

OPTIMAL CAPITAL INCOME TAXATION IT IS

ROBERT WALDMAN

ABSTRACT. It is well known that, in standard growth models, the optimal tax on capital income goes to zero asymptotically. Even some serious economists (cough Glenn Hubbard cough) seem to have decided this means that capital income taxes should be eliminated right now. However, asymptotically we'll all be dead. One can tell more if one looks at the simplest cases of standard growth models: an aK model with optimizing consumers who have logarithmic utility or a Cass-Koopmans model with Cobb-Douglas production and logarithmic utility (OK that was wonky but it's about assumptions so I have to be honest). Consider such an economy in which the distribution of wealth is unequal at time 0. A utilitarian state would want to redistributed income. The first best way to do this is with a lump sum transfer. Let's rule that out by setting an upper limit on the tax on wealth (or an upper limit on the tax on capital income). What is the best policy?

The best policy would be to tax capital income at the maximum allowed by the assumption and use the revenues to reduce inequality until perfect equality is achieved. After perfect equality is achieved, there is no more reason to tax and taxes are zero. Thus, as noted above, the optimal tax goes to zero in the long run. The reason is that it is optimal to tax as much as is allowed by assumption so long as there is any reason at all to tax. Then stop.

A simple minded application of the model to policy would suggest that we should tax as much as we can until everyone is perfectly equal. Now the model is not the world and this would be a terrible policy. However, the argument for roughly the opposite policy is based on the same silly model plus totally turning its implications upside down by pretending that time has already gone to infinity.

The assumption of logarithmic utility is not at all innocent. It makes a huge difference. A more general assumption (which is standard in the literature) is constant elasticity of substitution utility. In this case the optimal policy depends on the intertemporal elasticity of substitution *and* on whether the state can pre-commit to a policy that it would like to change later if it could.

If the state can't precommit, the implications are just like those for logarithmic utility described above: tax as much as possible so long as there is any reason to tax at all, then stop when everyone is perfectly equal.

With precommitment, if the intertemporal elasticity of substitution is less than one (as are all empirical estimates of said elasticity) the result is to tax even more, that is to tax as much as possible until the initially rich are as poor as the initially poor, then tax them some more.

These conclusions follow from analysis using standard techniques of the simplest case of the standard model. I think that they are not noted in the literature. I conclude, as always, that to the economics profession mathematical analysis of stylized models is taken seriously exactly so long as the conclusions fit the prejudices of economists. Thus when an economist says "Mathematical analysis which you wouldn't understand of my model shows that X is a bad policy" you should hear "I don't like X."

It is a well known result that, in standard economic models, the optimal tax on capital income is zero in steady state, that is, that the tax on capital income should go to zero as time goes to infinity.

A very common response is to assume that time has gone to infinity and eliminate taxes on capital income. This conclusion is drawn by real economists not just Arthur Laffer.

Of course it doesn't follow. This post presents the simplest possible model of optimal capital taxation, motivated by a desire to redistribute income, to show just how much it doesn't follow. If one were the government of this model economy, the optimal policy would be to tax capital income at the highest possible rate until all inequality is eliminated.

Thus a dynamic model of capital income taxation can imply optimal redistribution which, in the long run, is much greater than that in a standard static model of taxation of labor income. Even though the tax is distortionary, the optimal policy is to tax as much as possible until all inequality is eliminated.

This special case is consistent with the general result, since the state which has eliminated all inequality at say time T stops taxing at time T . However, it does not because it only has access to distortionary taxes, but because it has no desire to tax. If the lump sum tax fairy arrived (by surprise) after T and made it possible for the state to redistribute at will with no distortions, the state still wouldn't tax.

There are costs due to distortionary taxes, but with optimal policy the consequence is lower average consumption and not inequality.

