Tax Mechanisms and Gradient Flows

Stefan Steinerberger
(joint work with Aleh Tsyvinski)

WIFO, June 2019

Yale University

The main idea

- Classic Mirrlees static nonlinear income tax problem
  - we still do not have a full solution for the static problem
  - even in the simplest case (iso-elastic, Pareto) Diamond-Saez formula is still a coupled system of differential equations
- Main idea: Study a static optimal tax from the **dynamic point of view**.
  - Starting from a given tax (either optimal or suboptimal): associate with the tax a dynamical system the analysis
- Key object:
  - a **heat equation** and the associated **heat kernel**
  - heat equation is intimately connected to the optimal tax problem
  - allow to derive new properties of tax mechanisms

Main result: optimal tax

- Associate with the first order conditions for the optimal tax a heat kernel:
  - a family of averaging functions at different times (or scales)
- *Main results* – **a fairness principle**:
  - **Optimal tax at a given income** is equal to the **weighted average of the optimal taxes** at other incomes **at any scale** and of a weighted mechanical effect.
- Weighting is given by the heat kernel
  - show that it behaves as a **Gaussian**
  - The weights at different scales are tightly linked – **one unified weighting scheme at any scale**

Fairness principle

- Not externally imposed
  - It is revenue maximizing to treat agents similarly: a tax at a given income equal to the weighted average of taxes around it
- It is a smoothing result
  - a way to smooth deadweight loss and raised revenues across incomes
- Intuition about connection to the heat equation:
  - Nature wants to distribute heat evenly on the background

Main results: Gradient flows

- Construct a gradient flow for any initial (optimal or suboptimal) tax function:
  - changes the underlying tax function in the direction of maximal revenue increase
  - steepest descent in the space of functions
  - The optimal tax is a stationary point of the gradient flow
- Show that a version of the fairness principle holds along the trajectory of the gradient flow, for short time asymptotics
- Gradient flow gives rise to the heat equation
  - and all of its nice properties, e.g., smoothing

Main contributions

- A new way to look at the classical problem:
  - static system from a dynamic point of view
- Fairness principle – a new property of the tax systems
  - at the optimum
  - and of the tax reform
- Unifying a variational (tax reform) and an optimal tax point of view
- Mathematical contribution:
  - heat equation arises in a completely new context

Literature

- Tirole and Guesnerie (1981): tax reform from a gradient descent point of view for linear taxes - ODE
  - our work: nonlinear taxes, PDE
- Sonnenschein (1981, 1982) and Artzner, Simon, and Sonnenschein (1986): derive a heat equation as a gradient process of the firms adjusting the commodity they produce by maximizing the rate of change in profit
- Optimal transport:
  - gradient flow is one of the main objects: Villani (2003), Jordan, Kinderlehrer, and Otto (1998)
  - in (macro)economics, optimal transport is starting to be used: Chiappori, Eeckhout-Kircher, Fajgelbaum-Schaal, Galichon, Lindenlaub, Salanie. A great book by Galichon (2016).

Classic static nonlinear income tax setting

- **Productivity**: $\theta \in \Theta \subset \mathbb{R}_+$, c.d.f. $H(\theta)$, density $h(\theta)$[1]
- **Preferences**: $U(c, l) = c - v(l)$
- **Tax**: $T : \mathbb{R}_+ \to \mathbb{R}$, non-linear function of income $y = \theta \times l$
- **Agent**:

$$\max_{y \geq 0} y - T(y) - v(\frac{y}{\theta})$$

- $y(\theta, T)$ is the argmax, c.d.f. $\Phi(y)$, density $\phi(y)$[2]
- **Government revenue**: $R(T) = \int_\Theta T(y(\theta, T)) dH(\theta)$

[1] Assume zero density at the boundaries.

