This is simply not practical. A similar shortcoming is true of risk-based health care premiums (though in that case, taste heterogeneity is not a problem); see expression (25). If government is constrained – for practical considerations – to use one-size-fits-all policy parameters, then the different interventions we have examined under the TFUHC regime are not necessarily equivalent in an economy with heterogeneous agents, and they will not be able to implement the first-best solution. Both the food tax and risk-based premiums will yield only second-best results, and one may outperform the other.

The second key limitation of the representative agent analysis is the absence of distributional considerations. Concerns over distribution are typically at least as important as efficiency issues from the perspective of policy-makers in practice, and this is likely to be especially true with respect to health care policy in countries with TFUHC systems, where a primary motivation behind these systems is a distributional one. There are two possible approaches to these distributional issues. One approach would be to integrate the distributional goals of government directly into the analytical model by allowing heterogeneity across incomes. This approach has two limitations. First, allowing for multiple income types renders the analysis much less intractable. Second, this approach requires the specification of an explicit distributional goal. Instead, our approach in sections 5 and 6 following is to retain the focus on surplus maximization as the criterion for the policy design problem, and present the distributional implications of the various policy options in the context of a calibrated model using actual income distribution data. Any tradeoff between social surplus and distributional goals can then be judged by policy-makers on the basis of those results.

5. CALIBRATING THE MODEL WITH CANADIAN DATA

We approach the calibration of the model in two parts. The first part involves constructing an analytical description of the distribution of Canadian household income. The second part involves fitting the TFUHC base case from section 3.2 to Canadian data on diet quality and diet-related disease. The disease on which we were able to obtain the richest data is type 2 diabetes, and so we focus exclusively on that disease. Type 2 diabetes is one of the most important diet-related diseases, and about 90% of all diabetes cases in Canada are of type 2 [Health Canada (1999)]. Our measure of diet quality is fat content. In particular, we interpret $(1 - q)$ as the percentage of calories from fat in the diet, or the “fat intensity” of the diet.

5.1 Canadian Household Income

We use household income data rather than individual income data because it seems more reasonable to examine dietary choice as a household-level decision. The data we have on type 2 diabetes and related variables is for years in the period 1998 – 2001, so we work with income distribution data for 2000.⁶ The mean income is \$66,160 and the median is \$55,016.

The data on income is decile data. In principle, this data could be used directly for calibrating the model but this necessitates the use of numerical integration methods in the simultaneous solution of equations (34) – (38) below. We were not able to obtain a convergent solution via this method. Instead, we fitted a lognormal distribution to the income data, over a truncated income range of \$5,000 to \$10m, making the appropriate scaling adjustment to ensure that we retained a legitimate density. The parameters of the calibrated lognormal distribution are $\mu = 4.00758$ and $\sigma = 0.6074$. The mean of the calibrated distribution – henceforth denoted $f(y)$ – is \$66,160 and the median is \$55,015.

The goodness of fit to the actual data can be judged by the Kolmogorov-Smirnov test. The test statistic is calculated as the greatest discrepancy between the actual cumulative density and the null cumulative density (in this case the lognormal) across the

⁶ The data source is *Statistics Canada, Income of Canadian Families, 2001 Census, Catalogue No. 96F0030XIE2001014*.

income classes. The observed value of the Kolmogorov-Smirnov test statistic is 0.0602. This implies an extremely good fit to the data; the lognormal hypothesis cannot be rejected at the 20% level.

5.2 Type 2 Diabetes and Dietary Fat Intensity

The calibration requires data on treatment costs and productivity losses. The only available Canadian study on diet-related disease that reports both of these values is Dawson *et. al.* (2002). Their study focuses on diabetes. Dawson *et. al.* estimate that productivity losses in Canada in 1998 due to premature death from diabetes and its chronic complications were \$1.593b. They do not separate out losses from the two types of diabetes so we assigned 90% of that cost to type 2 diabetes, based on the relative incidence of type 1 and type 2 in the population. The Dawson *et. al.* estimate does not include losses due to morbidity, and these are likely to be at least as great as those due to premature death.⁷ Based on various pieces of mostly anecdotal evidence on diabetes patient survivorship, we assume that productivity losses due to morbidity are 125% of those due to premature death. This implies an overall annual productivity loss from type 2 diabetes of \$3.226b.