This result is, of course, model specific and the model is extremely stylized even by the standards of economic models (I plan to solve it typing away here and have not yet written anything down). A key assumption is that utility is logarithmic, this makes things much simpler. For a general CES utility function, the optimal policy depends on the intertemporal elasticity of substitution in consumption. For elasticities lower than one (that is all the ones that fit any data) the result is even more extreme as the optimal policy makes those who start rich poorer than those who start poor (he who is first shall later be last for the income effect is stronger than the substitution effect).

The Model (I warned you it was simple)

$$(1) \quad Y = aKsor = a$$

GNP is a constant times capital. No one works. This assumption is not critical; similar results hold for a Cobb-Douglas production function with labor.

Agent i chooses C_{it} to maximize the present discounted value of the log of consumption subject to a lifetime budget constraint. The rate of time preference is ρ .

$$(2) \quad \rho = r = a$$

So with no taxes the consumption of each agent is constant. This is not needed for the results.

At time 0, there are two groups of people each with measure one of members: The rich who own the capital and the poor who own no capital. For simplicity only assume each rich person owns the same amount of capital.

From now on r is the representative rich person and p is the representative poor person so consumption at t is

$$(3) \quad C_t = C_{rt} + C_{pt}$$

The rich own K_{rt} and the poor own K_{pt} .

There is a utilitarian state which aims to maximize the sum of utils. Clearly it will do something as the poor consume nothing without its help so laissez faire welfare is negative infinity. Government consumption is zero, the state just taxes and transfers. It has wealth K_{gt} so

$$(4) \quad K_t = K_{rt} + K_{pt} + K_{gt}$$

In the model K_{gt} will be positive so public debt is negative, that is, there is a public endowment. This is very important.

If the state can tax at will, it will immediately seize half of initial wealth K_o and give it to the poor and then never do anything again. This is the first best outcome with optimal lump sum taxation.

To make the problem interesting, assume that the state can't confiscate capital. In particular, I assume so for tax rate τ_t with τ_t less than or equal to 1, revenues are $\tau_t r(K_{rt} + K_{pt})$. The exact maximum rate is not critical. It can be 50% or 200% and the qualitative results don't change.

I also have to assume that τ_t is a integrable function of t and I will assume that it is continuous except at a finite number of points of discontinuity (one in the optimum). This last is a standard assumption in optimal control, but it is usually not stated.

The state uses the revenues to finance transfers to the poor subject to an inter-temporal budget constraint that the present value of tax revenues is equal to the present value of transfers.

In this model (and in the more general models studied in the literature) the optimal public debt is negative, that is the state builds up a public endowment (this is critical for the results including the result that τ_t goes to zero).

The poor are liquidity constrained at time zero. The utilitarian state is very paternalistic and uses this fact to control the consumption of the poor, that is, to make it what it would be without distortionary taxation. Formally, the state can control the consumption of the poor and so it will be optimal.

Optimal consumption of the poor is constant (give $\rho = r$) and equal to the present value of transfers from the state. This can only happen if the poor are liquidity constrained so the state makes sure that they are. Optimal $K_{pt} = 0$.

The state gives a constant transfer to the poor who rationally choose to consume exactly all of it.

The non-trivial problem is how does the state get the money?

Now the assumption of logarithmic utility has a very convenient property. The income and substitution effects of different after tax interest rates exactly cancel so consumption is a constant times wealth.

$$(5) \quad C_{rt} = \rho K_{rt} = r K_{rt}$$

One very convenient implication of this is that there is no possible problem of dynamic inconsistency of the optimal policy. Threats and promises about future taxes have no effect on the rich so a state which can pre-commit to some tax plan will impose exactly the same tax plan as the state which can't.

This also means that the state can control K_{rt} so long as the restriction $\tau_t \leq 1$ is not binding and that the state can always change policy to increase C_{rt} .

Now consider the effects of capital income taxation on capital formation.

$$(6) \quad \frac{\partial K_{rt}}{\partial t} = -\tau_t r K_{rt}$$

K_{rt} is decreasing in τ_s for all s less than t .

C_{pt} is chosen by the state. C_{rt} is increasing in K_{rt} .

