

Tax Mechanisms and Gradient Flows

Stefan Steinerberger*

Yale University

Aleh Tsyvinski

Yale University[†]

April 30, 2019

Abstract

We demonstrate how a static optimal income taxation problem can be analyzed using dynamical methods. Specifically, we show that the taxation problem is intimately connected to the heat equation. Our first result is a new property of the optimal tax which we call the fairness principle. The optimal tax at any income is invariant under a family of properly adjusted Gaussian averages (the heat kernel) of the optimal taxes at other incomes. That is, the optimal tax at a given income is equal to the average of optimal taxes at other incomes and income densities, where the weights are given by the heat kernel. Moreover, this averaging happens at every scale tightly linked to each other, providing a unified weighting scheme at all income ranges. The fairness principle arises not due to equality considerations but rather it represents an efficient way to smooth the burden of taxes and generated revenues across incomes. Just as nature wants to distribute heat evenly, the optimal way for a government to raise revenues is to distribute the tax burden and raised revenues evenly among individuals. We then construct a gradient flow of taxes – a dynamic process changing the existing tax system in the direction of the increase in tax revenues – and show that it takes the form of a heat equation. The fairness principle holds also for the short-term asymptotics of the gradient flow, where the averaging is done over the current taxes. The gradient flow we consider can be viewed as a continuous process of a reform of the nonlinear income tax schedule and thus unifies the variational approach to taxation and optimal taxation. We present several other characteristics of the gradient flow focusing on its smoothing properties.

*S.S. is supported by the NSF (DMS-1763179) and the Alfred P. Sloan Foundation.

[†]We thank Manuel Amador, Andy Atkeson, Felix Bierbrauer, Patrick Bolton, Jaroslav Borovicka, Pierre Boyer, Hector Chade, Alfred Galichon, Martin Hellwig, Oleg Itskhoki, Felix Kubler, Ilse Lindenlaub, Benny Moldovanu, Chris Phelan, Georgii Riabov, Larry Samuelson, Florian Scheuer, Philipp Strack, Kjetil Storesletten, Juuso Valimaki, and, especially, Nicolas Werquin.

Introduction

The classic static optimal nonlinear income taxation problem (Mirrlees 1971) is one of the prototypical examples of mechanism design problems.¹ Diamond (1998) and Saez (2001) derive expressions for the optimal taxes that are widely used in public finance. Yet, even in the simplest cases, those do not provide closed-form solutions or a full characterization for the optimal income tax. The solution is a coupled system of differential equations and has relatively few known characteristics.²

A new aspect of our analysis is in the application of dynamical methods to static mechanism design problems. Starting from a given tax (either optimal or suboptimal), we associate with it a dynamical system – the analysis of which yields a new view of the problem. This dynamical system is tightly connected to a heat equation, one of the most well-studied and well-behaved partial differential equations.

Our main result shows that the optimal tax satisfies a certain invariance relationship which we call a fairness principle. The fairness principle states that an optimal tax at a given income is equal to the properly weighted average of the taxes at all other incomes. We call it “fair” not because the planner has a social welfare function that places a weight on equality but rather that it is efficient to treat agents at various incomes similarly. In this sense, fairness is a different version of smoothing the burden of the deadweight losses of taxes and raising revenues.

We start by analyzing the static optimal nonlinear income tax. We associate with the equation determining the optimal tax a dynamical object – the heat kernel which determines a family of averaging functions at different times (or scales). We show that the optimal tax is invariant under the heat kernel. Specifically, the optimal tax at any given income is determined as a weighted average of the optimal taxes at other incomes and of income densities. Importantly, we show that this averaging is done at *all* times (scales). What we mean by this is that as the heat kernel is spreading wider (with, say, time), the optimal tax is equal to the average tax at the wider sets of incomes and income densities: from the weight placed on incomes just around the point to more and more incomes. In other words, the planner at the optimum ensures that the deadweight loss or the behavioral effect of taxes is distributed fairly – as the weighted average of taxes at all scales. This behavioral effect is added to the statutory impact of taxes or the mechanical effect which is also a weighted average over all scales of income densities. A fact of crucial importance is that the fairness principles for different scales are tightly linked via the so called semigroup property and is thus far from arbitrary. That is, there is one unified weighting scheme across all of the income scales. We then show that the behavior of the heat kernel is essentially Gaussian. For any time (scale), we provide a Gaussian upper bound that shows that the heat kernel decays as the Gaussian. For the small time asymptotics, the heat kernel is exactly a Gaussian scaled by the density and the elasticity at the given income at the

¹See, e.g., the book by Salanie (2011).

