1. The short answer : cycling. Since legislators can change the taxes in each district, coalitions are unstable, as districts which suffer high taxes under some proposal can put together new coalitions.

Let m be the minimum number of votes a piece of legislation needs in order to pass, $m = (N+1)/2$. (For example, if $N = 99$, then $m = 50$.) If the legislator representing district i does not like the taxes t_i paid in her district, what she has to do is propose a bill which lowers those taxes, and which also will get the votes of $m-1$ other legislators. Under the assumption that legislators care only about the payoff $f(g) + (y_j - t_j)$ to their own district j, that's easy to do. For example, if the proposal that is currently the law has $t_1 > 0$, the legislator representing district 1 can propose a new policy : the same spending level g, but lower taxes for each of districts $1, 2, 3, \ldots, m$. For example, lower each of $t_1, t_2, t_3, \ldots, t_m$ by \$100, and raise each of the taxes in regions $m+1, m+2, \ldots, N$ so as to keep g unchanged (that is, raise each of their taxes by $\frac{m}{m-1}100$ dollars). Representatives of the first m districts will all prefer this new policy to the status quo, since it lowers their taxes (and keeps g constant), so they will vote for the proposal, and it will pass.

But then a representative of one of the "losing" districts can make a counter–proposal : lower taxes by \$100 in districts $m, m+1, \ldots, N$, and raise them by $\frac{m}{m-1}100$ in districts $1, 2, \ldots, m-1$. This new proposal is preferred by representatives of a majority of districts to the proposal that has just been passed, so that it will win in a vote.

But some other legislator could propose a further change : a tax reduction in some of the districts and an increase in others. As long as this newest proposal lowers taxes in a majority of districts, it too will pass.

So endless cycling seems likely.

Now if taxes cannot be negative, this cycling might stop when taxes have fallen to 0 in many (at least m) lucky districts. But even then, legislators in these lucky districts can propose further increases in g, paid for by higher taxes in the unlucky districts which are paying positive taxes. As long as some district i is not paying all its income in taxes, the representative of some other district can propose raising t_i a little, and using the money to increase g. That new proposal would be approved by a vote of $N-1$ to 1, since every district but i gains, due to the increase in public spending. So the only possible stopping place is to confiscate all the money in every district, and spend the money on public expenditure. Since $f'(g) < 1$ when $g = y_1 + y_2 + \cdots + y_N$, then decreasing $t_1, t_2, \ldots, t_m$ a little below $y_1, y_2, \ldots, y_m$ would make legislators in districts $1, 2, \ldots, m$ happier than spending all the money on public expenditure, so it would win a vote against a policy of spending all money on public expenditure.

Thus, **any** policy can be defeated buy a change which makes legislators in at least m districts better off. Theory predicts endless cycling in this legislature. Of course, we do not usually observe endless cycling in legislatures. Why this is so is an interesting question. But this stability cannot be explained simply because legislators can form coalitions, or practise log–rolling. The question remains : why are the coalitions in the real world so stable? Why don't people outside the coalition succeed in breaking it up by some sort of proposal which makes some of the members of the coalition better off?

2. In this case, what legislators can propose is just one number, g, since they are required to set each t_i equal to g/N. In this case, the payoff a representative from district i gets from having a public expenditure level of g (financed equally across all districts) is

$$f(g) + y_i - \frac{g}{N}$$

This payoff increases with public expenditure g if and only if $f'(g) > 1/N$. Therefore, the most preferred policy of the representative of district i is the level g^* such that

$$f'(g^*) = \frac{1}{N}$$

But this most preferred level g^* is the same for each district i. Therefore, the outcome of the legislative process here will be a public expenditure level of g^*, which will be preferred unanimously to any other. [Not surprisingly, perhaps, this is the level g^* which satisfies the Samuelson condition, since here $MRS_i = f'(g)$, so that $\sum MRS = Nf'(g) = MRT = 1$ when $g{=}g^*$.]