So, for fixed transfers to the poor, C_t is decreasing in τ_s for all $s < t$.

The standard argument shows that τ_t goes to zero as t goes to infinity. In fact, there is a T such that optimal $\tau_t = 1$ if $t > T$ and optimal $\tau_t = 0$ if $t > T$.

Recall that the state can control K_{rt} so long as the restriction $\tau_t \leq 1$ is not binding and τ is continuous at t , and that the state can always change policy to increase C_{rt} . The state's problem can be re-interpreted as maximize welfare choosing C_{pt} and C_{rt} (except when the restriction is binding) subject to a social budget constraint that the present value of total consumption discounted at rate r is equal to K_0 . This means that for any t_1 and t_2 , if τ_t is an integrable function of t with a finite number of points of discontinuity which don't include t_1 and t_2 and $\tau_{t_1} < 1$ and $\tau_{t_2} < 1$ then

$$(7) \quad \frac{1}{C_{rt_1}} = \frac{e^{-r(t_2-t_1)}}{C_{rt_2}}$$

This is a standard result in optimal control.

Now given optimal savings by the rich

$$(8) \quad \frac{1}{C_{rt_1}} = \frac{e^{\int_{t_1}^{t_2} -r(1-\tau_s) ds}}{C_{rt_2}}$$

So if $\tau_{t_1} < 1$ then for any t_2 such that $\tau_{t_2} < 1$ then the integral of τ_s from $s = t_1$ to $s = t_2$ is zero.

Given continuity of τ at t_1 if $\tau_{t_1} < 1$ then there is a τ_2 very close to t_1 such that $\tau_{t_2} < 1$ so τ_t must be zero in the interval from t_1 to t_2 so, using continuity again, τ_{t_1} must be zero.

Thus the optimal tax must be 1 for a while, the jump to zero and stay at zero for a while.

So far I haven't excluded the possibility that it is switched back up to 1. The math so far says that τ_t must be zero or one.

This is standard optimal control (including the continuity except for a finite number of jumps assumption which is rarely stated explicitly).

Recall also that τ_t can always be cut so the state can always cause C_{rt_2} to increase. This means that if $t_2 > t_1$ then the integral of τ_s as s goes from t_1 to t_2 must be less than or equal to zero. Since τ_s must be zero or 1, the only way this is possible is for τ_s to be zero for all s $t_2 > s > t_1$.

That is the optimal policy is of the form There is a T such that if $t < T$ $\tau_t = 1$ and if $t > T$ $\tau_t = 0$.

So far this is standard analysis. Note that τ_t does indeed go to zero as t goes to infinity. In this simple model it goes all in one jump from the maximum allowed to zero.

Now consider $t > T$ so $\tau_t = 0$. It is possible to prove that for all $t > T$

$$(9) \quad C_{rt} = C_{pt}$$

This is actually easier than proving that τ_t jumps down from 1 to 0.

The argument is very simple. If τ_t is zero for $t > T$ then the deadweight cost of introducing a tiny tax on (or subsidy to) capital income with $|\tau_t| = \epsilon$ for $t_3 < t < t_4$ with $t_3 > T$ and $t_4 > T$ is second order in ϵ .

If C_{rt} is not equal to C_{pt} and the tax revenues are given to the poor (or taken from them if the tax is negative) there is a first order in ϵ effect on welfare due to redistribution. Therefore, unless $C_{rt} = C_{pt}$, there is some ϵ so small such that a tiny tax or subsidy is better than taxing exactly zero.

Thus the optimal policy is to set the maximum tax on capital income which is allowed by the assumptions and keep the tax at the maximum until perfect equality is obtained.

Then, as noted by Judd and Chamley, the state stops taxing capital income. It does this because its aim was equality and it has achieved perfect equality as fast as possible given the upper limit imposed by assumption on the tax on capital income.

1. GENERALIZATIONS

The math below will be even less rigorous and less comprehensible (if possible) than the math above.

