An Agent-Based Model of Taxes and their Effect on Wealth Distribution

Presented to the faculty of Lycoming College in partial fulfillment of the requirements for Departmental Honors in Physics

by

Matthew Velardi

Lycoming College

May 4, 2022

Approved by:

Christopher W. Kulp, PhD

David G. Fisher, PhD

Michael D. Kurtz, PhD

Andrew Brandon, PhD

1 Abstract

We created two separate models of an economy to measure the effects of different tax schemes on their wealth distribution, The first model contained non-interacting agents who passively earned money over their life span and paid estate taxes before bequeathing their wealth to the next generation. We found that increasing the tax rate and decreasing the proportion of people exempt from the tax, along with targeting payouts from collected taxes to the poorest 25% optimized our social welfare functions, though there were diminishing returns after about a 40% tax rate and 90% exemption. The second model used exchanges between agents to model income and mostly recreated the results of the first model, with the exception of the fact that targeting the payouts reversed its effect on the distribution. Our attempt to correct this failed but led to an interesting new avenue for further research.

2 Introduction

In this project, I examined how taxes affect the distribution of wealth in a model economy which consists of agents exchanging money with each other in a random fashion. The taxes examined are estate taxes (taxes on wealth inherited when an individual agent dies) and income taxes (annual taxes based on the income an agent earned in the previous year). I measured the taxes' effect on the agents' wealth distribution using the Gini coefficient, which measures inequality on a scale from zero (all wealth equally distributed) to one (one person owns everything).[1] I have also utilized other Social Welfare Functions, which attempt to quantize the well-being of a society through different methods, from simply showing the wealth of the least wealthy person, to showing total wealth of all people, to more complicated combinations of these. In summary, I have worked to determine the effect of estate and income taxes on the wealth distribution of an economy as measured by the Gini Coefficient and other Social Welfare Functions.

In this examination, I aimed to complete previous work and to make good progress on branching this out. The previous model I aimed to complete was examining how estate taxes

(taxes on wealth inherited when an individual agent dies) affect the distribution of wealth in a model economy which consists of many non-interacting individuals who earn money innately, unlike the aforementioned final model. We had good baseline results, but there were still some technical issues to be worked out when it comes to the tax structures. After completing these fixes, I created a model which combines the non-interacting model with another model which examines how income taxes (based on the difference in wealth from one year to the next) effect an economy of immortal agents who gain money through exchange with each other, with no innate income. To this end, my objective was to introduce the dynamics of dying agents and estate taxes into an exchange-based economy, and measure the effects of both income and estate taxes on the wealth distribution of this system.

3 Methodology

3.1 Accumulation Model

To accomplish this, I used an approach called Agent Based Modeling in the programming language Python, starting with a model I had already constructed alongside Drs. Kulp and Kurtz. In this prior research, we model an economy of people as a list of numbers, representing the wealth of each person; these are the "agents." Each of these agents also have randomly assigned attributes called "lifespan" and "skill." The model iterates over time, and in each time-step (representative of a year) the agents earn an amount of wealth based on their skill, and if they have reached the end of their lifespan, they "die" and a new agent is "born" in their slot, inheriting the wealth of the dead agent. The new agent is assigned a new lifespan and skill (the skill value can either be randomly selected or correlated to the agent's predecessor based on the model run). Crucially, though, the wealth inherited is processed through one of several estate tax systems — each full run of the model uses a different system for comparison. The models that I used in this project are all flat taxes with an exemption below a certain dollar amount. So a flat percentage of only inheritance over the exempted amount is taken in taxes. In the model, this amount is given by the wealth of an agent with more wealth than some percentage of agents, which is chosen to match historic exemptions. The collected taxes are then redistributed amongst agents, either targeting the payout of collected taxes to the least wealthy 25% of agents or to all agents. The model is stopped after a simulated 700 years and the Gini Coefficient and Social Welfare Functions are calculated to be compared to other simulated tax policies. The other SWFs accounted for are the Median Sen (a combination of the median wealth value and wealth inequality as measured by the gini), the Rawlsian (simply the value of the wealth of the least wealthy agent), and the Utilitarian (the sum of all wealth in the system). Lastly, to account for the random nature of the simulations, each run is actually 20 runs with the same settings, called Monte Carlo runs, from which the average results are taken.

The accumulation model also includes a rudimentary application of the Carnegie Conjecture, the idea that people who receive large inheritances work less ardently. This is modeled by assigning each agent an "effort function" which reduces their skill score by an amount that depends on their anticipated inheritance, loosely matching the work presented in "Heterogeneity of the Carnegie Effect."[2]

As stated, my first step was to finish this model, and this mostly involved fixing some issues with how taxes and exemptions were calculated. There were some issues regarding updating former US tax codes for inflation (i.e. prior to the fix, a value of $10,000 from the tax code from 2000 would be calculated using $10,000 in 2020, instead of the inflation-adjusted $15,000 in 2020). There was also an issue of calculating how much inherited wealth is exempted when calculating taxes. These issues were adjusted to more accurately emulate real tax codes: using properly defined monetary amounts and only taxing the amount of wealth an individual inherits above the exempted value. Once these corrections, among a few other minor ones, were in place, we gathered final results form this model, and moved to the next stage.

3.2 Modified BDY Game

This next step was to create a new model, incorporating ideas developed in [3]. In this work, the authors utilize a similar agent-based model, but they use agents who interact to model the effects of income and estate taxes on wealth distributions, without considering the effects of the death of agents and estate taxes on the wealth distribution. Thus, I set out to build a model which incorporates both systems to create an even more accurate representation of an economy.

I started with a base of interacting agents and income taxes and added in death and thereafter inheritance of wealth and estate taxes based upon this inheritance. The interactions work by collecting a unit of money from each agent who is able to give (i.e. each agent who owns at least one unit) and paying out each collected dollar to a randomly selected agent. Thus a number of dollars are collected from able agents and dispersed randomly amongst all agents. This is an accelerated version of the Bennati-Drăgulescu-Yakovenko (BDY) game [4], such that all able agents will interact enough to calculate taxes within 365 time steps in order to calculate income taxes properly. Thus, I increased the number of time steps from 700 to 700×365, still 700 years, but now each time step is a day to allow for many interactions per year. I used four different progressive income taxes, based on US tax codes from 1981, 1988, 2016, and 2018. After implementing this exchange-based income system, I added in a system for agents to die and thus inherit wealth, and of course taxes on this inherited wealth, in a manner similar to the previous estate tax model. I ran this new model using both, neither, and each tax system separately to get some initial results.

Part of this process of taking initial data was implementing targeted payouts of collected taxes. Our previous models targeted the payout based on how much wealth people had. Actual welfare systems use income to determine this. This shouldn't be too different in real life, and it is easier to check for welfare agencies. However, in our random-exchange-based economy, having money is entirely disconnected to earning money; the richest person is just as likely to receive money as the poorest. So if someone accumulates money and begins to lose it, they will be taxed less and likely receive payouts, but someone trying to "climb the ladder" will be taxed and not paid. Thus, targeting the payout keeps the rich rich and the poor poor, lessening the impact of tax and redistribution instead of compounding it.

