Wealth effects of maximizing palatability subject to a budget

John F. Raffensperger, PhD

Dept. of Management, Private Bag 4800,
University of Canterbury, Christchurch, New Zealand
john.raffensperger@canterbury.ac.nz

28 May 2004

Abstract

Objective: If a person eats more than the daily energy requirement, how does the palatability of the diet increase? Can the value of this palatability be quantified?

Design: We used a form of math modelling called linear programming, to find the minimum cost diet, and to find the most palatable diet subject to budget constraints. Palatability was determined in two ways, first by using the author's preferences, and second with random preferences under different scenarios for the distribution of preferences and for upper limits on consumption of each food.

Setting: The study was done in New Zealand, using the U.S.D.A. food database, and New Zealand local prices.

Subjects: The study was primarily economic, and did not require subjects. A small part of the study used the author's food preferences.

Interventions: None.

Conclusions: We found that cheating on one's diet has economic value, because a person can raise the palatability of the diet without increasing total cost. We quantify this economic value. Furthermore, we found that the value depends on the person's budget. We conclude that middle and upper income people (or poor people in a wealth society) face a high economic temptation to cheat on their diets. To our knowledge, this is the largest diet problem solved to date.

Sponsorship: None.

Descriptors: nutrition, economics, linear programming.

Introduction: the Minimum Cost Diet Problem (MinCDP)

Posed by George Stigler (Stigler 1945), the Minimum Cost Diet Problem was first solved exactly by Dantzig (1963), and has since been studied by others. Garille & Gass (2001) give a nice overview of the literature, while revisiting Stigler's original model in detail. A key problem with the MinCDP is that it tends to select unpalatable foods. The unpalatability of MinCDP solutions is well-known (Dantzig 1963, Garille & Gass 2001). Researchers have usually addressed palatability by adding bounds on foods (Dantzig 1963, Foytik 1981). Smith (1959) used a quadratic utility function. Balintfy (1964) used an integer program to select menu items. Recently, Darmon, Ferguson, & Briend (2002), taking a different approach, changed the model to minimise the difference to a typical diet, while varying the budget parametrically. While this produced diets that are close to what people actually eat, the diets did not always satisfy nutritional constraints.

The questions we aimed to answer are: (1) How does palatability change with respect to the budget? (2) How does palatability change with respect to relaxing the upper bound on kilocalories? In other words, by how much can a person improve palatability of their diet by cheating on it?

This paper's key contribution is to show that breaking one's diet has economic value, and we quantify this economic value at its optimum. People eat more than they should because they can improve the palatability of their diet without spending more money. We also give a solution for the diet problem over a substantial subset of foods in the US Department of Agriculture's database. To our knowledge, this is the largest diet problem solved to date. To help develop these results, we give a new method to measure palatability.

We chose to use the MinCDP as a framework to answer these questions for several reasons. The MinCDP neatly encapsulates nutritional and budget information in one model. The model allows exploration of many diets, while ensuring feasibility with respect to nutrients and the budget. Furthermore, we are interested primarily in information at the margin. The linear programming solution produces dual prices, which are the rates of change of interest. (A dual price is the amount that the objective will improve, if the constraint were relaxed by one unit. For example, in minimizing diet cost, the dual price on the energy constraint is the dollars that could be saved if the diet required one less kilocalorie.) The drawback, of course, is that the MinCDP is not a behavioural model. The MinCDP does not accurately reflect how people actually choose food, mainly due to problems of palatability. Therefore, we used two different measures of palatability to develop general economic principles of diet.

We used an Excel database of foods and their nutrients from the US Department of Agriculture (USDA 2002). The Excel file required quite a bit of attention for use in the MinCDP. First, it contained many near-duplicate records, such as eight records for chicken wings (roasted, stewed, raw, batter, flour, for each with skin or without skin). We removed many of these duplicates. The source file also contained a few spreadsheet errors – some data that appeared numerical (e.g. sodium in each food) was actually coded as a string.

As part of preparing the database, we obtained local prices for 693 foods in the database. Collecting these prices was a major undertaking. To our knowledge, this is the largest food database with prices. (Foytik 1981, for example, obtained prices for 160 foods.) We removed many foods with American brand names because they are unavailable in New Zealand (e.g. Breyer's ice cream). Significant differences between the US food supply and the NZ food supply became apparent, especially for foods native to the Americas, such as amaranth, squash, trout, and turkey. US grocery stores sometimes give away whole turkeys as holiday promotions, but turkey is expensive in NZ. Turkey baloney is almost non-existent. (There might have been a reverse situation with the moa bird!) Trout, for example, can be caught in NZ rivers, but cannot be sold legally. Fewer types of squash are available. Okra mercifully went unquoted; originally an African plant, it appears to be unavailable in NZ. Many differences in terminology also appeared. For example, Swiss chard is known locally as silver beet. Similarly, a rutabaga is known as a swede. In the end, we obtained 1,223 useable price quotes for 693 foods. For foods with more than one quote (usually from different stores), we used the lowest price.

The standard model for the Minimum Cost Diet Problem

Following compilation of the food and price database, we set up a standard linear programming formulation for the MinCDP. The MinCDP may be written as a set of linear inequalities (Dantzig 1963) as follows.