Productivity loss in the context of the model is the aggregation of expected income loss across households. There were $N_2 = 8,371,020$ households in Canada in 1998. Thus, for the base case calibration we require

$$(34) \quad 8.37102 \int_{0.005}^{10} \pi_0(m, y) y(1-m) f(y) dy = 3226$$

where $\pi_0(m, y) = \bar{\pi} + \delta[\bar{q} - q_0(m, y)]^2$, and $q_0(m, y)$ is given by expression (14) in section 3.2, with the functional dependence on y now made explicit.

Our estimate for treatment cost is also drawn from Dawson *et. al.* (2002). Their estimate for 1998, based on all diagnosed diabetes cases, is \$3.896b. We assigned 90% of that cost to type 2 diabetes, based on its incidence relative to type 1. These costs were

⁷ A recent U.S. study – Flegal *et. al.* (2005) – finds that premature deaths are likely not the most important consequence of diet-related disease because medical treatment (at huge cost) keeps the diseased patients alive. A recent study from the Heart and Stroke Foundation of Canada (2006) reports a similar finding with respect to trends in treatment and premature death rates for cardiovascular disease in Canada.

spread across $N_1 = 22,241,050$ adults aged 20 years and older in 1998. This implies the following relationship in the context of the model:

$$(35) \quad 22.241050 \int_{0.005}^{10} \pi_0(m, y) h(m) f(y) dy = 3506$$

Our estimates for the incidence of type 2 diabetes in the adult population, and the fraction of that incidence due to poor diet, are derived from data from Health Canada (1999) and World Health Organization (2002); see Appendix B for the details. The implied values for the calibration are $\bar{\pi} = 0.01312$ and

$$(36) \quad \int_{0.005}^{10} \pi_0(m, y) f(y) dy = 0.0369$$

The average dietary fat intensity in Canada in 1998 is believed to have been about 38%. [Health Canada (1999)]. In the context of the model, this means

$$(37) \quad \int_{0.005}^{10} [1 - q_0(m, y)] f(y) dy = 0.38$$

Average annual household expenditure on food in Canada in 1998 was \$5,908.⁸ Thus, we require

$$(38) \quad \int_{0.005}^{10} e_0(m, y) f(y) dy = 0.005908$$

where $e_0(m, y)$ is given by (5) evaluated at $q = q_0(m, y)$.

The final four parameters we specify are household caloric intake, the “risk-free” dietary fat intensity, the Canadian income tax rate, and the taste parameter, β . The average Canadian household in 1998 comprised one adult male, one adult female and 1.6 children. We assume an intake of 2200 calories per day for the adult male, 1800 for the adult female and 1250 for each child. Thus, the average household consumes 6000 calories per day over 365 days in a year, yielding $C = 2.19$ million in the context of the model. The “risk-free” dietary fat intensity relates to the value of \bar{q} in equation (6); it is the dietary fat intensity below which diet does not contribute to the risk of disease. For the calibration we use the fat intensity value recommended by Canadian health authorities

⁸ The data source is *Statistics Canada, Survey of Household Spending, Household Spending on Food, CANSIM II, SERIES V13873702*.

(as indicated on mandated nutrition labels): $1 - \bar{q} = 0.27$. The income tax rate we assume for the calibration is the average Canadian income tax rate in 2000: $t = 0.25$. (Note that we abstract from the nominal progressivity of the Canadian income tax system). As noted earlier, the taste parameter is likely to differ widely across individuals, and there is simply not enough data available to construct a distribution for it. We therefore set $\beta = 0$, but we discuss the key implications of doing so in section 7.

We used *Maple (Version 11)* to solve equations (34) – (38) simultaneously for $\{\rho, \theta, \delta, \alpha, m\}$. The fitted values are reported in Table 1 along with values for the specified parameters. Note that we have allowed m – the universal treatment level – to be fitted freely; we have *not* imposed the assumption that m is currently chosen optimally (as described in Policy 1) under the existing TFUHC system. The calibrated treatment level is 0.768.