We wanted to keep the realism of the income-based payouts, but also wanted our model to keep true with the intuition that targeted payouts will decrease inequality. Fortunately, we believed the next step I had planned to take in the process of building this model would help with this. This next step was to connect the amount of money one has to their likelihood to get money. This was accomplished by adding askill score to each agent, distributed according to a distribution which closely approximates the distribution of income. We did this so that when picking pairs of agents to interact, the agent with a higher skill score is more likely to take a dollar from the lower skilled agent than vice-versa. Specifically, when deciding who receives payouts, we added in a step which compared a randomly generated number to an odds ratio composed of the potential giver's and potential taker's skill scores. If the potential receiver wins, they receive the dollar, while if the potential giver wins, they get to keep their dollar. This method was used in order to keep agents debt free; any agent can receive any number of dollars each round, but they can never give away more than one. This indirectly ties income to wealth, since high skill agents are more likely to have high income and therefore high wealth levels, as opposed to previously when income was purely random. Doing this would increase inequality in general, but it would also allow targeting to decrease it. Since the wealthy should always earn money, they would be taxed and not paid, while the poor who can't easily earn money in this system are propped up.

4 Results and Discussion

4.1 Accumulation Model

The first four figures show the Gini Coefficients and other SWFs at the end of runs with specified exemptions and flat estate tax rate. Also noted in the figures is whether the skill value is inherited by new agents or if it is instead randomly assigned. These figures are all charts of an SWF as a function of the flat tax rate of the model system, with curves given for each of four exemption levels, corresponding to historic exemptions at the 87th, 97th, and 99th (current) percentiles, as well as one with no exemption.

Figures 1 and 2 show the SWFs for the first two models. These models both do not account for the Carnegie effect, but they are different in that Figure 1 does not have targeted redistribution while 2 does. Figures 3 and 4 do the same, but for models with the Carnegie effect. I will first comment on the general trends we will see in all such figures. First, all exemptions seem to have the same general shape in each chart, except for that of no exemption. This is because, in these runs, the poorest agents are not protected from being taxed, so at high tax rates, they tend to stay poor easier. This is easily seen in the Gini and Rawlsian curves, where it leads to greater inequality and lower wealth to the poorest agent, respectively. Other than this anomaly, we see that lower exemptions tend to perform better in these SWFs, as they allow for more money to flow from those with much wealth to those with very little; if you only tax 1% of people, even if they're the richest 1%, you won't get enough money to change the world. The last general trend is that the Utilitarian SWF tends to be stable around 3.8×10^8, meaning that nothing in the tax structure is affecting the total wealth accumulated enough to be visible.

As for trends between charts and figures, the first is that inheritance of skill tends to cause models to perform worse in our SWFs, since it allows for intergenerational inequality. This is where children of poor agents stay poor (since their skill stays low), and vice-versa for rich agents, while with random skill assignment, a poor agent could have a child with much greater skill to become rich, and a rich agent could have a not very skilled offspring which squandered the inheritance it receives after taxes. The largest effect visible between figures is that the aforementioned no exemption high-tax anomaly shrinks with a targeted payout of taxes. This is simply due to the fact that the poorest agents, while still taxed, receive more of that back (raising the Rawlsian), and the richer portions receive nothing, so inequality drops (lowering the Gini), though it is worth noting that the effect is still visible, so targeting does not completely undo the anomaly. The other major inter-figure observation is that the Carnegie Effect has little impact on any of the SWFs. We wouldn't expect to see much change in the Rawlsian, since the Carnegie Effect only affects the richer portions of the population. But we might expect something in the Gini and Utilitarian, since richer agents should earn less money, lowering both inequality and total wealth, especially with inherited skill, yet we see very little if any difference from Figure 1 to Figure 3 and from Figure 2 to Figure 4. This is perhaps explained by the fact that the Carnegie effect noted in [2] is quite small and only affects a small portion of the agents.

4.2 Modified BDY Game

Figures 5-9 display results from this model prior to adding in skill-based exchanges, while the rest show results with this addition. I will treat each figure separately, starting with Figure 5. This figure shows the Gini of three runs of the model with both taxes, sampled every 100 years. These runs each start with a different starting distribution of money amongst the agents, flat (everyone the same), normal (bell curve), or Pareto (a lot with very little, a few with a lot). These charts are helpful to determine how long to let a model run before it hits equilibrium, or before the behavior is stable. In each case, equilibrium starts within 100,000 time steps, or 300 years, so my 700 year model should be fine. The only other thing to notice are the jumps in the Pareto chart, as these are where the first few generations die and thus we see the effect of estate taxes all at once. This is because all agents start at age 0, which I later corrected by starting all agents at a random age less than their assigned lifespan.

Figure 6 pertains to the exchange-based model with only estate taxes, so we see similar charts to the first few figures, but only for the Gini and Rawlsian. The Utilitarian does not make sense in this model since no money enters or leaves the system. We see that the Gini follows the same trend, lowering somewhat with increased tax rate, and a small uptick at the end from the no exemption model. The Rawlsian, on the other hand, tells us something different. All values are below 1, so what this says is that there seems to always be someone with no money at all.