Indices: i foods, j nutrients.

Parameters

a_{ij} = amount of nutrient j in 100 grams of food i,
b_j = required daily amount of nutrient j,
c_i = cost per 100 grams of food i.
d_j = maximum daily amount of nutrient j,
m = the number of nutrients,
n = the number of foods.

Decision variables: x_i = 100 gram servings to eat per day of food i.

MinCDP: Min $\sum^{n}_{i=1} c_i x_i$ subject to (1)

$$\sum^{n}_{i=1} a_{ij} x_i \geq b_j, j=1,\ldots,m, \quad (2)$$

$$\sum^{n}_{i=1} a_{ij} x_i \leq d_j, j=1,\ldots,m, \quad (3)$$

$$x_i \geq 0. \quad (4)$$

For the nutrition lower bounds b_j and upper bounds d_j in constraint sets 2 and 3, we used the nutritional requirements of a 44-year-old male (the author). We solved MinCDP with AMPL/CPlex (Fourer, Gay & Kernighan 2003). The "Actual" column shows the level of nutrients in the optimal solution to model MinCDP. The minimum cost diet is shown in Table 1.

Table 1. Minimum cost solution, for the author's diet requirements.

Food	Grams/day
Wheat flour, whole-grain	210.2
Rice, white, glutinous, raw	179.3
Dandelion greens, raw	52.3
Oil, soybn, salad or cooking, (hydr)&cttnsd	36.7
Wheat flour, white, all-purpose, enr, bleached	33.4
Leavening agents, baking pdr, low-sodium	17.8
Chicken, liver, all classes, cooked, simmrd	15.5
Beans, french, mature seeds, raw	13.6
Wheat, durum	12.5
Lamb, var meats & by-products, heart, cooked, brsd	7.3
Acerola juice, raw	4.2
Cereals rte, kellogg, kellogg's complete wheat bran flakes	2.6
Wheat bran, crude	0.8

This benchmark solution serves the purpose of finding the lowest cost diet. The diet cost NZ$0.952/day. (Raw dandelion greens were given a cost of NZ$0.07/100 grams, based on the time required to collect it. This "food" is in the USDA database and all over the local soccer field, but not in local stores.) Unfortunately, as with other solutions to this problem, many people would not find this diet to be particularly palatable.

The rest of this study concentrates on two modifications to the above model. Suppose we were willing to spend, say, $1/day, rather than only $0.952. What diet should we choose? With more money, the model has a larger set of feasible diets available. Given that we are willing to spend more than the minimum, we can choose a different objective than minimizing cost. A reasonable objective is to maximize palatability.

A model for the Maximum Palatability Diet Problem (MaxPDP)

To maximize palatability subject to a budget, we modified the MinCDP in two ways. First, we changed the objective from that of minimizing cost to maximizing palatability. Second, we added an inequality constraint on the total amount of money that could be spent on the diet. The new formulation, MaxPDP, is as follows.

Indices: i foods, j nutrients.

Parameters

a_{ij} = amount of nutrient j in 100 grams of food i,
B = budget per day.
b_j = required daily amount of nutrient j,
c_i = cost per 100 grams of food i.
d_j = maximum daily amount of nutrient j,
f_i = upper limit on food i,
m = the number of nutrients,
n = the number of foods,
p_i = palatability of food i.

Decision variables: x_i = 100 gram servings to eat per day of food i.

MaxPDP: Max $\sum^n_{i=1} p_i x_i$ subject to (5)

$$\sum^n_{i=1} a_{ij} x_i \geq b_j, j=1,\ldots,m, \tag{6}$$

$$\sum^n_{i=1} a_{ij} x_i \leq d_j, j=1,\ldots,m, \tag{7}$$

$$\sum^n_{i=1} c_i x_i \leq B, \tag{8}$$

$$0 \leq x_i \leq f_i. \tag{9}$$

Note that the upper bound on energy is one of the constraints in set 7. We will denote the dual price for this constraint as the Greek letter epsilon, ε, suggesting energy. Thus, ε is the marginal increase in palatability for an additional kilocalorie.

The value of p_i is the palatability of food i. This parameter p_i is therefore the driver of the model, and the most personal aspect, since even people with nearly identical nutritional requirements can have quite different tastes. What, then, is a useful value of p_i to use in this model?

We chose to find p_i in two different ways. First, we laboriously determined the food preferences for one person, in a rather complicated way, which we describe next. While laborious, we anticipate that future personal diet-planning software would work in just the way that we specify. Second, we used random food preferences, which we discuss in a later section.

Finding p_i iteratively

How should one select a numerical value of palatability for a food? If I think I like coffee twice as much as cucumber, does that mean I have to eat twice as much? Given my coffee intake, that would be too much cucumber! If I assign a palatability of zero to liver, will I have to eat it anyway? And how should we take into account the trade-offs of palatability to cost and palatability to nutrition?

We have little hope of anyone exactly capturing correct p_i values, even when allowing for the rather strict limitations of a linear programming model. Palatability is too subjective and dynamic. We change our minds too quickly. Besides, diet and food are creative processes, not restricted to a finite list. Writing a linear program to find a maximum palatability diet is a bit like writing a computer program to determine and create the finest painting (Raffensperger 2004).