[2] Assume: (1) a unique global maximum; (2) ϕ_t is continuous and bounded away from zero

Variations of taxes

- Gateaux differential of revenue $R(T) = \int_{\Theta} T(y(\theta, T))dH(\theta)$:

$$\delta R(T, \hat{T}) = \int_{\mathbb{Y}} \hat{T}(y)\,\phi(y)\,dy + \int_{\mathbb{Y}} T'(y)\,\delta y(\theta)\,\phi(y)\,dy$$

- Gateaux differential of income[3]:

$$\delta y(\theta) = -\varepsilon(y(\theta))\,\hat{T}'(y(\theta))$$

Lemma

The Gateaux differential of revenue in the direction $\hat{T}$ is:

$$\delta R(T, \hat{T}) = \int_{\mathbb{Y}} \hat{T}(y)\,\phi(y)\,dy - \int_{\mathbb{Y}} T'(y)\,\varepsilon(y(\theta))\,\hat{T}'(y(\theta))\phi(y)\,dy$$

[3] $\varepsilon(\theta) = \frac{y(\theta)}{(1-T'(y(\theta)))}\frac{e(\theta)}{1+p(y(\theta))e(\theta)}$ is the elasticity along the nonlinear budget constraint

Optimal tax

$$\delta R(T,\hat{T}) = \int_{\mathbb{Y}} \hat{T}(y)\,\phi(y)\,dy - \int_{\mathbb{Y}} T'(y)\,\varepsilon(y(\theta))\,\hat{T}'(y(\theta))\phi(y)\,dy$$

- Integrate by parts

$$\delta R(T,\hat{T}) = \int_{\mathbb{Y}} \hat{T}(y)\,\phi(y)\,dy + \int_{\mathbb{Y}} \hat{T}(y)\,\frac{d}{dy}\left(T'(y)\,\varepsilon(y)\,\phi(y)\right)dy$$

- Optimal tax T_* satisfies (Mirrlees 1971, Diamond 1998, Saez 2001):

$$\boxed{0 = \phi_*(y) + \frac{d}{dy}\left(T_*'(y)\,\varepsilon_*(y)\,\phi_*(y)\right)}$$

Optimal tax is not in a closed form

$$0 = \phi_*(y) + \frac{d}{dy}\left(T'_*(y)\,\varepsilon_*(y)\,\phi_*(y)\right)$$

- **Saez**:

$$T'_*(y) = \frac{1}{\varepsilon_*(y)} \frac{1 - \Phi_*(y)}{\phi_*(y)}$$

  - $\varepsilon_*(y)$, $\phi_*(y)$ are endogenous objects

Optimal tax is not in a closed form

$$0 = \phi_* (y) + \frac{d}{dy} \left(T'_* (y) \varepsilon_* (y) \phi_* (y) \right)$$

- **Diamond** (assume isoelastic):

$$\frac{T'_* (y (\theta))}{1 - T'_* (y (\theta))} = \left(1 + \frac{1}{\varepsilon}\right) \frac{1 - F (\theta)}{\theta f (\theta)}$$

  - rhs is a closed form as a function of θ
  - lhs is the marginal tax rate paid faced a type θ
  - lhs is not a closed form expression for the *income tax schedule* $T(\cdot)$ as we do not know y as a function of θ
  - **this is the essence of the Mirrlees problem**

Plan: optimal tax

- Consider the static problem from the dynamic point of view
  - Associate a dynamic object: a heat kernel
- Derive a new characterization of the optimum:
  - a fairness principle (an invariance relationship for the optimum)

Heat kernel and the Operator

$$0 = \phi_* (y) + \frac{d}{dy} \left(T_*' (y) \varepsilon_* (y) \phi_* (y) \right)$$

- Second order differential operator $L = \frac{\partial}{\partial y} \left(\varepsilon_* (y) \phi_* (y) \frac{\partial}{\partial y} \right)$:

$$L T_* = -\phi_*$$

- Associate the heat kernel $q_t (x, y)$ given by Kolmogorov forward equation:

$$\frac{\partial}{\partial t} q_t (x, y) = \frac{\partial}{\partial y} \left(\varepsilon_* (y) \phi_* (y) \frac{\partial}{\partial y} q_t (x, y) \right)$$

$$\lim_{t \to 0} q_t (x, y) = \delta (x - y)$$

Invariance relationship: proof

$$\frac{\partial}{\partial t}\int q_t(x,y)\,T_*(y)\,dy = \int \frac{\partial}{\partial t} q_t(x,y)\,T_*(y)\,dy$$