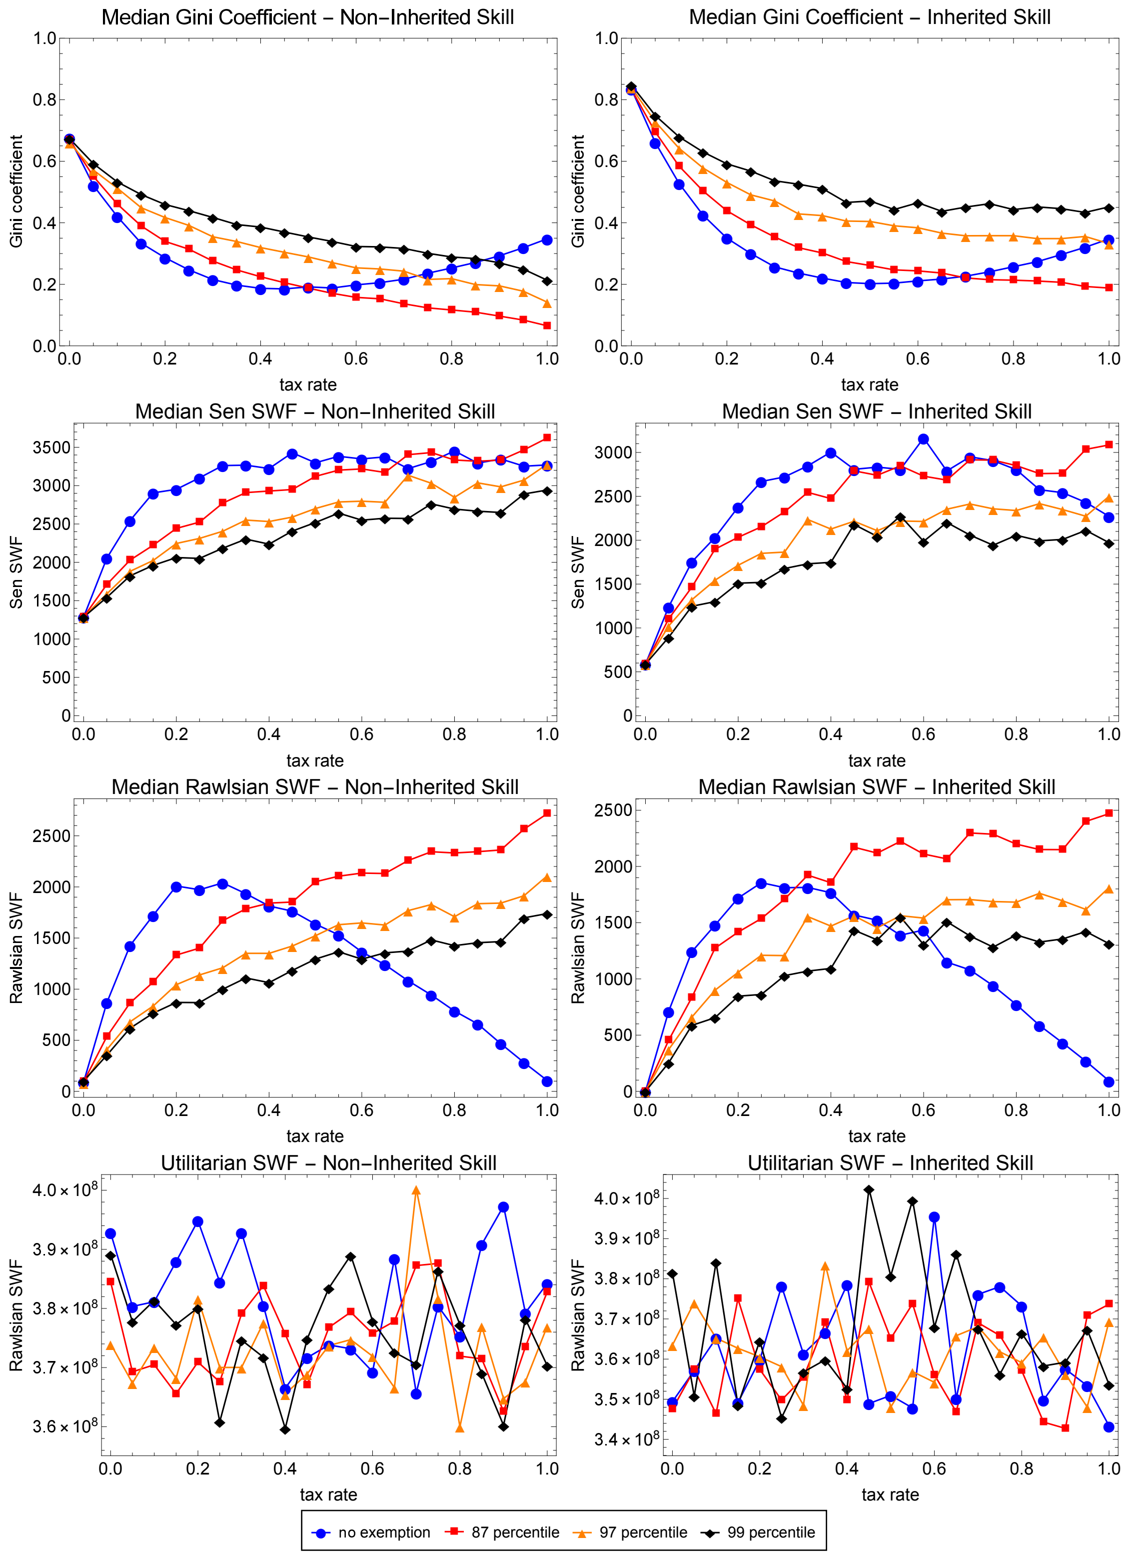

Figure 1: Accumulation Model with non-targeted payouts and without the Carnegie Effect