Nevertheless, we can ask people for binary choices: Which do you like better, *A* or *B*? The method we developed is to find approximate palatability in a question-and-answer manner to a set of proposed feasible diets, taking into account some of the sensitivity information from the linear programming output. The method is intended to reflect, to some extent, how a person might specify their tastes over time. Our first measure of preference p_i was developed for one subject, the author, using the following steps.

1. The parameter p_i was first initialized to the number of 100-gram servings that a person would intuitively want to eat per year. For example, the author entered initial values of 730 for coffee (meaning about two cups of coffee per day for a year), 365 for red wine, and zero for lamb liver. Since entering palatability values for 693 foods was time consuming, a default value of 100 was entered for many foods. The model was then set to maximize total preference, subject to a budget constraint, as formulated in the previous section. At this point, the model did not include upper bounds on the foods, implying $f_i = \infty$ for all i.

2. We used an approach like subgradient optimization (Held & Karp, 1970), with respect to upper bounds f_i on the foods. Subgradient optimization is a step-wise approach to adjusting penalties on relaxed constraints, to try to find feasibility of those constraints. As mentioned, the upper bounds on foods were relaxed at this point; rather than adding constraints, the palatabilities were adjusted. In response to a given solution with too much or too little of food i, the parameter p_i was adjusted after inspection of the solution. If too much of food i appeared in the solution, the p_i was lowered. If too little of a desirable food i appeared in the solution, then p_i was raised. The model was solved again and the p_i values adjusted many times.

3. Even with penalties adjusted in this way, for a sufficiently large budget, the palatability-maximizing model tended to produce solutions with too much of a given food, however palatable, such as a kilogram of strawberries. Over dozens of runs, it was found difficult to have precisely the right mix for palatability. There was often too much tomato juice, too little red wine, too many olives, or too much raw lemon. As mentioned, other researchers have simply added upper bounds on the amount of food per day. After a long while of attempting to avoid it, we found it useful to add upper bounds. The upper bound depended on the food; for example, dried bay leaf was given an upper bound of 20 grams, while cooked squash was given an upper bound of 250 grams. Upper bounds on the amount of each food had the pleasant effect of increasing the number of foods in the diet, as a result of enlarging the linear program's basis.

4. An upper bound implies that palatability is a non-linear function of quantity. When a food was at its upper bound, the dual price for the upper bound constraint (constraint set 9) was often positive, implying that increasing the upper bound would increase total palatability. This cannot be the case, since the dieter actually has *lower* palatability from raising the upper bound. We therefore adjusted the palatability p_i downwards by an amount equal to the dual price. The model returned the same solution, but with a dual price of near zero. Unfortunately, the dual prices depend on the budget, and we are interested in studying this problem over a range of budgets. We can correct the palatabilities for one budget, but then they will not be exactly right for a different budget. The adjusted palatabilities were checked over a range of budgets, and the upper bound's dual price was reasonably close to zero for modest budgets, but was higher for higher

budgets. But already, we have a strong indication that marginal palatability depends on the budget available.

At the end of this iterative process, the resulting p_i parameter is approximately one person's sense of palatability for each food i, adjusted for nutrients and budget. The iterative process seems similar to how people actually choose their food. When we have tired of cabbage, we intuitively lower our palatability for it. We expect that future diet planning software would work in this way. The user interface might recommend a meal with a lot of cabbage, and then a binary "this is too much cabbage" button would tell the software to decrease the palatability on cabbage, in a stepwise fashion. Over time, with decreasing steps, the relative palatabilities would approximate the user's true palatabilities, with only the assumptions of linearity and additivity.

For quite large budgets, the model tended to add many drinks, such as wine, beer, coffee, milk, and juices. We considered adding an upper bound on total mass, but decided against it at this point. The constraint would have been important only for very high budgets, and we have no nutritional data on recommended upper limits on mass.

With a budget of NZ$10/day, the palatability-maximizing model produced the diet in Table 1. To this researcher, such a diet would be quite palatable. Also shown below are final palatability values, found by adjusting p_i over many runs to achieve a more desirable mix. These final palatability values p_i are not the original values, but the adjusted values that produced a relatively palatable diet. Thus, they are naturally adjusted for nutrients and the budget.

Table 2. A palatability-maximizing diet, subject to a budget of NZ$10/day.