It is important to emphasize that our model is a highly abstract one and that we have calibrated it for just one diet-related disease. We view the numerical results as indicative of the likely comparative performance of the policy alternatives, and as useful in illustrating how policies that are equivalent in a narrow theoretical context can in fact be quite different in terms of their performance in practice, especially with respect to their distributional impacts. Our results should *not* be interpreted as reliable estimates of the absolute net benefits of the policies we examine.

5.3 THE CALIBRATED BASE CASE

The calibrated TFUHC regime from section 3.2 is the base case against which the alternative policy interventions are assessed. The properties of this base case are summarized in Tables 2 – 4 (Policy 0). Note from Table 3 that lower income households are predicted to have a higher dietary fat intensity, reflecting the fact that foods higher in fat are less expensive. As noted earlier, this propensity for lower income households to consume more fat is consistent with observed dietary profiles in practice. To put the predicted fat intensities from Table 3 in perspective, it is worth noting that a typical meal from a well-known North American family restaurant, comprising a cheeseburger, a large serving of fries and a cup of ice cream, currently contains about 1430 calories, of which about 46% are from fat.

5.4 The First-Best Solution

To simulate the first-best solution we set m and q for a household with income y according to equations (11) and (12) respectively, using the calibrated parameter values. The properties of this first-best solution are summarized in Tables 2 – 4 under “FB”.⁹ Several points emerge from these numbers. First, dietary fat intensity is significantly lower – by almost 19% on average – for all households under the first-best solution relative to the base case. The difference is greatest for lower income households.

Second, the first-best treatment level differs considerably across income groups. The median household in the lowest income quartile receives only 78% of the treatment received by the median household in the highest income quartile. Moreover, treatment should be lower than the base case level for the two lowest income quartiles, and higher than the base case level for the two highest income quartiles.

Third, productivity losses and treatment costs are dramatically lower than in the base case. The reduced costs stem from the lower disease incidence – especially for lower income households – and additionally, in the case of treatment costs, from lower treatment levels for the two lowest income quartiles. Note however that total treatment costs are 61% lower, while the average treatment level for treated patients is only 1.4% lower. The primary source of this dramatic decline in treatment cost is the reduction in the number of patients needing treatment, due to the lower disease incidence induced by better dietary choices. The cost against which these benefits must be weighed is the increase in food expenditures for all households. Average household food expenditures rise by \$175 (or around 3%) relative to the base case. Nonetheless, the overall difference in social surplus between the base case and the first-best solution is \$2056m.

This foregone social surplus under the current TFUHC regime comprises two parts: a deadweight loss due to the distortion of dietary incentives; and a deadweight loss due to the universality of treatment. The first of these two parts can be calculated by simulating the base case dietary choices for each household – as given by (14) – in response to the first-best treatment level for that household. The associated aggregate treatment costs, aggregate productivity losses and aggregate food expenditures can then

be calculated from those choices and compared to their first-best values. The implied loss of social surplus is \$1962m. The residual deadweight loss is \$94m; this is due to the universality of treatment.

The deadweight loss due to dietary distortions can be further decomposed into a component due to the moral hazard problem, and a component due to the fiscal externality. The latter can be calculated by simulating the first-best dietary choice for a given household – as given by (10) – but with α (the treatment cost parameter) set equal to zero. Setting $\alpha = 0$ in this way constructs a first-best dietary choice as if treatment cost did not matter, in which case there arises no moral hazard problem. The difference between this artificially constructed first-best choice and the base case dietary choice for each household – as given by (14) – in response to the first-best treatment level for that household is attributable to the fiscal externality. Making this calculation reveals that \$1020m (or 52%) of the deadweight loss from distorted dietary incentives is due to the fiscal externality. The residual \$942m is due to the moral hazard problem.

6. A COMPARISON OF POLICY OPTIONS

We now turn to policy options for addressing the dietary distortions. We examine each of the three options from section 4 in the context of the calibrated model. Each policy is assessed relative to the base case TFUHC. For each policy examined we hold the income tax rate constant and assume that any changes to net government revenues are redistributed equally across households through changes to expenditure on the public good (to the equal benefit of all households). It is important to note that if revenue changes were instead offset by changes in the income tax rate, or if expenditure on the public good benefited some households more than others, then the distributional impacts could be quite different from those we derive.