3. As in question 3, here legislators can proposed just 1 number, g. The rule requiring a proportional income tax means that

$$t(y_1 + y_2 + \cdots y_N) = g$$

or

$$t = \frac{g}{N\bar{y}}$$

where $\bar{y}$ is the average of the districts' incomes. That means that the payoff to district i from a public expenditure level of g is

$$f(g) + y_i - \frac{g}{N}\frac{y_i}{\bar{y}}$$

where I have substituted from the definition above of the income tax rate t. This payoff will increase with public expenditure g if and only if

$$f'(g) - \frac{1}{N}\frac{y_i}{\bar{y}} > 0$$

which means that preferences of each district's legislator are single–peaked, since $f'' < 0$.

The preferred policy g_i of the representative of district i is the g_i such that

$$f'(g_i) = \frac{1}{N}\frac{y_i}{\bar{y}} \qquad (*)$$

The right side of equation $(*)$ increases with a district's income y_i. The left side is a decreasing function of g_i (since $f'' < 0$). Therefore, the higher a district's income y_i is, the lower is its preferred level of spending g_i.

Since preferences of each legislator are single–peaked in g, the legislature should eventually agree on providing the median of the preferred levels g_i of the districts. That will be the preferred level of the district of median income, which is district m, from the convention that $y_1 < y_2 < \cdots < y_N$.

Therefore the level of public expenditure which gets chosen is the solution g to the equation

$$f'(g) = \frac{y_m}{\bar{y}} \frac{1}{N}$$

[This will equal the level g^* chosen under equal taxes in each district, in question #2, if and only if the median y_m of the districts' incomes equals the average $\bar{y}$ of the incomes.]

4. If committee members are rational, and clever, and can see ahead how the process will evolve, then the way to see what happens is to start with the last vote, after person 3 has got her chance to propose a policy. So consider what happens in that stage, when there have been two votes already, and one of the policies a, b or c is the winner so far.

From the 3 members' preferences, it can be seen that 2 of them prefer a to b, 2 of them prefer b to c, and two of them prefer c to a. Person 3, who makes the last proposal, would like to see her favourite policy c adopted. So if c was the winner of the second vote, she simply declines her chance to propose an alternative, since she likes nothing better than c. If a was the winner of the second round, then she proposes c, because she knows that both she and person 2 will vote for c over a. However, if b was the winner of the second vote, then she will **not** propose c, because she knows that c would lose in a vote against b. Instead, she would propose a. Why? She knows a will defeat b in a vote, and she prefers a to b. If b were the winner of the second vote, then a is the best policy that can win against b, from person #3's perspective.

What does all this say about the second vote? It says that if b wins the second vote, then a will wind up winning the last vote and being the policy chosen. If a or c were the winner of the second vote, then c will wind up being chosen.

Now person #2 gets to propose the policy in the second vote. Looking ahead to the third vote, he can see that his favourite policy, b, will **never** wind up the eventual winner. There is nothing he can do about that. But he can influence whether a or c actually is the eventual winner. He would rather have c as the eventual winner, because he prefers c to a. Of course, person #3 would also like to see c the eventual winner. So, in this second vote, person #2 and person #3 are in agreement : they would rather have a or c winning this vote than b. So if b were the winner of the first vote, person #2 would be better off proposing c (or a) when it is his turn, in the second stage. Proposing b, or proposing nothing if b is already the status quo, will just ensure that a is the eventual winner. So, for example, person # 2 could propose c, whatever was the winner of the first vote, and then ensure that it passes (by voting for it, along with person #3). [Person #2 could also propose a ; it doesn't matter.] Person #1, if she is clever, should actually vote for b, even if it were matched against her favourite a. Having a win the second vote just guarantees that c wins the third vote.

So, whatever happens in the first vote, person #2 can propose c (or a) in the second vote, and get person #3 to vote for it.

That means that the first vote has no influence at all on the eventual outcome. It also does not matter which policy was the status quo before the first vote. If voters #2 and #3 can look ahead from the second vote to the third vote, then a or c will win the second vote, and c will win the third vote, and there is nothing that person #1 can do about it.