Food i	Grams/day x_i	Palatability p_i
Coffee, brewed, espresso, rest-prep	400.0	122.1
Cabbage, cooked, boiled, drained, without salt	300.0	51.0
Milk, nonfat, fluid, without vit a (fat free or skim)	284.8	66.9
Alcoholic bev, wine, table, red	250.0	357.3
Bananas, raw	200.0	209.7
Potatoes, microwaved, cooked in skn, flesh & skin, without salt	180.0	172.2
Rice, white, long-grain, parbld, unenr, cooked	168.1	138.2
Apple juice, canned or bottled, unswtnd, with vit c	150.0	85.9
Boysenberries, frozen, unswtnd	150.0	112.5
Pineapple, canned, water pack, sol&liquids	150.0	89.0
Onions, cooked, boiled, drained, without salt	125.0	72.0
Ice creams, choc	103.2	225.0
Broccoli, flower clusters, raw	100.0	124.3
Brussels sprouts, raw	100.0	100.8
Brussels sprouts, cooked, boiled, drained, without salt	100.0	89.7
Pineapple juice, canned, with vit c, unswtnd	100.0	92.7
Beef, plate, inside skirt steak, ln, 0" fat, all grades, cooked, brld	86.4	437.6
New Zealand spinach, raw	80.0	150.0
Beans, snap, green, cooked, boiled, drained, without salt	50.0	91.1
Beans, snap, green, frozen, all styles, unprep	50.0	85.2
Pineapple, raw	41.0	108.1
Peas, green, frozen, unprep	25.7	135.9
Oil, veg, sunflower, linoleic, (60%&over)	17.4	80.0

After this model had been solved for the budget of $10/day, the model was solved 202 times, over a range of budget values from $0.96/day to $100. The initial $0.96 was chosen as it is just above the benchmark minimum cost value of $0.952. In each of the 202 solutions, the budget and total palatability were recorded. Figure 1, “Kcals ≤ 2000.1,” shows a graph of the budget versus total palatability.

Palatability p_i was not adjusted at each solution; the palatability coefficients p_i used were those deemed acceptable for the $10/day solution given earlier. Therefore, the diets with budgets different form $10/day may not have truly maximized palatability, so the true total palatability for each point in Figure 1, “Kcals ≤ 2000.1,” may be higher, except for the $10 point. What a difficult problem this is! However, for development of personal diet software, a person’s budget would be relatively fixed over a reasonable planning horizon, so this would not be an impediment to implementation, and the process affords us a reasonable approximation.

From Figure 1, we see the satiety of the wealthy – after about $20/day, more money does little to improve total palatability for this data. At these luxurious budgets, the marginal palatability for another dollar is near zero, as can be seen from the slope of Figure 1, “Kcals ≤ 2000.1.” The wealthy dieter may want to add finer wines, dark chocolates, and perhaps some foie gras to the database, but total palatability will still flatten out somewhere.

Palatability and the energy constraint

What, then, can a wealthy society do to improve the palatability of one’s diet? In solving these models, we noted that the upper bound on energy remained tight. This meant that a person could raise the palatability of their diet by eating more calories than they need, *while staying within their budget*. To observe quantitatively the marginal palatability per kilocalorie, we recorded the dual price on the energy constraint for each of the 202 models that were solved for Figure 1, “Kcals ≤ 2000.1.” This dual price is ε, the marginal palatability for an increase of one kilocalorie in the upper bound on energy.

Figure 1. Palatability versus budget

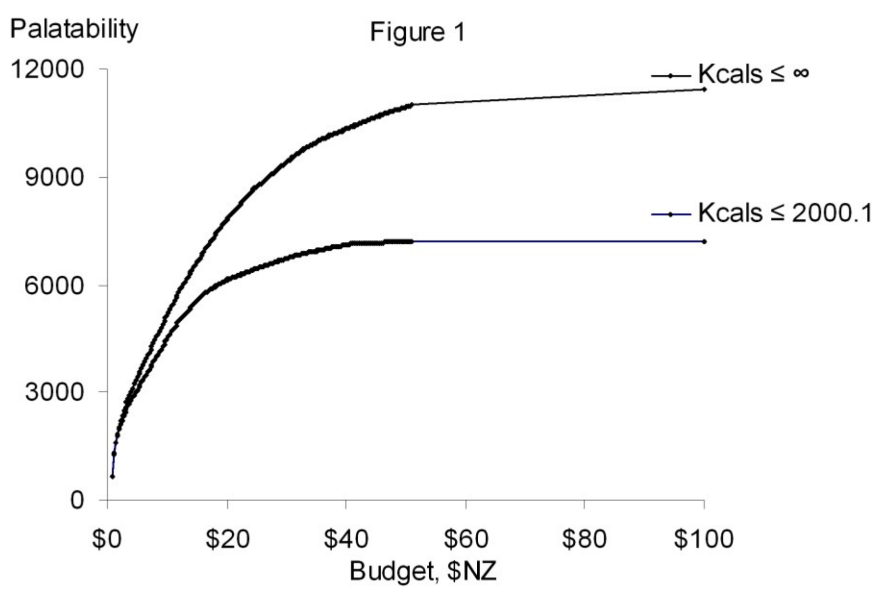

In Figure 2, we have graphed the marginal palatability of another kilocalorie, ε, with respect to the budget. Note the rounded stair step shape – there is no reason to expect this graph to be smooth or monotonic. The general shape, however, is that relaxing the energy constraint *always* results in greater palatability (since ε is positive), but *even more so* with an increasing budget. Thus, the way for a wealthy person to improve palatability is not to spend more money, but to cheat on one's diet.