⁹ Recall from section 3.1 that the first-best solution is interior if and only if $\alpha < y$. The calibrated value of α is 1.2996 (measured in thousands of dollars). Thus, the first-best solution is interior for any household with income greater than \$1300. This holds for all Canadian households.

6.1 Policy 1: Second-Best Treatment

Recall from section 4.1 that the second-best treatment level is chosen in response to the distorted dietary behavior, as given by (14) from section 3.2. This second-best treatment level is a function of income. In an economy in which households differ with respect to income, a *universal* level of treatment cannot be optimal for all households. Thus, the treatment level we derive here should be properly thought of as “third-best”, taking as given the dietary behavior reflected in (14), *and* constrained by the universality requirement. We derive this (third-best) optimal treatment level as that which maximizes aggregate social surplus. In particular, we construct the compensating variation (CV) – relative to the base case – for a household with income y , and then integrate over the income distribution to derive the aggregate net benefit as a function of m . We then maximize that net benefit with respect to m to find the optimum.

The optimal solution is $m_1 = 0.65$. Note that this is significantly lower than the base case treatment level of 0.768. This implies that the level of treatment currently offered in Canada (for type 2 diabetes) is excessively high, given the TFUHC system and the current absence of any incentive-based policies to correct dietary distortions.

The values of key variables under this policy are reported in Tables 2 – 4 (Policy 1). Three key points emerge from these numbers. First, dietary fat intensity and disease incidence are lower than in the base case, and the difference is most significant for lower income households. This difference in impact across households reflects the fact that lower income households have lower quality diets in the base case, and are therefore exposed to higher risk when the treatment level is reduced. Second, treatment costs are reduced substantially under this policy but productivity losses are actually *higher* than in the base case. This large potential loss of income – due to lower treatment – is what motivates households to reduce their dietary fat intake. Third, the overall net benefit from the reduced treatment policy (\$1142m) is substantial despite its bluntness, but nonetheless captures only 56% of the additional surplus potentially available (measured against the first-best benchmark).

How is this net benefit distributed across households? The distributional impacts of the policy are calculated as the CV for each household. This is reported by income quartile in Table 5 (Policy 1). The top half of the table identifies winners (W) and losers

(L) according to income level. Households with incomes lower than \$213,000 – who constitute 98.7% of all households – are winners; they are better off relative to the base case. These households gain \$1148m in aggregate. Households with incomes higher than \$213,000 – who constitute 1.3% of all households – are losers. In aggregate, this group loses approximately \$6.8m.

The lower half of Table 5 summarizes distributional impacts by income quartile. The first line in the lower half of the table lists the income quartiles (1 – 4). The next line reports the CV for the median household within the quartile. For example, the median household in the first quartile gains \$151 under Policy 1 relative to the base case. The next line reports the aggregate gain (or loss) for each quartile. For example, the first quartile as a whole gains \$299m under Policy 1. The last line of the table reports the fraction of the overall net benefit captured by each quartile. For example, the lowest quartile captures 26.2% of the overall net benefit from Policy 1.

The key message from these distributional impacts is that lower-middle income households benefit more than higher income households. This reflects the fact that a lower treatment level is more costly for higher income households than for lower income ones, because the former have more to lose from disease. On the other hand, the very lowest income households have a higher risk of disease than middle income households, so they do less well than middle income households when the treatment level is reduced. The very wealthiest households – those with incomes greater than \$213,000 – are actually made worse off by this policy (though the upper income quartile nonetheless derives modest gains in aggregate).

6.2 Policy 2: Risk-Based Premiums

We derive the optimal premium rule – from the class of rules specified in equation (18) – by maximizing the aggregate net benefit with respect to m and γ jointly, where γ is constrained to be the same across households. Thus, we have imposed a universality requirement on both the treatment level and the premium rule. The optimal solutions are $m_2 = 0.767$ and $\gamma_2 = 7.317$. The implied premiums cover only 15.6% of total treatment costs; treatment is still funded primarily out of tax revenue.

The implications of this policy for the key variables are reported in Tables 2 – 4 (Policy 2). There are two key points to note about these results. First, the optimal treatment level under this policy is significantly higher than under Policy 1 (second-best treatment level), and remarkably close to the base case treatment level (0.767 vs. 0.768). This means that the current level of treatment in Canada – as implied by the calibration – is almost exactly compatible with this policy.