5. To get a policy (Q, C) approved, the administrator must propose a policy that the legislature prefers to last year's policy. Otherwise the legislature will vote against the proposal.

Suppose, for simplicity, that the preferences of the legislature can be represented by some utility function

$$\phi(Q) - B$$

with $\phi'(Q) > 0$ and $\phi''(Q) < 0$. (That is, more services have positive benefits, but these marginal benefits decline with the level of service.)

Then the administrators would like B as large as possible, subject to $B \geq C(Q)$ and subject to

$$\phi(Q) - B \geq \phi(Q_0) - B_0$$

where $C(Q)$ is the cost of providing the service level Q, and (Q_0, B_0) is last year's policy.

If the legislature preferred the proposal (Q, B) strictly to last year's policy, then the administrator could ask for a little more more B, and still win the vote. So the administrator should always propose at (Q, B) with

$$\phi(Q) - B = \phi(Q_0) - B_0$$

Let Q^* be the legislature's most preferred level of service, the level such that $\phi'(Q^*) = C'(Q^*)$. Now if last year's service level Q is greater than or equal to Q^*, then the administrator cannot do any better than to propose $Q = Q_0$, $B = C(Q_0)$. Increasing Q above Q_0, and requiring a budget of $B = C(Q)$ would lower the utility of the legislature, since $\phi'(Q) - C'(Q) < 0$ for $Q > Q^*$. (Asking for $B > C(Q)$ would make the legislature even less happy.) So the legislature would never vote for any $Q > Q^*$, if last year's policy were $Q_0 = Q^*$, and $B_0 = C(Q^*)$.

Would the administrator ever want $B > C(Q)$? Niskanen's argument, that this is not the case, still applies here. For suppose that there were some policy proposed, (Q, B), which passed the legislature, and which had $B > C(Q)$. I have already argued that the policy should leave the legislature on the margin of choosing last year's policy instead : $\phi(Q) - B = \phi(Q_0) - B_0$. Let Q' be such that $C(Q') = B > C(Q)$. Then the legislature would strictly prefer (Q', B) to last year's policy. So the administrator could actually propose (Q', B'), for some $B' > B$ and get it passed, leaving it with a bigger budget than B. Therefore, it was not optimal to have $B > C(Q)$. As in Niskanen's model, the administrator should have $B = C(Q)$.

So the administrator's choice is to get the largest possible Q such that

$$\phi(Q) - C(Q) \geq \phi(Q_0) - C(Q_0) \qquad (**)$$

If $Q_0 \geq Q^*$, the administrator should propose $Q = Q_0$ and $B = C(Q_0)$. But if $Q_0 < Q^*$, then there is some other $Q > Q^*$ such that equation (**) holds. Figure 1 shows that $\phi(Q) - C(Q)$ is single–peaked ; if $Q_0 < Q^*$, then there is some Q_1 to the right of Q^* such that $\phi(Q) - C(Q)$ is the same height as $\phi(Q_0) - C(Q_0)$. The administrator wants this highest possible Q which is at least as attractive to the legislature as last year's policy. But the farther left is Q_0, the farther right is this Q_1. The smaller was last year's level of service (if it was less than Q^*), the higher is the level of service that the administrator can get approved this year. Figure 2 illustrates how Q_1 varies with last year's service level Q_0 : if $Q_0 > Q^*$, then the administrator cannot get the legislature to approve an increase in spending over B_0, and so might as well propose the same budget as last year ; if $Q_0 < Q^*$, then the administrator can get the legislature to approve some $Q_1 > Q^*$, since the legislature is relatively unhappy with the previous year's small budget.