Figure 2. Marginal palatability per kilocalorie versus budget

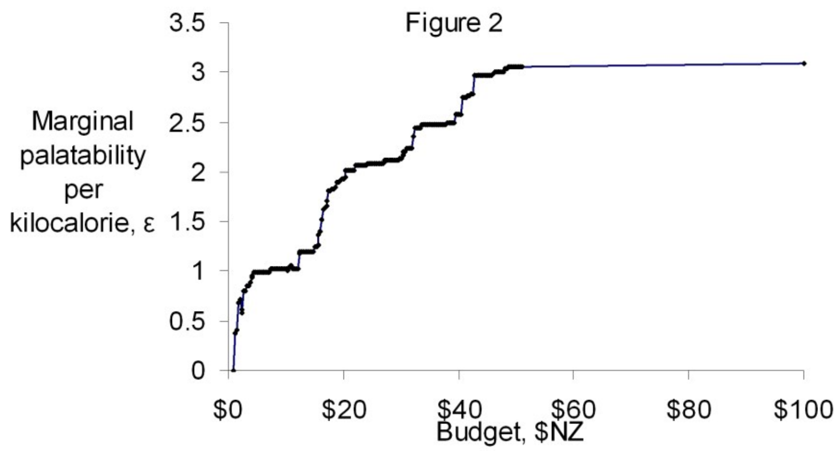

How large of an effect is this? How much can a person on a fixed budget improve the palatability of their diet by eating more kilocalories than they should? To answer this, we ran the same 202 models, but this time, we relaxed the upper bound on kilocalories. We hasten to point out that kilocalories are bounded by other upper bounds, especially the upper bounds on carbohydrate, fat, and protein shown in Figure 1, "Kcals $\leq$ 2000.1." But where will this take us? Figure 1, "Kcals $\leq \infty$," shows the answer – a great deal! We see that the dieter who cheats on the energy constraint could gain an improvement in palatability, and increasingly so with wealth.

But this study is for only one person's palatabilities. Can we expect these results to hold in general? We would expect the "increasing temptation with budget" result to hold over a wide range of tastes, because a cheap diet is so restricted that it is unlikely to be palatable (have some more liver!), hence relaxing the energy constraint a little bit will improve palatability only a little. An expensive diet is likely to be more palatable, so relaxing the energy constraint a little bit (just one more bite of chocolate éclair, please!) is likely to improve palatability a great deal. In short, relaxing the budget constraint allows foods with higher coefficients on palatability, which in turn raises the palatability per calorie. We shall see that when we switch to other measures of taste, these conclusions still hold.

Experiments with random palatability values: marginal palatability per kilocalorie ε

Since taste is highly personal, we must be careful not to generalize from a sample of one. To resolve this, the palatability study was done again many times with different random palatability values. We wrote a script in AMPL (Fourer, Gay & Kernighan 2003), which used the ODBC driver to read data directly from the spreadsheet.

Once we had a stable solution procedure, we solved thousands of LP models in the following way.

For k=1 to 100:

1. Choose a random palatability p_i for each food i in the database. We call this a *palatability set*.

2. For 100 different values of budget $B \in \{\$0.96, \$1, \$1.2, \$1.4, \ldots \$92, \$96, \$100\}$, solve a diet problem to maximize palatability subject to the budget constraint with budget B.

These steps resulted in 100x100 = 10,000 linear programs.

Furthermore, in step 1 above, we must choose a distribution from which to draw p_i. In fact, we used several different distributions. Furthermore, we do not know the upper bound f_i that people prefer on amount of each food i, nor the upper bounds on total food and water. We therefore ran the above experiment in six ways:

Experiment 1. No upper bounds on foods, total mass limited to 6 kg, preference distributed as max(0, normal(50,50)).

Experiment 2. No upper bounds on foods, total mass limited to 6 kg, preference distributed as max(0, normal(50,10)).

Experiment 3. No upper bounds on foods, water limited to 6 litres, preference distributed as max(0, normal(50,50)).

Experiment 4. No upper bounds on foods, water limited to 6 litres, preference uniformly distributed 0-100.

Experiment 5. No upper bounds on foods, preference uniformly distributed 0-100.

Experiment 6. Upper bounds on foods randomly distributed as max(0,Normal(3,1)), total mass limited to 3 kg, preference distributed as max(0, normal(50,10)).

Generally, we believe these experiments to be most valid for lower budgets, e.g., less than \$30 or \$40. Those diets look plausible, while the diets at very high budgets look implausible.

Perhaps the most unreasonable of these experiments is Experiment 5. Experiment 5 had no upper bound on each food or on total food. In this case, the model increased liquids with little or no nutrients, since those liquids tended to be inexpensive, and still had a positive palatability. The resulting diets were absurd, e.g. including 52 litres of water and 21 litres of tea. By constraining total mass or water in the other experiments, we sought to put an upper limit on the amount that a person would actually want to eat in a day. These constraints tended to be tight only with budgets over \$20. Interestingly, while Experiment 4 limited water intake to 6 litres, it resulted in diets with more calories, mainly from alcohol. Thus, constraining volume actually encouraged the consumption of calories. However, this was at the speculative high end of the budget.

We next present the results of these six experiments. From Figure 3, we see that in every case, the graph of ε rises sharply up to a budget of about \$10, then (sometimes after a bit of a fall), the graph of ε tends to flatten out, though the scale can change drastically. A preliminary set of experiments produced similar results, notably the eyehook shape similar to Experiment 4, but were deemed unbelievable and re-done.