Second, the reduction in dietary fat intensity induced by this policy – relative to the base case – is substantial. Note in particular from Table 4 that dietary fat intensity for the median household in each of the two lowest income quartiles is reduced to a level *below* first-best, while fat intensities for the higher income households are reduced to levels that are only marginally higher than first-best. The overall net benefit of the policy is \$1879m, which means that it captures around 91% of the additional surplus potentially available (measured against the first-best benchmark) despite operating within the confines of universal treatment.

The distributional impacts of risk-based premiums – measured against the base case – are summarized in Table 5 (Policy 2). Two points are noteworthy here. First, the median household in every income quartile is made better off by this policy. Second, while almost 98% of households benefit, households with incomes below around \$16,450 are made worse off. Thus, the policy does not Pareto-dominate the base case.

6.3 Policy 3: A Quality-Based Tax on Fat

We derive the optimal tax by maximizing the aggregate net benefit with respect to m and τ jointly. The optimal solutions are $m_3 = 0.763$ and $\tau_3 = 7.7586 \cdot 10^{-7}$, and the effective tax rate is

$$(39) \quad r_3(\phi) \equiv \frac{\phi \tau_3}{p(\phi)} = \frac{0.77586\phi}{3.1247 - 1.1237\phi}$$

For illustrative purposes, the effective tax rate (ETR) and after-tax price per 1000 calories (ATP) are reported in Table 6 for a selection of food items, based on a survey of food prices at local supermarkets. The average tax rate – the percentage tax payable on average food expenditure – is 7.46%.

The key summary variables under this policy are reported in Tables 2 – 4 (Policy 3). The tax policy induces an outcome broadly similar to that under risk-based health-care premiums (Policy 2) but is slightly inferior to the latter in three respects. First, fat intensities under the food tax are marginally further from first-best for all income quartiles except the lowest quartile. Second, a higher fat intensity for the lowest quartile – and an associated higher utilization of treatment – means that the optimal treatment level under the food tax policy is slightly lower than under risk-based premiums. Third, the food tax yields a lower overall net benefit than the risk-based premiums policy.

Recall that the food tax and risk-based premiums perform equally well in the representative agent setting. Why do they differ here? The difference stems from the fact that the food tax is not as sharp an instrument for addressing the moral hazard problem when households differ across income. The chain of causation between diet and treatment comprises two links: diet to risk, and risk to expected treatment cost. A policy of risk-based premiums targets the second link directly, while the food tax targets the second link only indirectly, via its effect on the first. This is inconsequential in the representative agent setting, because the food tax can be set to mimic the risk-based premiums exactly. That exact correspondence breaks down when households differ with respect to income, and the policy-maker is constrained to choose policy parameters that cannot differ across households. In that case, it matters that the food tax targets a relationship that is one link removed from the crux of the moral hazard problem, namely, the relationship between risk and expected treatment costs. Note that this difference in the performance of the two policies would be even greater if tastes are heterogeneous across households, because the optimal tax rate depends on the taste parameter, while the optimal premium rule does not.

The distributional impacts of the food tax are summarized in Table 7. These impacts are broadly similar to those under risk-based premiums (Policy 2) though the median household in every quartile is marginally worse off under the food tax than under risk-based premiums. However, the policies cannot be Pareto-ranked: the very poorest households are slightly better off under the food tax than under risk-based premiums. Nor is it possible to Pareto-rank the food tax policy and the base case, even though almost 98% of households are better off under the former. In particular, the very poorest

households are made worse off, albeit by only \$11.81m in aggregate (compared to an overall net benefit of \$1845m from the policy).

7. CONCLUSION

Our results suggest that policy intervention to improve incentives for a healthy diet in Canada may have substantial net benefits. While the absolute numbers derived here should be viewed with caution, they point to the potential for reducing health care costs, and improving health outcomes at the same time, by addressing the moral hazard problem associated with TFUHC. Although first-best outcomes cannot be achieved under *any* universal treatment regime in a setting with heterogeneous households, it appears that there is considerable room for improvement over the *status quo*. In particular, risk-based health care premiums and a quality-based tax on food are policies worthy of further consideration. While the gains would tend to favor higher income households, the vast majority of households appear likely to benefit. Moreover, the losses to the poorest households are small enough to be easily addressed via compensating measures.