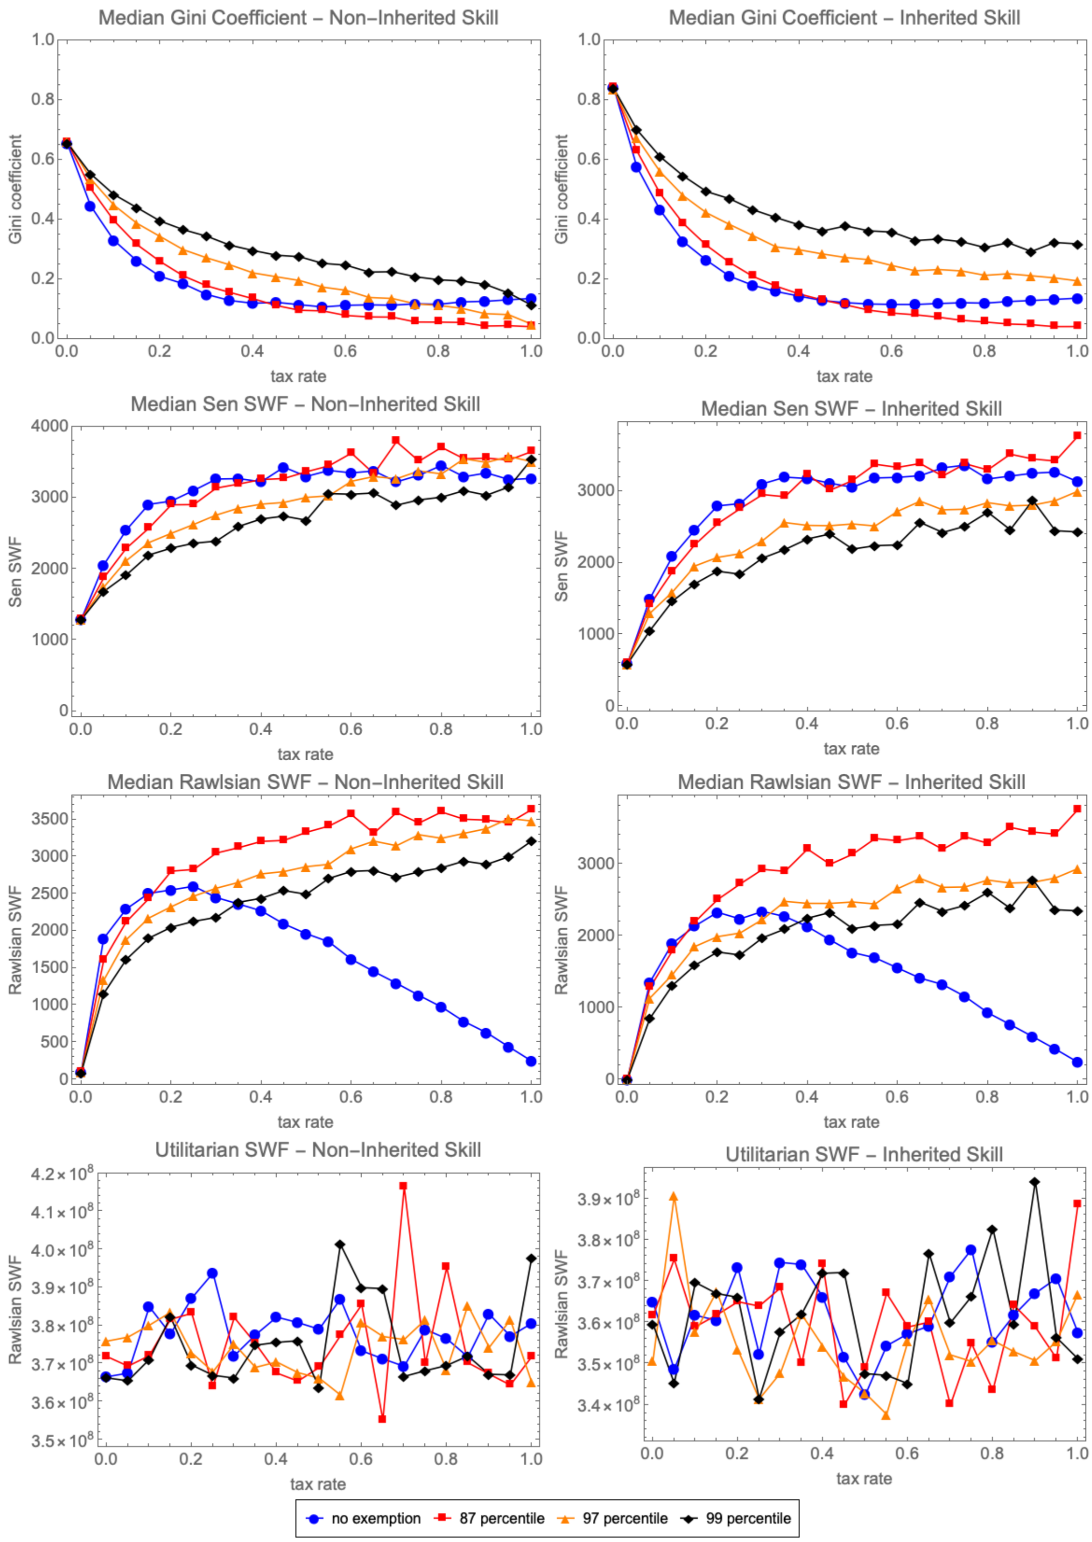

Figure 2: Accumulation Model with targeted payouts and without the Carnegie Effect