$$=_{[\textit{use KFE}]} \int \frac{\partial}{\partial y}\left(\varepsilon_*(y)\,\phi_*(y)\,\frac{\partial}{\partial y} q_t(x,y)\right) T_*(y)\,dy$$

$$=_{[\textit{integrate twice by parts}]} \int q_t(x,y)\frac{\partial}{\partial y}\left(\varepsilon_*(y)\,\phi_*(y)\,\frac{\partial}{\partial y} T_*(y)\right)$$

$$=_{[\textit{use } LT_* = -\phi_*]} \; -\int q_t(x,y)\,\phi_*(y)\,dy.$$

$$\int q_t(x,y)\,T_*(y)\,dy = T_*(x) - \int_0^t \int q_s(x,y)\,\phi_*(y)\,dyds$$

Fairness principle

Proposition

Fairness principle for the optimum. *The optimal tax $T_*(y)$ is invariant under the heat kernel $q_t(x,y)$ given by, for any $x \in Y$ and any $t > 0$:*

$$T_*(x) = \int_0^t \int q_s(x,y)\,\phi_*(y)\,dyds + \int q_t(x,y)\,T_*(y)\,dy$$

- $q_t(x,y)$ is an averaging function at scale t (very similar to Gaussian)
- An optimal tax $T_*(x)$ at any income x is equal to a (weighted by q_t) average of the taxes at all other incomes $T_*(y)$ **at any (!) scale** t
  - plus a mechanical effect

Why fairness?

- A planner wants to raise revenue:
  - no notion of treating people equally in any way
- The most efficient way to do it is such that
  - an agent at a given income x is paying roughly the average of the amount of taxes paid by people working just more or less than him
- More precisely: smoothing of the behavioral effect of taxes (distortions) and raised revenues at any scale t

What is $q_t(x,y)$?

Proposition

The heat kernel satisfies a Gaussian upper bound, ***for any*** $t > 0$*:*

$$q_t(x,y) \le \frac{c_1}{\sqrt{t}} \exp(c_2 t) \exp\left(-c_3 \frac{(x-y)^2}{t}\right),$$

for some positive constants c_1, c_2, c_3.

For $t \to 0$,

$$q_t(x,y) \sim \frac{1}{\sqrt{4\pi\sigma_*(x)\,t}} \exp\left(-\frac{(y - x - \sigma_*'(x)\,t)^2}{4\sigma_*(x)\,t}\right),$$

where $\sigma_*(x) = \varepsilon_*(x)\,\phi_*(x)$

Semigroup property

Proposition

For all points x, y and all times $t, s > 0$

$$q_{t+s}(x,y) = \int q_t(x,z) q_s(z,y) dz.$$

- A fact of crucial importance:
  - the behavior of q_t is tightly linked to both past and future behavior of the heat kernel and is thus far from arbitrary
- **One weighting scheme at any scale!**

Marginal taxes

$$T_*(x) = \int_0^t \int q_s(x,y)\,\phi_*(y)\,dyds + \int q_t(x,y)\,T_*(y)\,dy$$

Corollary

The optimal marginal tax satisfies, for all $t > 0$

$$\frac{\partial}{\partial x} T_*(x) =$$

$$\int_0^t \int \left(\frac{\partial}{\partial x} q_s(x,y) \right) \phi_*(y)\,dyds + \int \left(\frac{\partial}{\partial x} q_t(x,y) \right) T_*(y)\,dy$$

- The marginal tax is determined by the **global** behavior of the optimal tax code
  - a priori not obvious: understanding of the level does not necessarily translate to understanding the derivative
  - trivial for $t \to 0$ (def of a derivative), but we show **for all** t

What is $\frac{\partial}{\partial x} q_s(x, y)$?