As promising as these results may seem, it is important to emphasize that they are derived in the context of a simple model that abstracts from many important practical issues. First, we have not considered implementation issues. A quality-based tax on food would be best applied upstream in the food production chain where unhealthy calorie sources like saturated fat and refined sugar would be taxed as ingredients. This would avoid the administrative difficulty of taxing different downstream foods at different rates according to the quality of their calories. However, the tax policy would nonetheless have substantial administration costs.

Risk-based health care premiums would similarly face implementation problems. Chief among them would be the issue of observable risk. Pre-disease physiological indicators – including such highly visible indicators as obesity – are closely correlated with disease risk but that correlation is far from perfect. Moreover, society may not be comfortable with discrimination based on physical characteristics, even when those characteristics can be tied to risky behavior. This might be especially true when the risky behavior is correlated with low household income.

We have also abstracted from many individual-specific factors that are likely to affect dietary choices. For example, recall that tastes for different types of food were excluded from the calibrated model. If low quality foods – high in saturated fat and refined sugar – are indeed preferred by many people, then a failure to account for this could mean that the welfare benefits of improved dietary quality are over-stated. It would also mean that the response to policy intervention in practice could be smaller than the numerical analysis here suggests. This is can be seen most clearly in the context of the tax of food. Recall from (33) that the optimal tax rate in the representative agent model depends on the taste parameter β . By ignoring β in the numerical analysis, the optimal tax rate calculated is too low for an individual with $\beta > 0$, and would therefore induce a smaller dietary response than predicted. The converse is true if $\beta < 0$. Moreover, if β differs significantly across individuals then a one-size-fits-all tax rate could have more significant distributional impacts than predicted here. This would be especially true if taste for unhealthy food and income are somehow correlated.

Despite the absence of these considerations from our model, three useful lessons can be drawn from the analysis. First, the moral hazard distortion under TFUHC is potentially a key contributor to poor dietary habits, disease incidence and high treatment costs. Second, policy intervention to improve dietary incentives is worthy of serious consideration based on the likely net benefits. Third, the likely distribution of net benefits from policy intervention – while favoring high income households to some degree – may not be nearly as regressive as might be expected on the basis of current dietary habits across income groups.

APPENDIX A: POLICY 1

Let u_1 denote the maximand in (17). Differentiate u_1 with respect to m and evaluate the derivative at $m = m^*$ to obtain

$$(A1) \quad \frac{\partial u_1}{\partial m} = -\frac{(t + m^*)(\theta C + \beta)^2}{2\alpha\delta(1-t)^2}$$

This is negative. Since $\partial u_1 / \partial m = 0$ at $m = m_1$, and since u_1 is strictly concave in m , it follows that $m_1 < m^*$.

APPENDIX B: DIABETES INCIDENCE AND EXOGENOUS RISK

Let z_i denote the incidence of type i diabetes in the adult population in 1998. According to Health Canada (1999), $z_1 + z_2 = 0.041$ and $z_2 = 9z_1$. It follows that $z_1 = 0.0041$ and $z_2 = 0.0369$. This explains equation (36). Now let λ_i denote the fraction of type i cases that are due to excessive body fat, and let λ denote the fraction of *all* diabetes cases that are due to excessive body fat. It follows that $0.0041\lambda_1 + 0.0369\lambda_2 = 0.041\lambda$. According to World Health Organization (2002), $\lambda_1 = 0$ and $\lambda = 0.58$. Thus, $\lambda_2 = 0.6444$. Then the incidence of type 2 diabetes that is *not* due to excessive body fat is $(1 - \lambda)z_2 = 0.01312$. This is our estimate for π_0 .