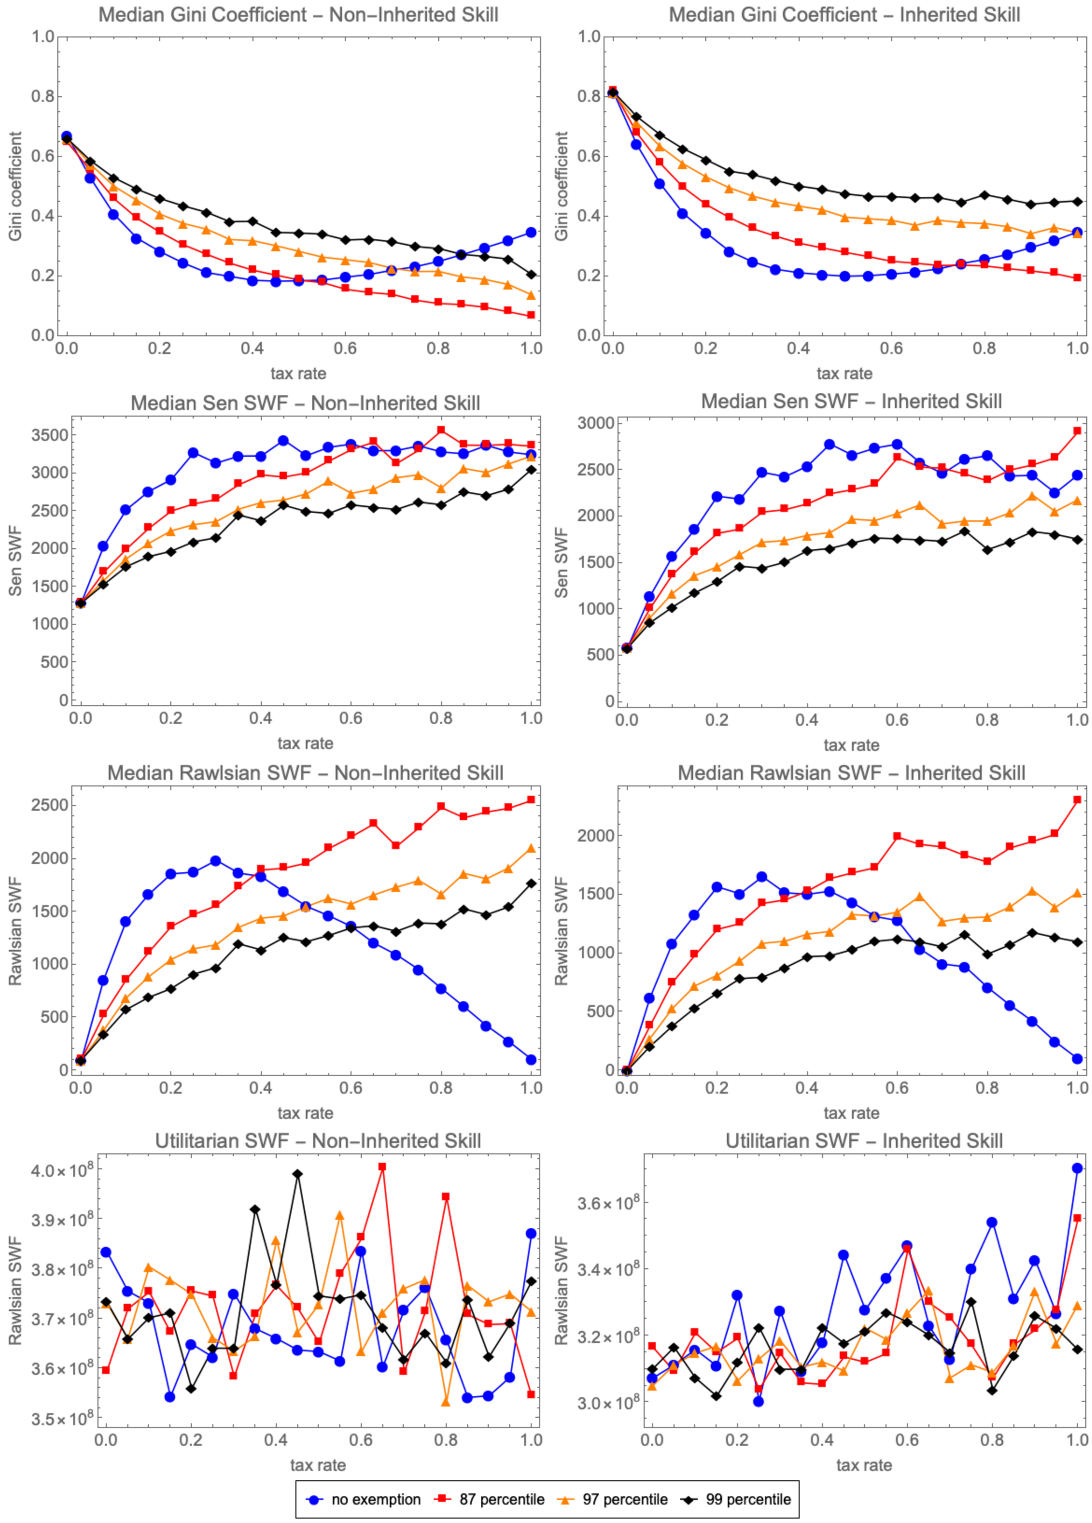

Figure 3: Accumulation Model with non-targeted payouts and accounting for the Carnegie Effect

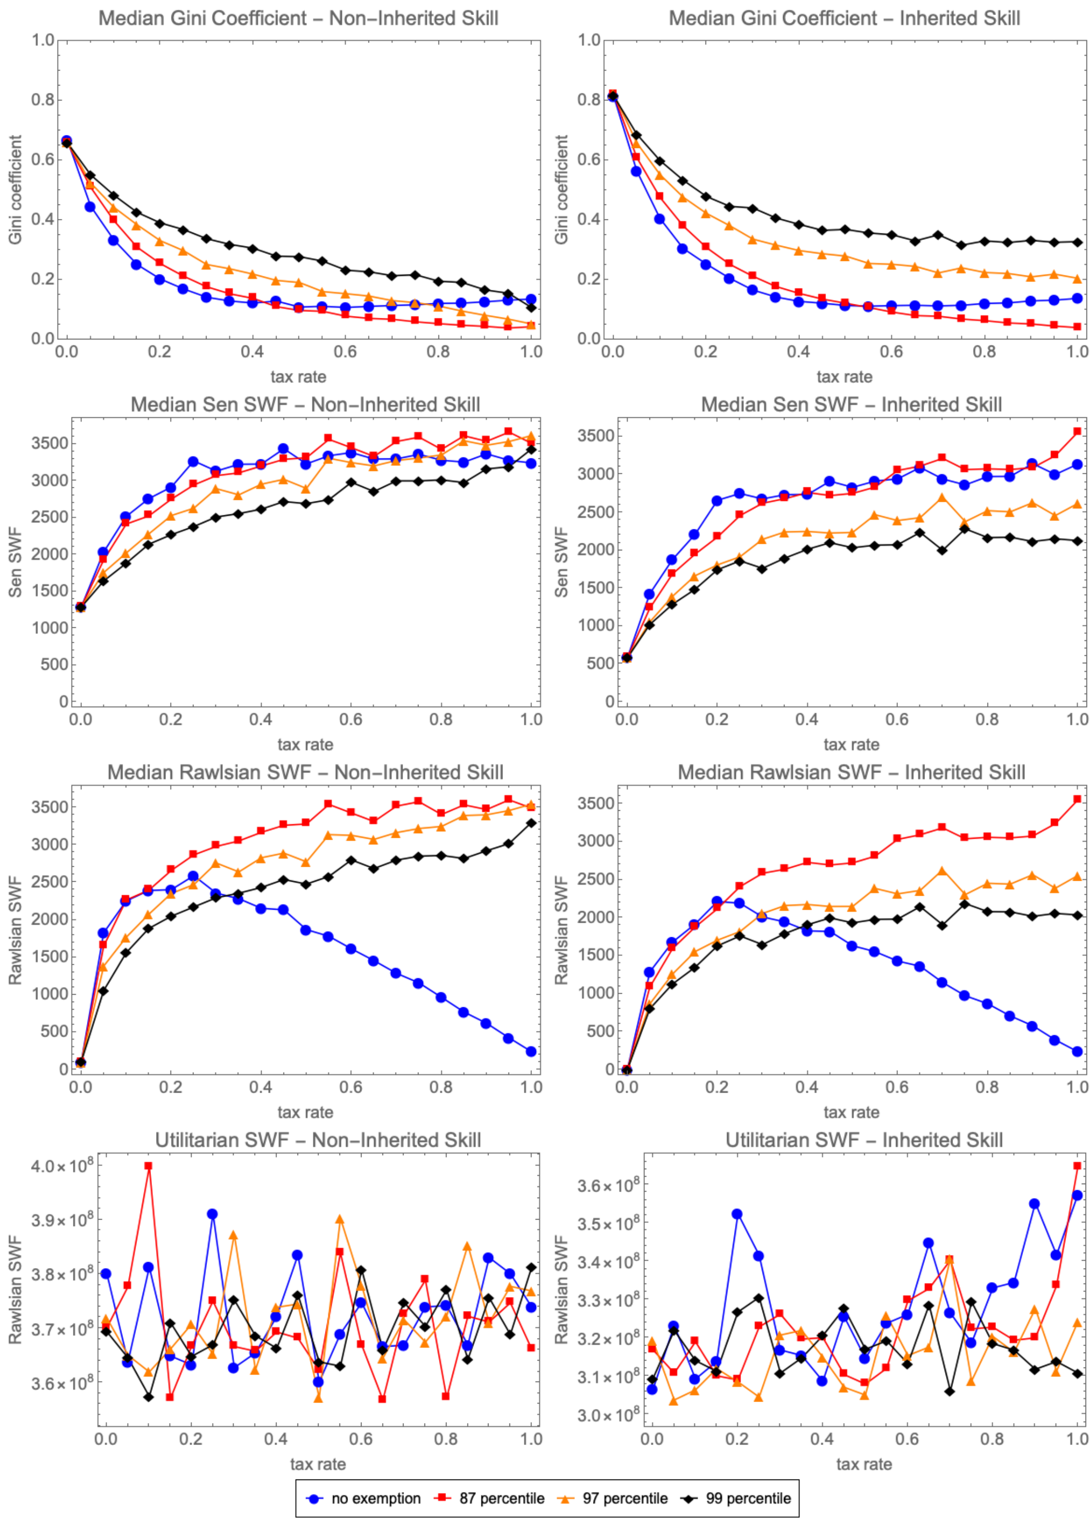

Figure 4: Accumulation Model with targeted payouts and accounting for the Carnegie Effect