A recent survey of New Zealand food budgets (Otago 2003) suggests that a nutritious New Zealand diet costs about $7 per person per day (close to the author's actual food budget). However, a news report (XtraMSN 2003) suggested that budgets of poor people are less than this, as low as $3.60 per person per day. Given that the minimum possible is slightly less than $1, the range of $3 to $10 seems important for people's behaviours.

Figure 3 shows that for this range of budgets, a little increase in the budget provides large absolute gains in palatability. More interestingly, over this important range of $3 to $10, we observe in Figure 3 a strongly increasing "temptation to cheat" on the energy constraint with an increasing budget. The marginal palatability per kilocalorie is increasing over this range. It appears that those who can spend more on food may be more tempted to cheat on their diets.

Figure 3. Marginal palatability versus budget over six esperiments

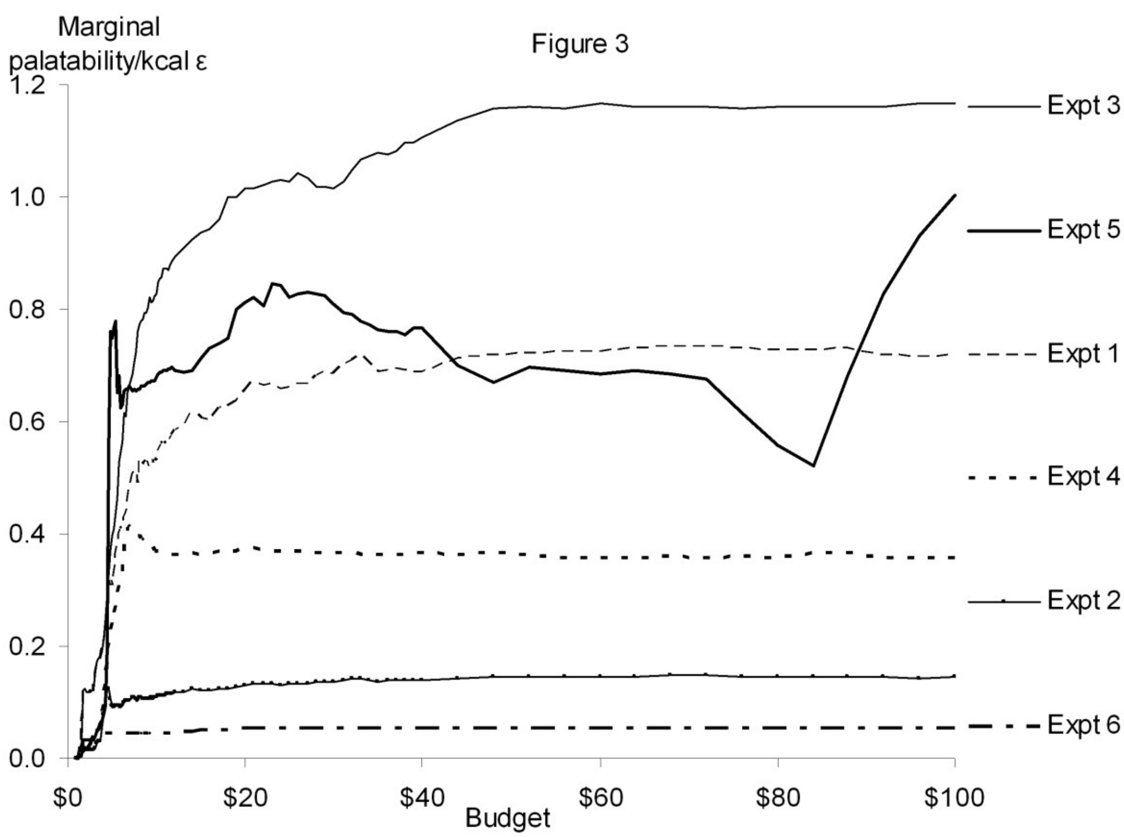

It is important to understand that, within the restrictions of good nutrition, more money allows a person to buy a better-tasting calorie for *almost any* person's taste. A severely restricted budget severely restricts choice, and this restricted choice is very likely to be an unappealing set of foods. At the very low end of the budget, the model has very little choice. A palatability-maximizing diet on a very low budget is almost identical to a pure cost-minimizing diet. There are just too few choices. With just a little more money, around $2, we have access to a somewhat wider selection of foods, a few of which are likely to be very palatable.

For budgets of more than $5 to $10, the gentle fall in ε makes more sense. The model has a limited set of foods (in spite of the empirically large database), much more limited than people really have available to them. For example, people eat in restaurants at very high cost for greater palatability, and for other factors such as enjoyment of the environment, convenience, and to socialize. If many foods (such as fine wines, dark chocolates, and foie gras) were added to the database, it is speculated that palatability per kilocalorie would tend to increase rather than decrease. Also, other constraints, especially upper bounds on other nutrients, become more binding.

So we have seen what happens at the margin. But what of total palatability? Figure 4 shows total palatability by budget for experiment 4, which is quite representative of the others, varying only in scale. The graph is similar to Figure 1. We again see the satiety of the wealthy – after about $10, little more palatability is gained for more money.