REFERENCES

- Bhattacharya, J. and N. Sood (2006), "Health Insurance and the Obesity Externality", in K. Bolin and J. Cawley (eds.), *The Economics of Obesity*, Emerald: New York.
- Birmingham, C.L., J.L. Muller, A. Palepu, J. Spinelli and A.H. Anis (1999), "The Cost of Obesity in Canada", *Canadian Medical Association Journal*, 160(4), 483-488.
- Brownwell, K.D. and K.B. Hogen (2004), *Food Fight: The Inside Story of the Food Industry, America's Obesity Crisis, and What we Can Do About It*, Contemporary Books: New York.
- Canadian Population Health Initiative (2006), *Improving the Health of Canadians: Promoting Healthy Weights*, Canadian Institute for Health Information: Ottawa.
- Cawley, J., C.D. Meyerhoefer and D. Newhouse (2005), *The Impact of State Physical Education Requirements on Youth Physical Activity and Overweight*, National Bureau of Economic Research Working Paper 11411.
- Chou, S., M. Grossman and H. Saffer (2004), "An Economic Analysis of Adult Obesity: Results from the Behavioral Risk Factor Surveillance System", *Journal of Health Economics*, 23(3), 565-87.
- Critser, G. (2003), *Fat Land: How Americans Became the Fattest People in the World*, Mariner Books: New York.
- Cutler, D.M., E.L. Glaeser and J.M. Shapiro (2003), "Why Have Americans Become More Obese?", *Journal of Economic Perspectives*, 17, 93-118.
- Dawson, K.G., D. Gomes, H. Gerstein, J.F. Blanchard and K.H. Kahler (2002), "The Economic Cost of Diabetes in Canada, 1998", *Diabetes Care*, 25(8), 1303-1307.
- Drewnowski, A. and S. E. Specter (2004), "Poverty and Obesity: The Role of Energy Density and Energy Costs", *American Journal of Clinical Nutrition*, 79(1), 6-16.
- Douketis, J.D., G. Paradis, H. Keller and C. Martineau (2005), "Canadian Guidelines for Body Weight Classification in Adults: Application in Clinical Practice to Screen for Overweight and Obesity and to Assess Disease Risk", *Canadian Medical Association Journal*, 172, 995-998.
- Ehrlich, I. and G.S. Becker (1972), "Market Insurance, Self-Insurance, and Self Protection", *Journal of Political Economy*, 80, 623-48.

- Flegal, K.M., B.I. Graubard, D.F. Williamson and M.H. Gail (2005), “Excess Deaths Associated With Underweight, Overweight, and Obesity”, *Journal of the American Medical Association*, 293, 1861-1867.
- Health Canada (1999), *Diabetes in Canada: National Statistics and Opportunities for Improved Surveillance, Prevention, and Control*, Public Works and Government Services Canada: Ottawa.
- Heart and Stroke Foundation of Canada (2006), *Tipping the Scales of Progress: Heart Disease and Stroke in Canada 2006*, Heart and Stroke Foundation of Canada: Ottawa.
- Kenkel, D.S. (2000), “Prevention”, in: A.J. Culyer and J.P. Newhouse (eds.), *Handbook of Health Economics (Vol. 1)*, Elsevier: Amsterdam.
- Kirkpatrick, S. and V. Tarasuk (2003), “The Relationship Between Low Income and Household Food Expenditure Patterns in Canada”, *Public Health Nutrition*, 6(6), 589-597.
- Kuchler, F., A. Tegene and J.M. Harris (2005), “Taxing Snack Foods: Manipulating Diet Quality or Financing Information Programs?”, *Review of Agricultural Economics*, 27(1), 4-20.
- Le Petit, C., and J.M. Berthelot (2005), *Obesity: A Growing Issue*, Statistics Canada: Ottawa.
- Philipson, T. and R. Posner (2008), *Is the Obesity Epidemic a Public Health Problem? A Decade of Research on the Economics of Obesity*, NBER Working Paper 14010.
- Variyam, J. N. and J. Cawley (2006), “Nutrition Labels and Obesity”, National Bureau of Economic Research, Working Paper 11956.
- World Health Organization (2002), *The World Health Report, 2002: Reducing Risks, Promoting Healthy Life*, World Health Organization: Geneva.
- Yaniv, G, O. Rosin and Y. Tobol (2009), “Junk-Food, Home Cooking, Physical Activity and Obesity: The Effect of the Fat Tax and the Thin Subsidy”, *Journal of Public Economics*, 93, 823-30.
- Zweifel, P. and W.G. Manning (2000), “Moral Hazard and Consumer Incentives in Health Care”, in A.J. Culyer and J.P. Newhouse (eds.), *Handbook of Health Economics (Vol. 1)*, Elsevier: Amsterdam.