The key assumption is that utility is logarithmic. This means that future taxes have no effect on behavior so there is no possible dynamic inconsistency problem. This happens because income and substitution effects exactly cancel. In a more general model the instantaneous utility is a CES function of consumption

$$(10) \quad u_t = \frac{1}{1-\sigma} C_t^{1-\sigma}$$

All existing empirical estimates of σ are greater than one (and believe me I published a paper in which it would have been very very nice to be able to assume $\sigma < 1$).¹

This means that for someone who has only capital income, the income effects of future taxation of that income are greater than the substitution effects. That is if they know they are going to be taxed more in the future, they consume less.

I think the analysis that shows that τ_t is 1 from $t = 0$ to some T and then zero works fine for any CES utility function. Just substitute the marginal utility of consumption where you see $1/\text{consumption}$ and it all goes through.

It is also clear that, if the state can not precommit, then for $t > T$ $C_{rt} = C_{pt}$. That only depends on the utility function being concave.

¹Alessandra Pelloni and Robert Waldmann (2000) "Can Waste Improve Welfare?" The Journal of Public Economics. vol. 77 pp 45-79.

However, if the state can pre-commit, it might pre-commit to a T such that C_{rt} is not equal to C_{pt} for $t > T$. Define T_1 as the t such that if τ_t is 1 if $t < T_1$ and 0 if $t > T_1$ then for $s > T_1$ $C_{rs} = C_{ps}$.

In the model, the inefficiency due to the restriction on tax policy which rules out lump sum taxes and transfers is that the consumption of the rich is too high for all $t < T$. Choosing T slightly greater than T_1 will reduce this consumption (income effect stronger than substitution effect). This will mean that, for $t > T$, the initially rich are poorer than the initially poor and $C_{rt} < C_{pt}$. This imposes a cost which is second order in $T - T_1$ while the benefit of making the rich consume less for $t < T$ is first order in $T - T_1$. Therefore, if $\sigma > 1$, the optimal policy is more extreme than taxing the rich as much as possible until perfect equality is obtained. It is to tax the rich as much as possible until they are slightly poorer than those who were poor at $t = 0$. Thus, for $\sigma > 1$, for optimal taxation of capital income, there is more than 100 % redistribution. Optimally, he who was first shall later be last.

In contrast many of the other assumptions aren't critical at all. Below I return to the assumption that utility is logarithmic so the consumption of the rich is a constant times their wealth and the state has no reason to threaten to make them poorer than the poor in the long run.

As noted above, the reasoning of the last section is much more general than the model. It is not necessary to assume that the poor start with exactly zero wealth. If they start with some wealth, they consume all of it and then live on the transfers. The state would like to prevent this but it can't. During the period when the poor have privately owned wealth, the analysis that says τ_t must stay one or jump permanently down to zero is based on them too, but it still holds. The analysis that after T $C_{pt} = C_{rt}$ goes through fine even if the poor still have private wealth at T .

Now go back to the assumption that the poor start with zero wealth.

The results obtained in the previous section obtain if there is a Cobb-Douglas production function and the poor are endowed with labor which they supply inelastically (and numerous enough compared to the rich to be poorer than the rich). In this case, for optimal tax and transfer policy, the poor consume their labor income plus the transfer. With a Cobb Douglas production function with constant labor supply wages are a constant times K so the consumption of the poor is a constant times K and their interests are served by maximizing K by minimizing the consumption of the rich. The way to do that is to tax capital income at the highest rate possible. In the Cobb-Douglas case, less redistribution is needed to achieve perfect equality so T is lower.

I think the analysis also works, for the Cobb-Douglas production function, if the rich have labor income from inelastically supplied labor as well as capital income. I'm not sure of this but I don't see anywhere in the argument anything about the rich not having labor income.

From now on I return to the aK model. In the discussion below I assume that no one works and that all production is pre-tax capital income.