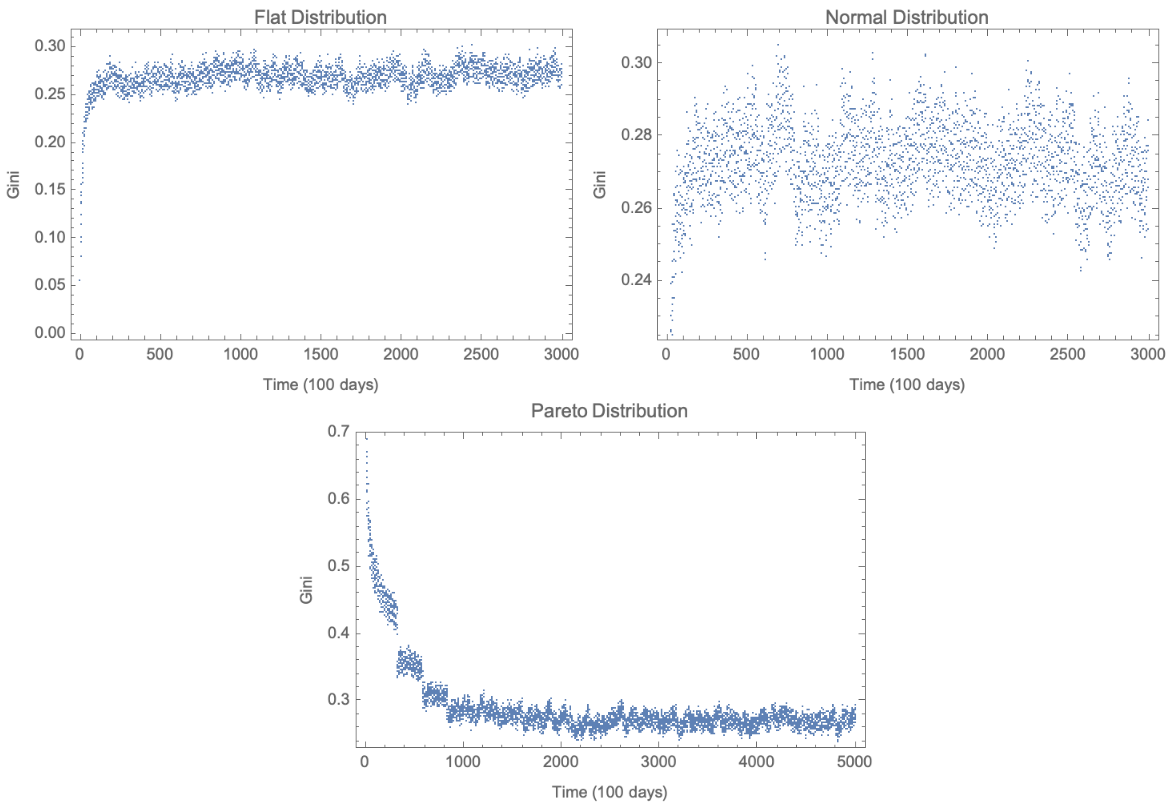

Figure 5: Gini Coefficient over time of exchange based models starting with different initial distributions.

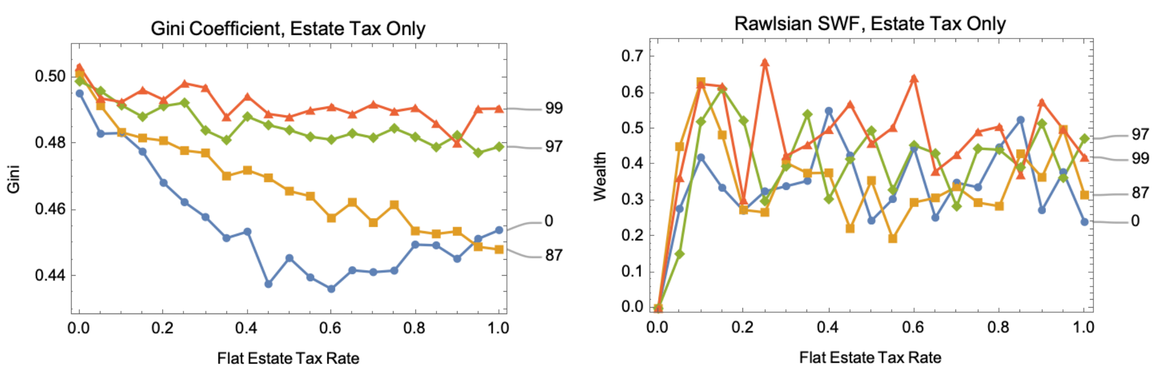

Figure 6: Gini and Rawlsian for basic Exchange based model, only estate taxes

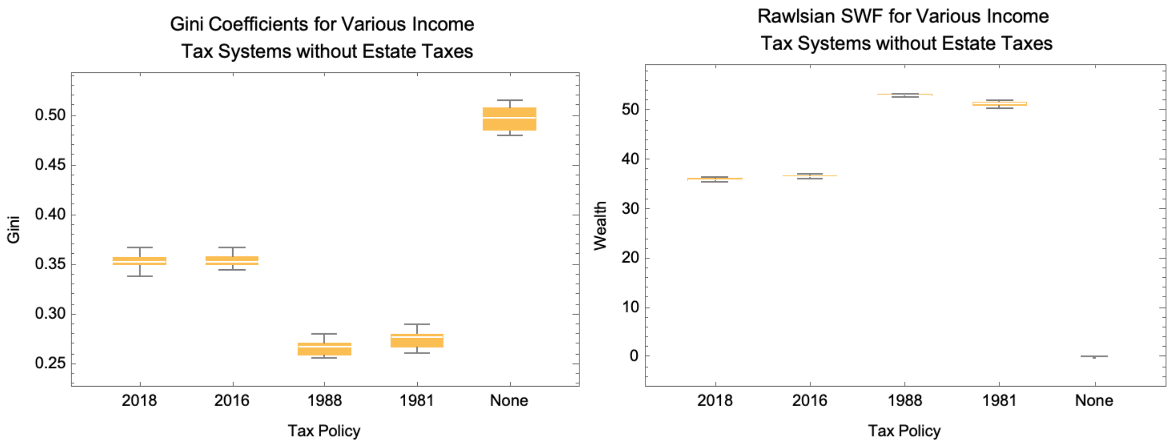

Figure 7: Gini and Rawlsian for basic Exchange based model, only income taxes

Figure 7 once again shows only the Gini and Rawlsian, but this time for a model with only Income taxes. In this case we see the tax code on the x-axis again, but now they are not flat. Instead they are demarcated by the year of US tax code they are meant to exemplify. We see that the pair of years in the 2010s and the pair in the 1980s are similar, which is strange because both reflect what could be huge changes. The 1980s runs include a model with 16 brackets (1981) and one with only two (1988), which shows that tax codes don't necessarily have to be complicated to produce desired results. The change in the 2010's is one from democrat to republican, as 2018 represents President Trump's tax code, while 2016 represents President Obama's; these, too, are similar, despite the hubbub surrounding the change. Also in this figure is the results for no tax at all, i.e. a natural system of exchanges. Interestingly, this is the only one that reproduces the Rawlsian of the Estate tax only model of 0, and has much more inequality than either pair of tax codes.