- Has a positive part and a negative part with the same total mass
- **Marginal tax** = weighted taxes paid by individuals with income below **minus** weighted taxes paid by individuals with income above
- Similar formulas for the higher derivatives

Key to the results

- View the static optimal tax from the dynamic point of view
  - by associating the heat kernel $q_t(x, y)$ with the operator L determining the optimal tax
- The optimal tax $T_*(x)$ is of course time-independent
- Yet, can think of it as being invariant under a dynamic system that starts from this tax and applies the heat kernel $q_t(x, y)$ to it
  - t can be time or scale

Summary: optimal tax

- Static optimal tax can be studied from the dynamic point of view
  - key object: a heat kernel
- New characterization: optimal tax satisfies an invariance relationship, the fairness principle
  - optimal tax is determined by the weighted average at all scales of optimal taxes at all other incomes
  - weighting behaves like a Gaussian and is tightly linked at all scales
  - behavioral effect and raised revenues are smoothed at all scales
  - applies to the higher derivatives
- Not a closed form solution but a very different view on the classic problem

Gradient flows

- Construct a dynamic system, **a gradient flow**,
  - starts at any (optimal or suboptimal) tax function
  - changes the tax system in the direction of the steepest increase in revenues.
  - The optimal tax is a stationary point of this system.
- Show that a version of the fairness principle holds for the gradient flow
- Gradient flow is a heat equation
  - all kinds of nice properties

Variations of taxes

- Start with any initial tax system T_t
- Change in revenue from a perturbation

$$\delta R(T_t, \hat{T}_t) = \int_{\mathbb{Y}} \left\{ \phi_t + \frac{\partial}{\partial y}[T_t' \varepsilon_t \phi_t] \right\} \hat{T}_t(y)\, d$$

- The **gradient flow**

$$\frac{\partial T_t(y)}{\partial t} = \phi_t(y) + \frac{\partial}{\partial y}[T_t'(y)\, \varepsilon_t(y)\, \phi_t(y)]$$

  - maximizes the change in revenue in L^2 norm over infinitesimally short period of time
  - a continuous version of the steepest descent in a functional space
  - evaluated at the current tax system: T_t, ε_t, ϕ_t are given.

Fairness principle for the gradient flow

Proposition

*(**Fairness principle for the gradient flow**). For small $\tilde{t}$, the tax $T_{\tilde{t}}(y)$, generated by the gradient flow, is given by a weighted Gaussian average of the incomes:*

$$T_{\tilde{t}}(x) \sim (\tilde{t} - t)\,\phi_t(x) + \int_{\underline{y}}^{\bar{y}} \frac{1}{\sqrt{4\pi\sigma_t(x)(\tilde{t} - t)}} e^{-\frac{(y - x - \sigma_t'(x)(\tilde{t} - t))^2}{4\sigma_t(x)(\tilde{t} - t)}} T_t(y)dy$$

- **Key idea of the proof**: for the short $\tilde{t}$, the coefficients $\varepsilon_t(y)$ and $\phi_t(y)$ are essentially "frozen": get the short time asymptotics of the Brownian motion.
- Extends the notion of fairness to the trajectory of the gradient flow of taxes
- Everything in the closed form but unlike the optimum holds for short $\tilde{t}$.
  - of course, also holds for the optimum

Operator Splitting

- Gradient flow:
  - changes the tax schedule in favor of increasing government revenue,
  - letting ϕ_t and ε_t being endogenously driven by T_t
  - density and elasticity change in response to the evolution of taxes.
- Propose to evolve the system separately:
  - fix the distribution of incomes ϕ_t and the elasticity ε_t for a short period of time δt,
  - evolve T_t,
  - then re-compute $\phi_{t+\delta t}$ and $\varepsilon_{t+\delta t}$ based on the new tax function $T_{t+\delta t}$
- Government evaluates the changes in revenues under the current information given by the **exogenous sufficient statistics** evaluated at a given initial time.