We see in a quantitative way the need for personal discipline in maintaining a healthy diet for a relatively wealthy person. The wealthy will have greater difficulty resisting temptation, because, on average, they can have a greater increase in palatability by giving into to their waistline without an increase in the budget.

Figure 4. Results of Experiment 4.

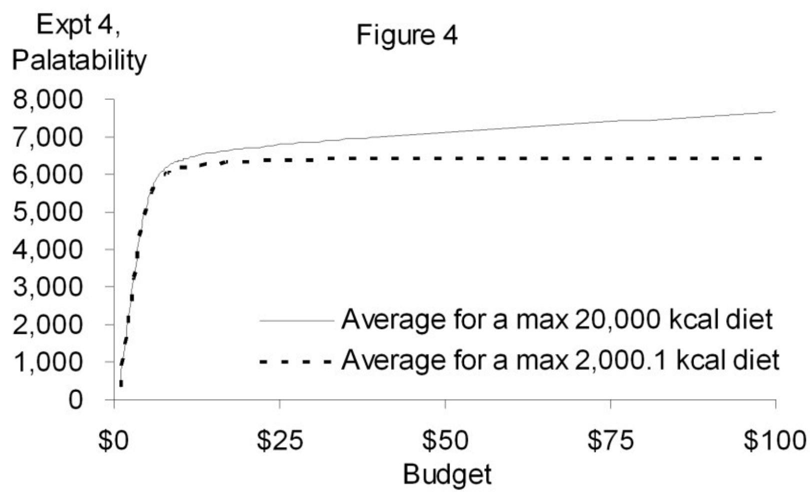

The economic value to palatability of cheating on one's diet

We have seen that someone willing to cheat on their kilocalorie constraint can increase the palatability of their diet. This has economic value, and we can quantify it. Graphically, this is the horizontal distance between the curves of palatability to budget for a restricted diet and an unrestricted diet.

Now here is a fascinating result: This difference can tend to infinity. For example, consider Figure 4, where a person is spending about $12 per day. If they cheat on their diet, total palatability is about 6,500. But – for this data – palatability of 6,500 is higher than the person could obtain on a restricted diet for *any* budget! This is shown graphically in Figure 5.

Consider the relatively poor person, spending $3 to $10. This person has an incentive to eat more kilocalories, because cheating on their diet allows greater palatability within their budget constraint. The extra calories they eat have economic value to them, because to get the same palatability with a restricted diet would cost more money. These values are shown in Table 3 for a range of budgets. We see that for a $7 budget, relaxing the energy constraint is worth between $0.26 and $0.95 for this data.

Table 3. Value (cents) to palatability of relaxing the energy constraint, for different budgets/day.

Budget	$3	$4	$5	$6	$7	$8	$9	$10	$15	$20	$25
Expt 1	1¢	3¢	6¢	16¢	26¢	41¢	62¢	90¢	312¢	838¢	4190¢
Expt 2	0¢	4	6	16	26	41	62	90	312	839	4205
Expt 3	1¢	2	9	21	45	73	123	192	808	4034	NA
Expt 4	0¢	1	10	31	95	258	751	1491	NA	NA	NA
Expt 5	0¢	1	9	23	30	30	30	31	31	33	39
Expt 6	1¢	7	17	29	51	76	115	176	NA	NA	NA

This incentive to cheat increases with budget, up to a point. Consider the wealthy person spending $20 or more. From Figure 4, we see that their palatability will level off, and they already have a high palatability. They face various upper bounds, such as on the amount of each food they want to eat, the total amount of food they are willing to eat, etc. When money is not a constraint, something else eventually becomes binding.

On the other hand, since people do not solve computer models when selecting their diets, they have no guarantee that their diets will satisfy *any* dietary constraint. Thus, the values shown here may be considered *lower bounds* on the palatability and economic value of cheating.