Figures 8 and 9 show the results (Gini and Rawlsian, respectively) of the full exchange-based model. They are in the same form as the previous estate tax charts, and broken down by which income tax code they use. The Gini curves follow the same basic trend, and much like the only income tax models, the decade pairs are similar in value in both figures. The Rawlsian figure is interesting in that it seems the inclusion of income taxes alleviates the issue seen in the estate-tax only figure, where the Rawlsian was essentially zero.

Implementing targeting in the income tax only model shows the problem noted earlier with

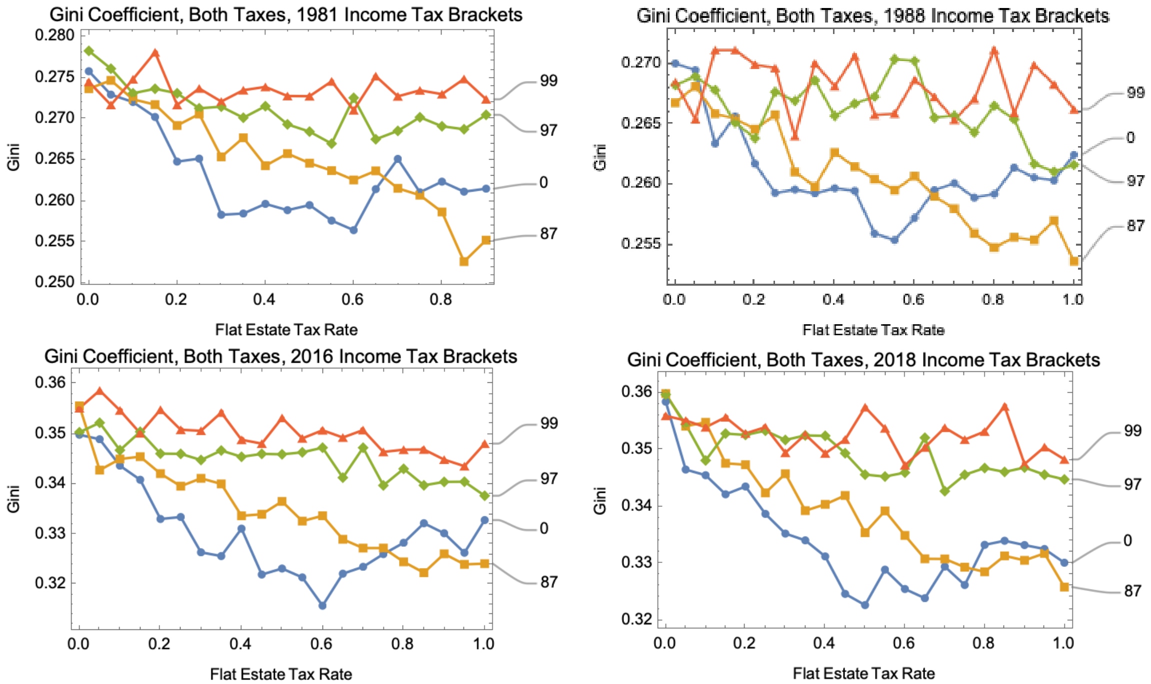

Figure 8: Gini Coefficient for basic Exchange based model, both taxes

Figure 9: Rawlsian SWF for basic Exchange based model, both taxes

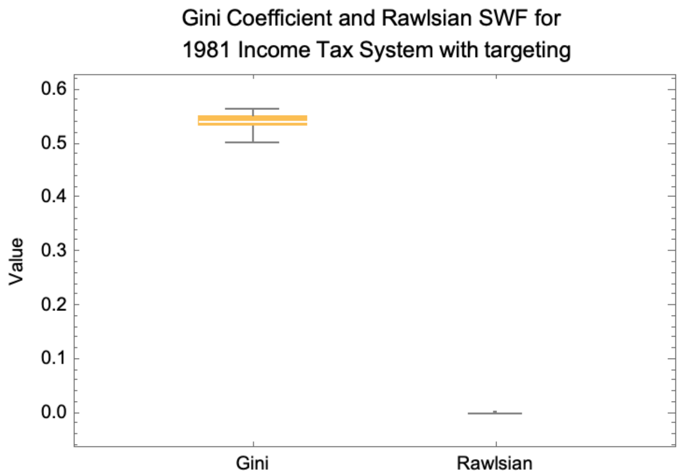

Figure 10: Basic exchange model with targeting

income-based targeting. In Figure 10, we see that the Gini coefficient of a model utilizing the 1981 tax model (around 0.55) is significantly higher than that of this model without targeting (around 0.28). In this figure, we also see that the Rawlsian is now 0. These both indicate that targeting based on income keep the poor poor and the rich rich, as discussed previously.

The remaining figures pertain to the latest version of the model, after inclusion of a skill-comparison, our intended fix for the targeting issue. These results are shown in figures 11 and 12. In Figure 11, we see that the Gini in models both with and without targeting are significantly higher than without the skill comparison, and visa versa for the Rawlsian. However, we also note that this did not fix the targeting issue, with a significantly higher Gini and lower Rawlsian in the model with targeting.

The explanation for this lies in Figure 12. This figure shows the actual wealth distributions for the same models as Figure 11. In each, there is a large number of people with very little money and very few with a lot of money. However, the difference lies in how the money is distributed amongst those who have very little. In the case of an economy without targeted payouts, all of these people have about an equal amount of little money, though not zero, since they all receive the payouts from income taxes. However, in the economy with targeted payouts, we see that approximately a quarter of these people have received payouts, and that these payouts are thus larger, allowing them to attain some wealth. The Gini coefficient sees the

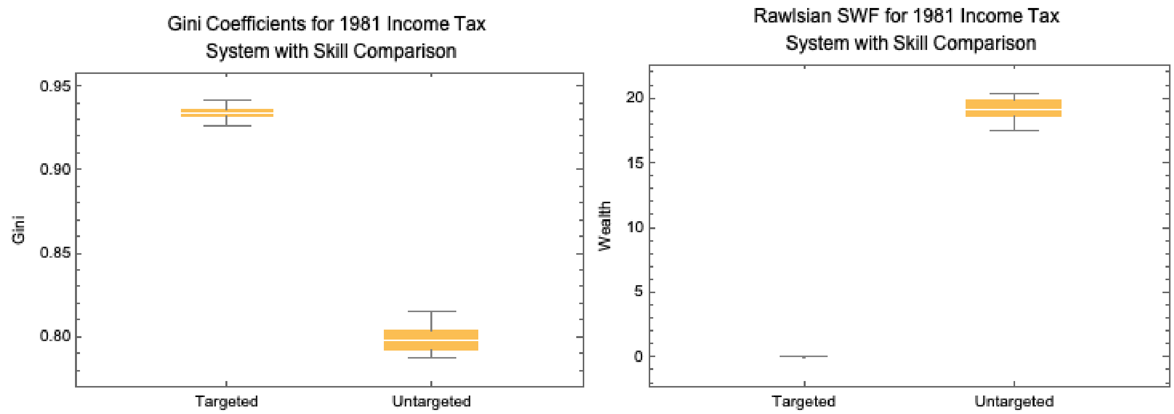

Figure 11: Gini and Rawlsian for Exchange based model with skill comparisons, only income taxes