Figure 1

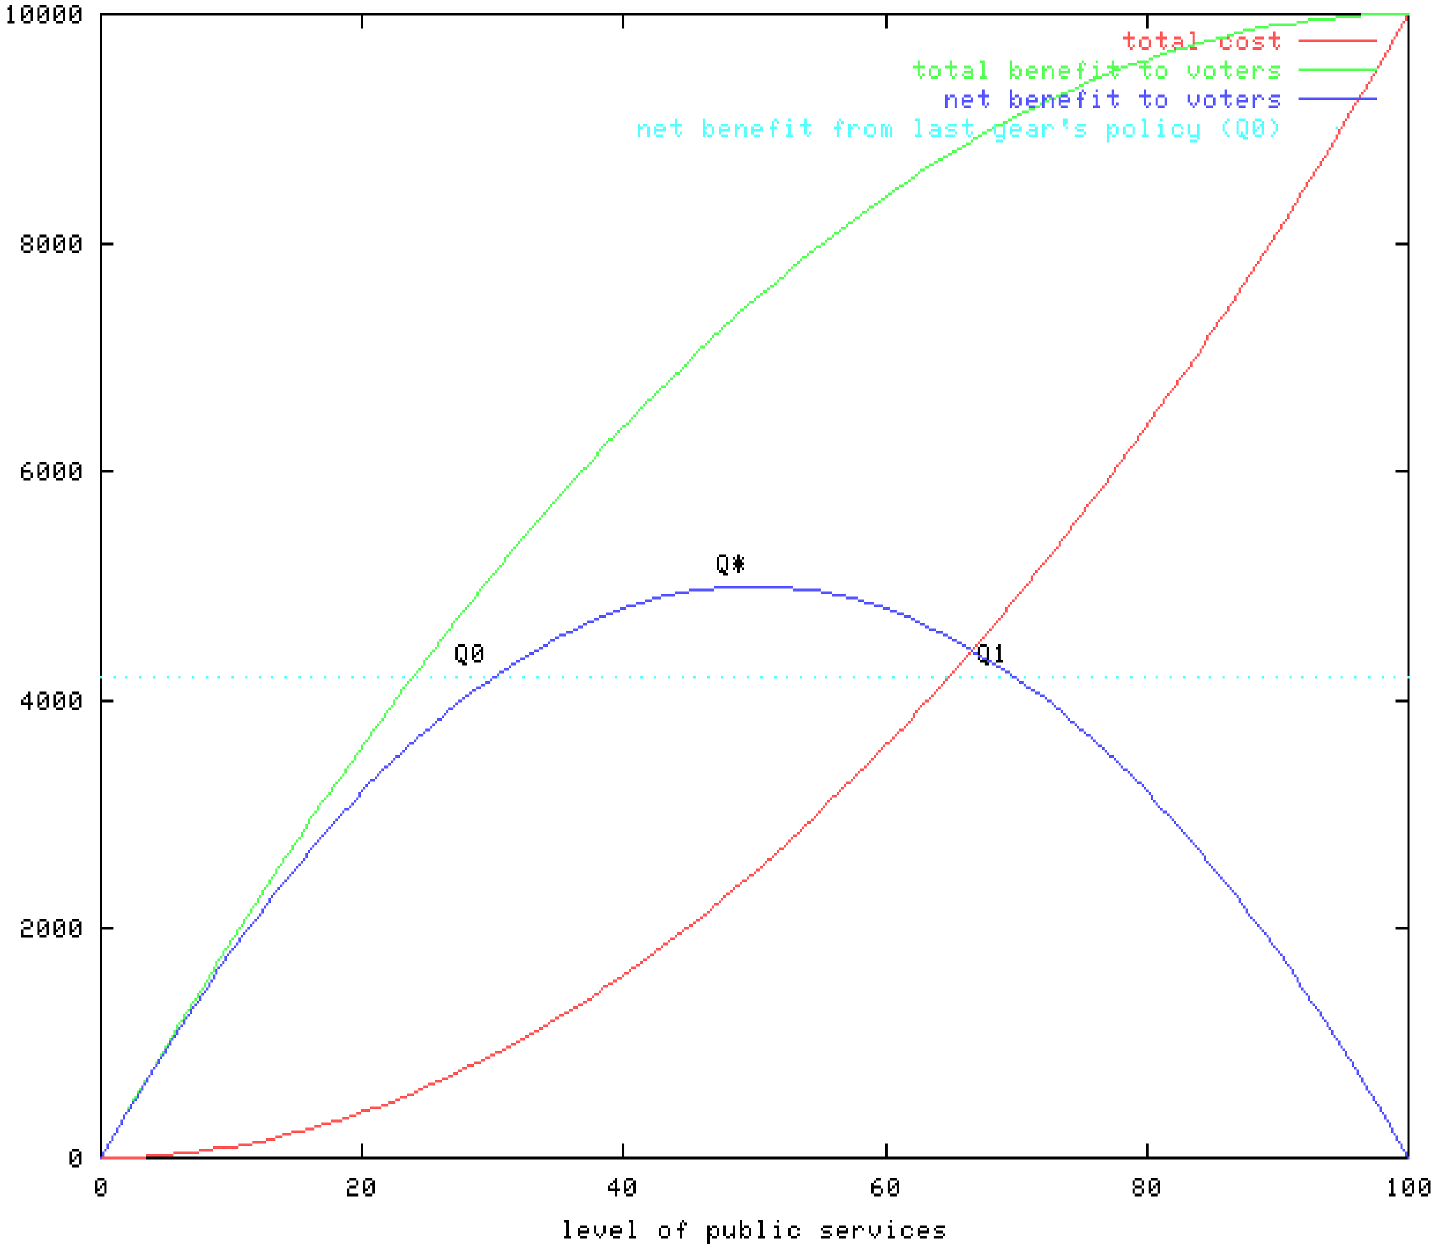

total cost
total benefit to voters
net benefit to voters
net benefit from last year's policy (Q0)
Q*
Q0
Q1
10000
8000
6000
4000
2000
0
0
20
40
60
80
100
level of public services