Similarly, I think a similar analysis works if the state must make the transfer equally to all citizens and so must give back to the rich. Like wages in a Cobb Douglas model this means that the income of the rich is a constant larger than $r(1 - \tau_t)$ times K , however they still get $r(1 - \tau_t)$ on their savings. I don't see how

this changes the result that τ_t must be one for a while and that if it falls to zero it stays there. I certainly don't see how it can change the result that if $\tau_t = 0$ then $C_{rt} = C_{pt}$. I think the effect of the restriction that the transfer must be equal to everyone is that T is infinity, that is the state taxes capital income at the maximum possible forever.

Now go back to the assumption that the state can give to people it indicates by name but must tax a function of an observable quantity (this is what the US constitution says). Thus the state can give to the poor and not give anything back to the rich.

I don't know where to put this, but a radical change in the assumptions of the model has no effect at all on the analysis. It works if the tax is used to finance government consumption not transfers. Consider a model with all citizens starting with perfectly equal wealth. Individual utility is the log of individual consumption plus a concave function of government consumption. Note I assume that utility is additively separable in personal and government consumption. This is important. For one thing, it implies that individual wealth will remain equal. Note also that the shape of the function of government consumption doesn't matter much and doesn't have to be the same for different people. The argument that τ_t is one for a while then falls to zero and stays there didn't depend at all on the reasons the state wanted the money. The conclusion that after T $C_{rt} = C_{pt}$ now becomes after T the marginal utility of individual consumption for each and every individual is equal to the sum over citizens of their marginal utility of government consumption. That is, if $r = \rho$, the fact that the state can't tax lump sum and must rely on distortionary taxes has no effect on the constant value of $\frac{G_t}{C_t}$ for $t > T$. For $r = a > \rho$ I need to know about the income expansion path of demand for G . The result holds if utility is the log of personal consumption plus a constant times the log of government consumption.

Now I go back to the model of redistribution. From now on $G = 0$ and there are poor people who start out with zero wealth.

The assumption that the state can build up an endowment is critical. One might argue that such a policy would reduce economic efficiency as the state would meddle with management of firms it owns. This assumption is critical to the published analysis that leads to the conclusion that τ_t must go to zero as t goes to infinity. However, I think the results hold even if the state is not allowed to build up an endowment.

I now assume a balanced budget constraint so that transfers to the poor at time t equal capital income tax revenues at time t . This means that the state can no longer keep the poor liquidity constrained and control their consumption. That was the point of building up a public endowment. If τ_t goes to zero (it will) then the poor must save. If they save their capital income plus a fixed fraction of the transfer, then their consumption is proportional to that of the rich so long as the transfer is positive. This means that it is optimal since so long as they save they have the same Euler equation as the rich. The fraction of the transfer which is saved must be positive if τ_t goes to zero as t goes to infinity.

This means that, so long as the transfer is positive higher taxes on the rich imply a more rapid increase in $K_t = K_{pt} + K_{rt}$. The argument that T must be 1 or zero works. The state has the same degree of control over their consumption that is total if $\tau_t < 1$ and can increase the consumption of the rich and decrease the

consumption of the poor if $\tau_s = 1$ for all $s \leq t$. The first order conditions (7) and (8) for rich and the analogous equations for the poor are the same. if $\tau_{t_1} < 1$ and the integral from t_1 to t_2 of τ_s is not zero for some t_2 , then the state can change policy in a way that satisfies its budget constraint and helps the rich and hurts the poor (the change which was good if the state can build up an endowment) or do the opposite and help the poor and hurt the rich. The second would increase welfare so long as the marginal utility of the poor is greater than that of the rich at time t_1 .

So I think the results hold even if the state is not allowed to build up an endowment. This restriction is very costly, because the poor like the rich will consume more than is optimal for $t < T$, but I think that the optimal policy is still there is a T such that τ_t equals one for $t < T$ and 0 for $t > T$ and that for $t > T$ $C_{rt} = C_{pt}$. The effect of the balanced budget restriction is to reduce the consumption of both for $t > T$ but it should still be equal.