²See, e.g., Hellwig (2010) for the state-of-the art analysis of incentive problems with unidimensional characteristics.

optimum. A corollary of the fairness principle is a formula determining the optimal marginal tax by the global behavior of the level of the optimal tax code. These results are new to the taxation and mechanism design literatures.

Second, we construct a gradient flow for any initial (optimal or suboptimal) tax function. The gradient flow is a dynamical system that changes the underlying tax function in the direction of the steepest descent of the tax revenue functional. We show that the fairness principle also holds along the trajectory of the gradient flow. Specifically, we prove that the small time behavior of the gradient flow (that is, the short-time asymptotics) is such that the evolved tax is equal to the Gaussian average of the current taxes at all incomes and income densities, where the Gaussian is adjusted by the density and the elasticity at the current tax system. The optimal tax is a stationary point of the gradient flow. That is, starting from the optimum both the averaging and the agents' behavior stays constant.

We then propose to use the operator-splitting methods corresponding to the changes in the tax revenues and to the agent's behavior and show that the gradient flow is a heat equation within each step of operator splitting. The operator splitting technique has a natural economic meaning. On the trajectory of the tax reform, the government, for an infinitesimally short time, evolves taxes in the direction of increased revenues, keeping the density of agents' incomes and elasticities fixed at their value observed in the current economy. Then the densities and the elasticities are recomputed. That is, the government evaluates the changes in revenues under the current information given by the exogenous sufficient statistics evaluated at a given time.³ Within each step of this process, the gradient flow of tax revenues is a heat equation. Having identified the heat equation as the principal object governing the evolution of the tax function, there is a wealth of additional results one can obtain to study its behavior. We primarily focus on showing that the gradient flow possesses a very nice underlying smoothing structure. First, we show that the gradient flow is trying to smooth out the rough irregularities in the tax schedule. That is, it acts first in the income regions which have the largest deviation of the marginal tax from the properly defined reference point. These deviations are weighted by the natural weighing measure that depends on the labor income elasticity and the density of the incomes. Second, we show that the reform viewed as the gradient flow leads to the continuous tax systems, mollifying any tax scheme instantaneously. Third, we use the Sturm-Liouville theory to describe further characteristics of the gradient flow and, in particular, to show exponential decay of the variability of the tax system compared to the reference point.

The fact that the gradient flow gives rise to a law of motion for the tax function characterized by a parabolic PDE (namely, a heat equation) is subtle and is at the essence of the economics of the problem. The well-known and much studied heat equation arises in a new, and possibly surprising, setting: in an environment where a government is interested in collecting income tax revenue and the agents are heterogeneous in their skills. In a broad sense, the agents and the government have opposing objectives: to maximize revenue for the government, and to

³See, e.g., Kleven (2018) for a recent review of the sufficient statistics approach in taxation.

maximize the utility for the agents, given taxes. We show, among other things, that a local process where the government tries to maximize its income via taxation leads, in a natural way, to the heat equation. Just as nature wants to distribute heat evenly on a background (leading to the spreading of heat), the optimal way for a government to raise revenues is to distribute the tax burden and raised revenues evenly among individual subjects – this is perhaps particularly surprising. This fairness is indeed a phenomenon coming from agents adapting to uneven taxation systems by working more or less, respectively. We believe this to be a new way of arriving at the heat equation and a new tax and revenue smoothing result in taxation.

The construction of the dynamical system that we propose also contributes to unifying the optimal and the variational approach to taxation. The variational approach considers a potentially suboptimal tax and proceeds with varying it locally to derive the formulas for the effects of the tax reforms.⁴ When considering the effects of the local tax reform, a natural question to ask is whether there is a process of reforms that may lead to the optimum. Tirole and Guesnerie (1981) construct such a process for linear taxes based on gradient projections, leading to an ordinary differential equation. The environment with nonlinear taxes is significantly more challenging as now the whole tax function is evolved as opposed to just one linear tax – we therefore have the steepest descent path in a space of functions. We have constructed such a process that leads, within each step, to a well-behaved partial differential equation. One advantage of this process is that it uses, similarly to the variational approach, only the current information about the economy, such as elasticities and the density of incomes. The second advantage is that it gives rise to one of the most regular and well-behaved mathematical objects – the heat equation – allowing for a wealth of characterizations. Additionally, our analysis shows that continuous version of the iterative fixed point method commonly used for computing optimal taxes (see, e.g., Brewer, Saez, and Shepard 2010) can be represented within each step as a heat equation.