TABLES

Specified Parameter	Value	Fitted Parameter	Value
t	0.25	ρ	3.1247 x 10^{-6}
$\bar{\pi}$	0.01312	θ	1.1237 x 10^{-6}
C	2.19 million	δ	1.3940
$1-\bar{q}$	0.27	α	1.2996
β	0	m	0.7684

TABLE 1: Specified and Fitted Parameter Values

Policy	0	1	2	3	FB
Av. dietary fat intensity (%)	38.0	34.3	30.7	30.3	30.9
Av. Type 2 incidence (%)	3.69	2.36	1.53	1.54	1.57
Av. treatment level (%)	76.8	65.0	76.7	76.3	73.8
Average food tax rate (%)				7.46	
Av. food expenditure (\$)	5908	5999	6094	6620	6083
Treatment costs (\$m)	3506	1269	1451	1433	1383
Productivity losses (\$m)	3226	3563	1904	1880	1838
Income tax revenue (\$b)	137.3	137.2	137.8	137.6	137.6
Fax tax revenue (\$m)	1	2	7	4	6
Net benefit (\$m)		1142	1879	1845	2056

TABLE 2: Key Summary Variables by Policy

Dietary Fat Intensity for the Median Household within the Quartile (%)	Q1	Q2	Q3	Q4
Policy 0	47.6	38.3	34.6	31.4
Policy 1	40.6	34.5	32.1	29.9
Policy 2	31.7	31.0	30.4	29.7
Policy 3	32.6	30.4	29.3	28.4
FB	32.5	31.1	30.3	29.5

TABLE 3: Dietary Fat Intensity by Income Quartile (Median Household within Quartile)

Treatment Level for the Median Household within the Quartile (%)	Q1	Q2	Q3	Q4
Policy 0	76.8	76.8	76.8	76.8
Policy 1	65.0	65.0	65.0	65.0
Policy 2	76.7	76.7	76.7	76.7
Policy 3	76.3	76.3	76.3	76.3
FB	64.5	72.4	77.2	82.3

TABLE 4: Treatment Level by Income Quartile (Median Household within Quartile)

Distributional Impact	Policy 1				Policy 2			
	W	L	L	W	W	L	L	W
Winners and losers by income group (\$1000)	< 213.00	> 213.00	< 16.45	> 16.45				
Percentage of households (%)	98.72	1.28	2.11	97.89				
Impact by income group (\$m)	1148.36	-6.80	-13.08	1891.97				
Impact on median household within the quartile (\$)	1	2	3	4	1	2	3	4
Impact by quartile (\$m)	151	160	150	111	146	227	150	111
Share of net benefits (%)	299	333	311	198	252	471	311	198
	26.2	29.2	27.2	17.3	13.4	25.1	27.2	17.3

TABLE 5: Distributional Impact by Income Quartile Relative to Base Case:
Policies 1 and 2

Food Item	Fat			
	Intensity (ϕ)	Price per 1000 Calories (\$)	ETR %	ATP \$
Margarine	1.00	0.55	38.8	0.76
Cream	0.88	1.17	32.0	1.54
Cheese				
Eggs	0.62	2.69	19.8	3.22
Pork (lean cut)	0.57	4.34	17.8	5.11
Chicken (canned)	0.44	6.54	13.0	7.39
Yogurt (low fat)	0.24	7.70	6.5	8.20
Oatmeal	0.11	11.64	2.8	11.97
Broccoli (frozen)	0	19.68	0	19.68

TABLE 6: Effective Tax Rates and After-Tax Prices under Policy 3

Distributional Impact	Policy 3			
Winners and losers by income group (\$1000)	L		W	
	< 16.37		> 16.37	
Percentage of households (%)	2.06		97.94	
Impact by income group (\$m)	-11.81		1856.34	
Impact on median household within the quartile (\$)	1	2	3	4
	140	220	259	289
Impact by quartile (\$m)	242	456	539	607
Share of net benefits (%)	13.1	24.7	29.2	32.9

TABLE 7: Distributional Impact by Income Quartile Relative to Base Case: Policy 3