Figure 5. Value of cheating on the diet

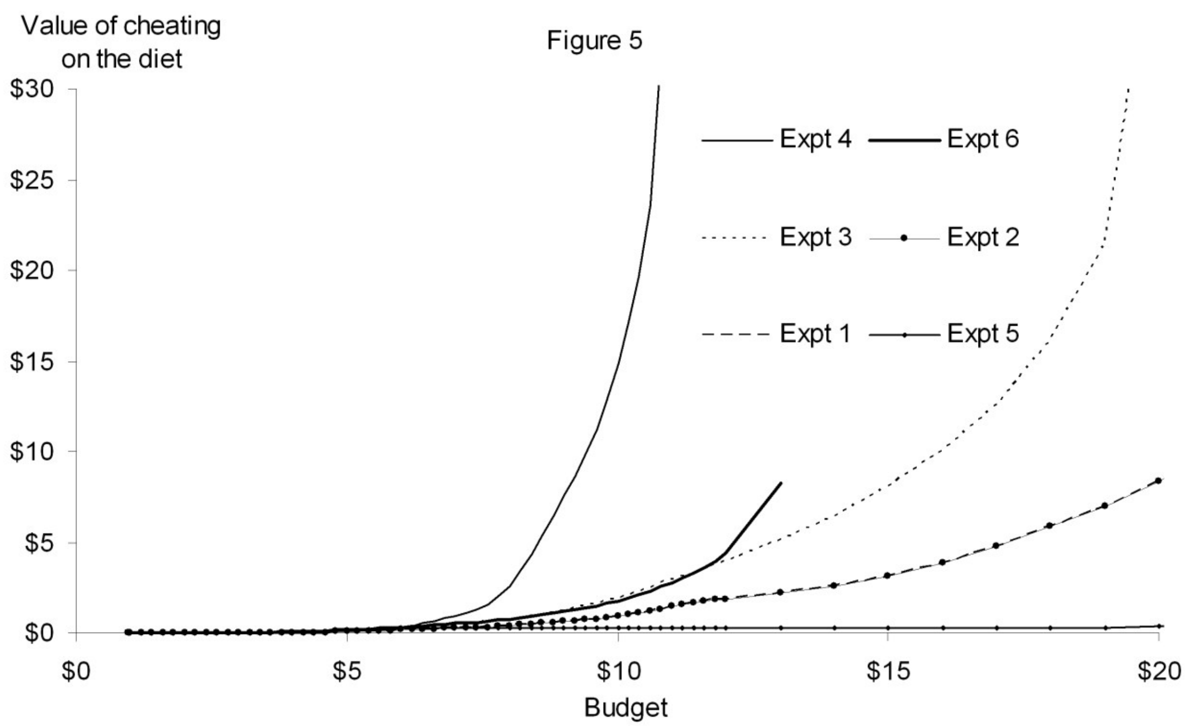

Conclusion: It is cheaper to eat more than to eat better.

This study shows that for a palatability maximizing model, an increase in the budget produces an increase in the marginal palatability per kilocalorie. The translation is that a wealthy person can raise palatability by cheating on their diet, and cheating on their diet has economic value. When people can spend more per calorie, the next calorie is more palatable than the previous. The model shows that this effect is quite strong, for a wide range of tastes, over an important range of budgets. Thus, wealthy societies are likely to have more difficulty maintaining discipline in a diet than will poorer societies, because wealthy societies have more choice.

Our model suggests the possibility of a strange exception. If the eyehook shape of Experiment 4 were found to be representative of actual diets, we would have a possible explanation as to why people of lower incomes tend to obesity. A subsistence society would be at the far left end of these curves. A poorer family within a wealthier society would likely be somewhat further to the right than a family in the subsistence society. Depending on their budget, they may actually have higher marginal palatability for the next kilocalorie than do wealthier people. We leave these questions for future study.

We speculate that some weight-reduction diets (such as the currently fashionable low carbohydrate diets) succeed partly because a large class of foods are proscribed. Restricting the allowable set of foods lowers the marginal palatability per kilocalorie, thus reducing the temptation to cheat on one's diet.

Taking this to a policy level, the results suggest that obesity in a wealthy population is not likely to improve much at the margin by taxing food unselectively, because people can improve palatability by changing their menu to increase kilocalories while staying within their budget. In fact, taxing food unselectively could make the problem worse.

The food database is available upon request from the author.

References

1. Balintfy, J.L. (1964), "Menu planning by computer," *Communications of the ACM*, v.7, n.4, pp. 255-259.
2. Dantzig, G.B. (1963), *Linear Programming and Extensions*, Princeton University Press.
3. Darmon, N., Ferguson, E.L., and Briend, A. (2002), "A Cost Constraint Alone Has Adverse Effects on Food Selection," *Journal of Nutrition*, v. 132, pp. 3764-3771.
4. Fourer, R., Gay, D.M., & Kernighan, B.W. (2003) *AMPL: A Modeling Language for Mathematical Programming*, Thomson Publishing, Pacific Grove, CA.
5. Foytik, Jerry, "Devising and Using a Computerized Diet: An Exploratory Study," *The Journal of Consumer Affairs*, Summer 1981, vol. 15, no. 1, pp. 158-169.
6. Garille, S., and Gass, Saul I., (2001), "Stigler's problem revisited," *Operations Research*, vol. 49, no. 1, pp. 1-13.
7. Held, Michael, and Karp, Richard M. (1970), "The Traveling-Salesman Problem and Minimum Spanning Trees" *Operations Research*, vol. 18, no. 6, Nov-Dec, pp. 1138-1162.
8. Raffensperger, Peter A., dinner conversation, 24 Feb 2004.
9. Smith, V.E. (1959), "Linear programming models for determination of palatable diet," *Journal of Farm Economics*, vol. 41, no. 2, May, pp. 272-283.
10. Stigler, G. (1945), "The cost of subsistence," *Journal of Farm Economics*, v. 25, pp. 303-314.

11. University of Otago, Dept. of Human Nutrition (2003), "Annual Food Cost Survey," unpublished, NZ.

12. US Department of Agriculture, Agricultural Research Service (2002), USDA Nutrient Database for Standard Reference, Release 15. Nutrient Data Laboratory Home Page, http://www.nal.usda.gov/fnic/foodcomp.

13. XtraMSN (2003), http://xtramsn.co.nz/homeandliving/0,,9764-2655831,00.html, 31 Oct. 03.