Our analysis of the static taxation problem from the dynamical point of view, the construction of the gradient flow, and the analysis of the optimum and the gradient flow from the dynamical perspective are new to the mechanism design and optimal taxation literatures.

More broadly, there are several papers that are connected to our work. Bolton and Harris (2010) is perhaps the most comparable in terms of the approach. They associate a dynamic risk sharing rule with that of the static problem. Furthermore, Bolton and Harris (2010) obtain an elegant asymptotic expansion of the dynamic problem around a myopic optimum showing how the static problem is modified by the dynamic correction terms.

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

⁴See, e.g., Saez (2001), Kleven and Kreiner (2006), and Golosov, Tsyvinski, and Werquin (2014) for the methodology; Kleven, Kreiner, and Saez (2009) and Jacquet and Lehmann (2015) for the analysis of the multidimensional types; Saez and Stantcheva (2016) and Bierbrauer and Boyer (2018) for the political economy context; Sachs, Tsyvinski, and Werquin (2016) and Scheuer and Werning (2016) for the analysis in general equilibrium; and Saez and Stantcheva (2018) for capital income taxation.

rate of change in profit subject to a quadratic cost of adjustment. McCann (2014) argues that this result is a precursor to some of the results on the gradient flows in the optimal transport literature.⁵

Some of the techniques that we use have parallels in the optimal transport literature (see, e.g., Villani (2003)) in which there is a renewal of interest in economics (see, e.g., early work of Chiappori, McCann, and Nesheim (2010), a comprehensive book by Galichon (2016), or a review in the context of matching models by Chiappori and Salanie (2016)). One important result in the optimal transport literature is showing that a Fokker-Planck equation arises as a gradient flow in a Wasserstein space (Jordan, Kinderlehrer, and Otto (1998)). In mechanism design, optimal transport has been recently used in the context of multi-dimensional screening by Figalli, Kim, and McCann (2011) and Daskalakis, Deckelbaum, and Tzamos (2017).

Our analysis is interesting also from the purely mathematical point of view. The heat equation is a classical object in mathematics and physics (indeed, so classical that Fourier’s specific way of solving it via trigonometric series is the origin of harmonic analysis, an entire subfield of mathematics). Its importance suggests that it should indeed appear in a wide variety of settings. Here, we present a new setting in which the heat equation naturally arises.

The paper is structured over two main parts that while connected are self-contained. Section 3 deals with the optimal nonlinear tax. Section 4 deals with the gradient flow of taxes.

1 Environment

For clarity of exposition we present the simplest economic environment of tax mechanisms. The model straightforwardly generalizes to a variety of other, more sophisticated environments – we discuss these extensions in the text and in the conclusion.

1.1 Individuals

Agents are characterized by an exogenous and fixed productivity type $\theta \in \Theta \subset \mathbb{R}_+$. Preferences over consumption c and labor effort l are represented by the utility function $U(c, l) = c - v(l)$, where the disutility of labor effort v is twice continuously differentiable, increasing and strictly convex. The government levies a tax liability $T : \mathbb{R}_+ \rightarrow \mathbb{R}$ which can be an arbitrarily non-linear function of the individual’s labor income $y = \theta \times l$. The agent’s budget constraint is $c = \theta l - T(\theta l)$.

The optimization problem of an individual with type θ reads:

$$\max_{l \geq 0} \theta l - T(\theta l) - v(l). \tag{1}$$

We denote by $l(\theta, T) \in \mathbb{R}_+$ the argmax of this problem and by $y(\theta, T) \in \mathbb{R}_+$ the agent’s labor

⁵See also Blanchet and Carlier (2015).

income. For ease of notation, when there is no ambiguity we remove the argument T from these variables and write them as $l(\theta)$, $y(\theta)$.

Characterization of individual behavior. Assuming that the tax function T is continuously differentiable, labor income $y(\theta) \equiv y(\theta, T)$ is characterized by the first-order condition:

$$1 - T'(y(\theta)) = v' \left(\frac{y(\theta)}{\theta} \right) \frac{1}{\theta} . \quad (2)$$

We assume that no individual θ is indifferent between two or more incomes in the initial equilibrium: for all θ , the individual problem (1) has a unique global maximum given the tax system T . It is straightforward then to show that there is a one-to-one map between productivity types θ and pre-tax incomes $y(\theta)$.