Heat equation

- Fixing ϕ_t and ε_t, the gradient flow

$$\frac{\partial T_t(y)}{\partial t} = \phi_t(y) + \frac{\partial}{\partial y}[T_t'(y)\,\varepsilon_t(y)\,\phi_t(y)]$$

  - is a heat equation, i.e., a second order pde of the form

$$\frac{\partial T}{\partial t} = \phi(y) + \frac{\partial}{\partial y}[\sigma(y)\,\frac{\partial T(y,t)}{\partial y}]$$

- One of the most well-known and well-behaved partial differential equation
  - obtain a number of textbook properties

Why heat equation?: loose intuition

- Heat tries to escape the larger is the gradient of temperature around the point
- Agents:
  - if the tax changes by $\frac{\partial T_t(y_t(\theta))}{\partial t}$
  - the agents move their income by $\frac{\partial y_t}{\partial t}(\theta) = -\varepsilon_t(\theta) \frac{\partial T'_t(y_t(\theta))}{\partial t}$.
  - individual θ adjusts his income in the opposite direction and proportionally to the change in the marginal tax rate that he faces
- If the tax is high on me compared to agents nearby, I will "escape" to a lower income

Reference point

- The initial tax system can be rather arbitrary
- For $\phi(y)$ and $\sigma(y)$ fixed, the solution will converge to unique fixed point $\tau(y)$ as $t \to \infty$:

$$0 = \phi(y) + \frac{\partial}{\partial y}[\sigma(y) \frac{\partial \tau(y)}{\partial y}]$$

- **Important**: we are not interested in letting the gradient flow for all $t \to \infty$ and thus finding τ but only evolving it for a very small time t (so as to respect the operator splitting).
- But, τ serves as a proper reference point to describe various smoothing properties of the gradient flow.

Smoothing properties I

Proposition

The gradient flow of the revenue functional coincides with the gradient flow of the functional $\mathcal{J}$ given by the weighted Sobolev-type seminorm H^1

$$\mathcal{J}(T) = \frac{1}{2} \int \sigma(y) \left(T'(y) - \tau'(y) \right)^2 dy$$

- This gradient flow is trying to smooth out rough irregularities in the difference between T and τ
  - a large value of $|T'(y) - \tau'(y)|$ implies to existence of a large value of $(T'(y) - \tau'(y))^2$
  - the flow is trying to decrease this as quickly as possible.
- $\sigma(y)$ serves as a natural weighting measure:
  - if $\sigma(y)$ is large, then irregularities in that region count even more severely and are dampened quicker.

Smoothing properties II

Proposition

Let $T(t, y)$ denote the solution of the heat equation associted with the gradient flow of the tax revenue functional. If $\sigma(y)$ is bounded away from 0 and $T(0, y)$ is bounded, then $T(t, y)$ is infinitely differentiable for any $t > 0$.

- Tax reform viewed as the gradient flow leads to the continuous tax systems.
- Moreover, the gradient flow has the effect of mollifying any tax scheme instantaneously.
  - If $T(y)$ has large amounts of strong oscillations or maybe even discontinuous jumps, then the gradient flow acts strongest on those parts first

Smoothing properties: Sturm-Liouville

Proposition

Let $T(t,y)$ denote the solution of the heat equation associted with the gradient flow of the tax revenue functional, $T(0,y)$ be an arbitrary initial tax schedule), and λ_1 be the first eigenvalue of the associated Sturm-Liouville operator. Then, $\forall t > 0$:

$$\int_a^b (T(t,y) - \tau(y))^2 dy \leq e^{-2\lambda_1 t} \int_a^b (T(0,y) - \tau(y))^2 dy.$$

- The gradient flow smoothes a measure of variability of the tax schedule, the squared deviation from the limiting stationary solution.
  - Moreover, such smoothing is exponential.
- This result is not about the convergence to the optimal tax T_* but rather about the smoothing behavior of the gradient flow at each step of the operator splitting.

Global convergence

- Outside of the scope of this paper:
  - would be a nice Ph.D. in math thesis
- But what we have is a classical operator splitting
  - similar to splitting convection from diffusion
  - which pretty much converges under very broad circumstances Glowinski, Osher, and Yin (2016)
- Results on the fairness principle do not need operator splitting.

Conclusion and extensions

- Dynamical systems can provide new insights on the analysis of static optimal taxation problems
  - Fairness principle for the optimum and for the trajectory
  - Evolution of taxes is a heat equation: unifying variational and the optimal tax approach
- Extensions (immediate):
  - Maximizing social welfare
  - Extensive margin
  - Non quasi-linear preferences
- Multidimensional types?
- More broadly: nonlinear pricing, matching, etc.