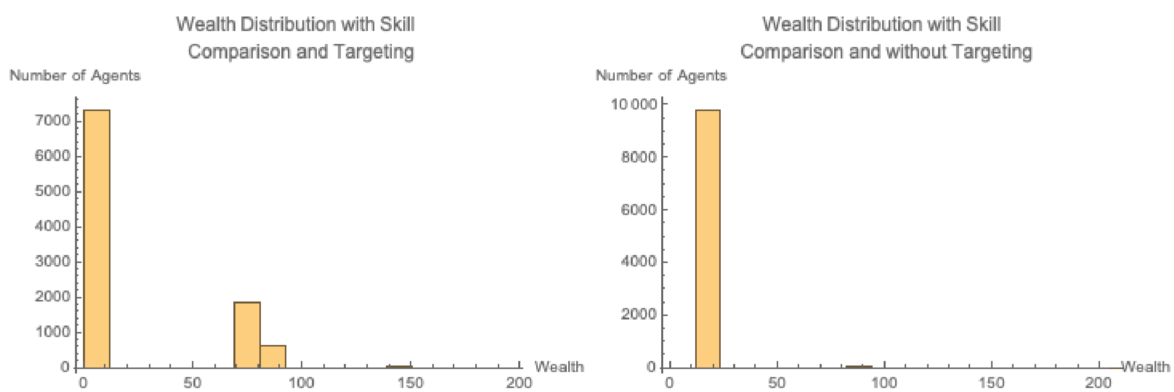

Figure 12: Wealth distribution for basic Exchange based model with skill comparisons, only income taxes

system without targeting as more equal, since most people do have the same amount of money, while in a system with targeting, since only some of the poor receive payouts, two unequal classes are created, indicating more inequality. And since with targeting many poor do not receive payouts, they remain having 0 money, leading to a 0 Rawlsian, while without targeting, all agents receive a payout, so the Rawlsian is essentially just a reflection of this payout.

5 Conclusions

The accumulation model produced expected results: increasing tax rates, decreasing the size of the exempted population, and targeting payouts of collected taxes all in general decreased inequality as measured by our SWFs. Our modified BDY game showed similar results for tax rates and exemptions, but its inherent randomness reversed the effect of targeting payouts. Attempting to fix this by tying income to wealth failed to fix this issue, and introduced significantly more inequality. However, we are looking at the effects of each of these systems over a period of 700 years. This is done to see what the equilibrium state is, but may not be helpful in determining the real world effects of different tax systems. Thus, the next step would be to pivot from long-term models to short-term. This would allow us to see the effects of changing from one tax structure to another within a few years, allowing us to determine what is realistic to expect in such situations, something which seems to be relatively unexplored in this field.

References

[1] C. Gini. *Variabilità e mutabilità*. 1912.

[2] Erlend E Bø, Elin Halvorsen, and Thor O Thoresen. Heterogeneity of the carnegie effect. *Journal of Human resources*, 54(3):726–759, 2019.

[3] Christopher W. Kulp, Michael Kurtz, Nathaniel Wilston, and Luke Quigley. The effect of various tax and redistribution models on the gini coefficient of simple exchange games. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(8):083118, 2019.

[4] A. Dragulescu and V.M. Yakovenko. Statistical mechanics of money. *The European Physical Journal B*, 17(4):723–729, Oct 2000.

Signature:

Email: kulp@lycoming.edu

Signature:

Email: kurtz@lycoming.edu

Signature:

Email: brandon@lycoming.edu

Signature:
DAVID FISHER (May 6, 2022 11:48 EDT)

Email: fisher@lycoming.edu

Signatures for Haberberger Paper

Final Audit Report 2022-05-07

Created:	2022-05-04
By:	Kiera Vinson (Peaceelf2000@gmail.com)
Status:	Signed
Transaction ID:	CBJCHBCAABAAi2hwNwLbLlLd871wVMDgWXX06F7UsKJH

"Signatures for Haberberger Paper" History

Document created by Kiera Vinson (Peaceelf2000@gmail.com)
2022-05-04 - 7:06:09 PM GMT- IP address: 64.189.222.154

Document emailed to CHRISTOPHER KULP (kulp@lycoming.edu) for signature
2022-05-04 - 7:09:09 PM GMT

Document emailed to MICHAEL KURTZ (kurtz@lycoming.edu) for signature
2022-05-04 - 7:09:09 PM GMT

Document emailed to ANDREW BRANDON (brandon@lycoming.edu) for signature
2022-05-04 - 7:09:09 PM GMT

Document emailed to DAVID FISHER (fisher@lycoming.edu) for signature
2022-05-04 - 7:09:10 PM GMT

Email viewed by CHRISTOPHER KULP (kulp@lycoming.edu)
2022-05-04 - 7:10:10 PM GMT- IP address: 71.58.191.98

Document e-signed by CHRISTOPHER KULP (kulp@lycoming.edu)
Signature Date: 2022-05-04 - 7:11:03 PM GMT - Time Source: server- IP address: 71.58.191.98

Email viewed by MICHAEL KURTZ (kurtz@lycoming.edu)
2022-05-04 - 8:22:23 PM GMT- IP address: 73.230.89.135

Document e-signed by MICHAEL KURTZ (kurtz@lycoming.edu)
Signature Date: 2022-05-04 - 8:22:48 PM GMT - Time Source: server- IP address: 73.230.89.135

Email viewed by DAVID FISHER (fisher@lycoming.edu)
2022-05-05 - 4:33:31 PM GMT- IP address: 199.91.148.226

Document e-signed by DAVID FISHER (fisher@lycoming.edu)
Signature Date: 2022-05-06 - 3:48:59 PM GMT - Time Source: server- IP address: 50.206.199.228

Email viewed by ANDREW BRANDON (brandon@lycoming.edu)

2022-05-07 - 0:25:35 AM GMT- IP address: 174.60.101.156

Document e-signed by ANDREW BRANDON (brandon@lycoming.edu)

Signature Date: 2022-05-07 - 0:26:44 AM GMT - Time Source: server- IP address: 174.60.101.156

Agreement completed.

2022-05-07 - 0:26:44 AM GMT