Productivity and income distributions. We denote by $H(\theta)$ the c.d.f. of $\theta \in \Theta$, and by $h(\theta)$ the corresponding density function. We assume that the set Θ is a compact interval of \mathbb{R}_+ , and that the density of types h is equal to zero at the boundaries of Θ . We also denote by $\Phi(y)$ and $\phi(y)$ the c.d.f. and the p.d.f. of incomes $y \in \mathbb{Y} \subset \mathbb{R}_+$. We assume that the density of incomes ϕ is continuous and bounded away from zero on any finite interval $[\underline{y}, \bar{y}] \subset \mathbb{Y}$ with $\underline{y} > 0$.⁶

1.2 Government

We define government revenue by

$$R(T) = \int_{\Theta} T(y(\theta, T)) dH(\theta) . \quad (3)$$

It is equal to the sum over agents $\theta \in \Theta$ of the tax liability on their income $y(\theta, T)$, taking into account their optimizing behavior given the tax system T .⁷

We denote by

$$p(y) = - \frac{\partial \ln(1 - T'(y))}{\partial \ln y} = \frac{yT''(y)}{1 - T'(y)}$$

the local rate of progressivity of the tax schedule. It is equal to (minus) the elasticity of the retention rate $1 - T'(y)$ with respect to income y (see, e.g., Musgrave and Thin (1948)).

⁶Our results can be straightforwardly generalized to the case of types and incomes in the whole space \mathbb{R}_+ by using an increasing sequence of compact sets $\Theta, \mathbb{Y} \subset \mathbb{R}_+$.

⁷It is immediate to extend our results to the case of the government maximizing a social welfare function.

1.3 Taxable income elasticity

The structural labor supply elasticity of agent θ with respect to the retention rate $r(\theta) = 1 - T'(y(\theta))$ is defined by

$$e(\theta) = \frac{\partial \ln y(\theta)}{\partial \ln r(\theta)} = \frac{v'(l(\theta))}{l(\theta)v''(l(\theta))}.$$

We define the compensated labor income response along the non-linear budget constraint by

$$\varepsilon(\theta) = \frac{y(\theta)}{r(\theta)} \frac{e(\theta)}{1 + p(y(\theta))e(\theta)}. \quad (4)$$

Assuming that there is a one-to-one map between types θ and incomes $y(\theta)$, we can equivalently denote the labor supply responses $\varepsilon(\theta)$ by $\varepsilon(y)$. In the sequel we use both notations interchangeably. We assume that $\varepsilon(y)$ is continuous and bounded away from zero on any subinterval $[y, \bar{y}] \subsetneq \mathbb{Y}$.

The variable $\varepsilon(\theta)$ is equal to the partial derivative of labor income $y(\theta)$ with respect to the retention rate $r(\theta)$, normalized by the term $1 + p(y(\theta))e(\theta)$. Intuitively, this term accounts for the fact that the agent's labor income adjustment in response to an exogenous increase in his retention rate $r(\theta)$ leads to an endogenous shift in his tax rate by $p(y(\theta))$, which in turn causes a further labor income adjustment $e(\theta)$. Solving for the fixed point leads to the total response (4).

2 Tax reforms and optimum taxation

In this section we define a notion of local tax reforms, and derive their effects on individual behavior and government revenue. We then derive a formula for the optimal tax. All of the results in this section are standard in the literature.

2.1 Variations of taxes

We define a direction of reform of the tax function T as a continuously differentiable function $\hat{T} : \mathbb{R}_+ \rightarrow \mathbb{R}$. The perturbed tax function is then $T + \mu\hat{T}$, where $\mu > 0$ is the size of the reform in the direction \hat{T} . In this section, we derive the first-order change in government revenue $R(T)$ in response to the tax reform $\mu\hat{T}$ as $\mu \rightarrow 0$. That is, we compute the Gateaux differential of the functional $T \mapsto R(T)$, formally defined by:

$$\delta R(T, \hat{T}) \equiv \lim_{\mu \rightarrow 0} \frac{R(T + \mu\hat{T}) - R(T)}{\mu}.$$

We first describe the impact of a tax reform \hat{T} of the tax schedule T on the government revenue.