Figure 2

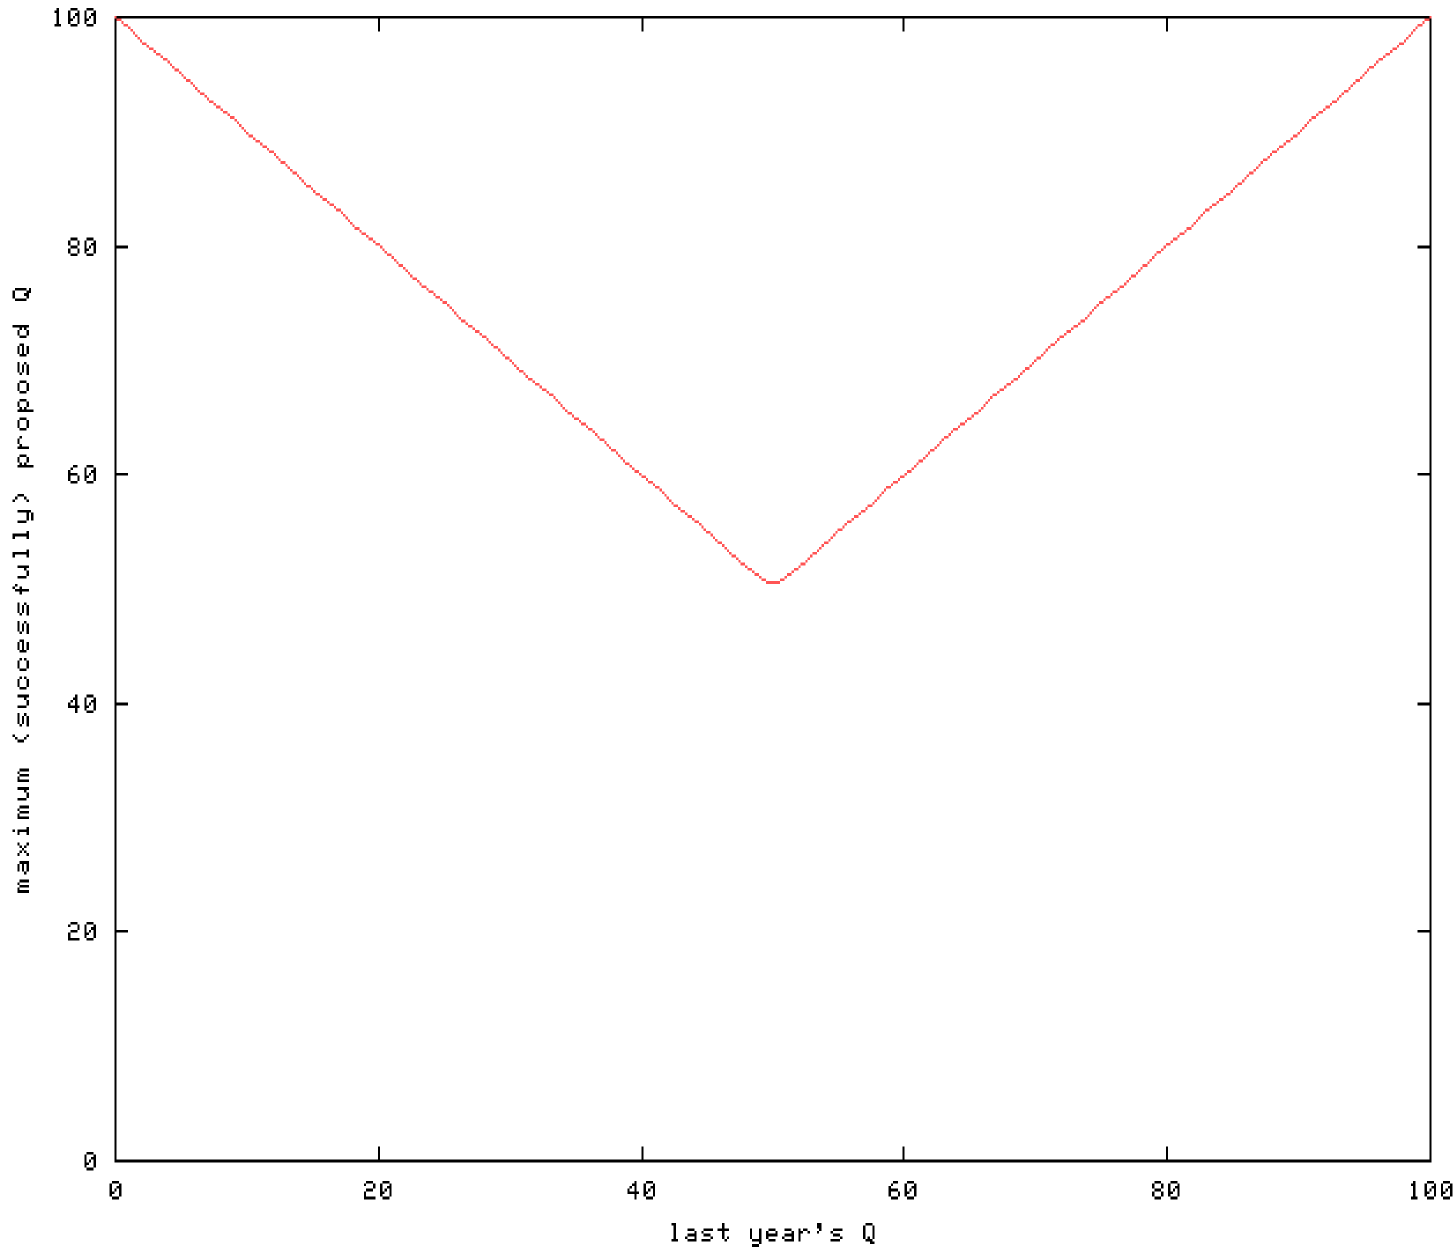

100
80
60
40
20
0
0
20
40
60
80
100
maximum (successfully) proposed Q
last year's Q