Lemma 1. *The Gateaux differential of government revenue in the direction \hat{T} is given by:*

$$\delta R(T, \hat{T}) = \int_{\mathbb{Y}} \hat{T}(y) \phi(y) dy - \int_{\mathbb{Y}} T'(y) \varepsilon(y) \hat{T}'(y) \phi(y) dy. \quad (5)$$

Proof. The Gateaux differential of the government revenue functional in the direction \hat{T} is given by:

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

where the differential of the agent's income $\delta y(\theta)$ is given by the Gateaux differential of the agent's first order condition (2):

$$-T''(y(\theta))\delta y(\theta) - \hat{T}'(y(\theta)) = v''\left(\frac{y(\theta)}{\theta}\right) \frac{1}{\theta^2} \delta y(\theta).$$

Solving for δy and using the definition of elasticity $\varepsilon(y)$ gives $\delta y(\theta) = -\hat{T}'(y(\theta))\varepsilon(y(\theta))$. \square

Equation (5) shows that the first-order effect of the tax reform \hat{T} on government revenue, $\delta R(T, \hat{T})$, is given by the sum of two terms. The first integral on the right hand side is the statutory impact of the reform, i.e., the amount of revenue raised mechanically by changing the tax payment of each agent with income y by $\hat{T}(y)$. This term simply sums these additional tax payments over the whole population, using the density of incomes $\phi(\cdot)$, assuming that everyone's behavior remains unchanged following the reform. The second integral is the excess burden, or the deadweight loss, of the tax reform. Specifically, consider the agents who earned income y in the initial equilibrium, i.e., before the tax reform. An increase in their marginal tax rate by $\hat{T}'(y)$ lowers the labor income of these agents by $\varepsilon(y)$, by definition of their labor income response along the non-linear initial tax schedule. This in turn reduces government revenue by a fraction $T'(y)$ of this income loss. Summing these effects over the population using the density of incomes $\phi(\cdot)$ leads to expression (5).

Note that this expression gives the effects of tax reforms $\hat{T}(y)$ in closed-form since all of the variables (the tax schedule $T(y)$, the labor supply elasticities $\varepsilon(y)$, the density of incomes $\phi(y)$) are observed or can be estimated empirically in the economy with the given tax $T(y)$.

2.2 Optimal Tax

Lemma 1 provides a formula for the revenue effects of any tax reform \hat{T} in the economy starting from any, optimal or suboptimal, tax schedule T .

We now obtain a characterization of the optimal (i.e., government revenue maximizing) tax schedule T_* by imposing that no tax reform has a non-zero first-order effect on government revenue, i.e., $\delta R(T_*, \hat{T}) = 0$ for all $\hat{T} : \mathbb{R}_+ \rightarrow \mathbb{R}$. Let $\varepsilon_*(y)$, $\phi_*(y)$, and $\Phi_*(y)$ denote the compensated labor income response and the p.d.f. and c.d.f. of incomes given the tax schedule

T_* .

Integrating by parts the Gateaux differential of the government revenue functional (5)

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

and setting it to zero for any $\hat{T}(y)$ yields:

$$0 = \phi_*(y) + \frac{d}{dy} (T'_*(y) \varepsilon_*(y) \phi_*(y)). \quad (6)$$

Integrating with respect to y yields the the following Corollary (due to Mirrlees 1971, Diamond 1998, Saez 2001).

Corollary 1. *The optimal tax schedule T_* satisfies: for all $y \in \mathbb{Y}$,*

$$T'_*(y) = \frac{1}{\varepsilon_*(y)} \frac{1 - \Phi_*(y)}{\phi_*(y)}, \quad (7)$$

Formula (7) shows that the optimal marginal tax rate at income level y is the product of two terms. First, it is proportional to the inverse compensated labor income response to tax rates at income y , $1/\varepsilon_*(y)$: the higher the disincentive effect of marginal tax rates, the lower the optimal tax rate. The second term is related to the hazard rate of the income distribution, $(1 - \Phi_*(y))/(\phi_*(y))$ and is a benefit-cost ratio that measures the fraction of agents whose tax liability increases by a lump-sum amount in response to a marginal tax rate increase at income y , relative to the fraction of agents whose labor supply is distorted.

It is important to note that, while certainly important to gain economic insights, the formula (7) is not a full solution for the optimum tax as y itself does depend on T in a nonlinear way since individuals optimize with respect to the tax code.

It is useful to discuss two different variants of the optimal tax formula used in the literature. Saez (2001) derives the formula that is identical to (7). On the left hand side of this formula is the marginal tax on income which is the main object of interest. The right hand side is, however, defined over endogenous variables – the elasticity along the nonlinear budget constraint $\varepsilon_*(y)$, the density $\phi_*(y)$ and the c.d.f. $\Phi_*(y)$ of incomes are evaluated at the optimum and hence themselves depend on the income tax schedule $T_*(y)$. The Diamond (1998) formula, in the case of iso-elastic preferences for labor $v(l) = \frac{l^{1+1/\epsilon}}{1+1/\epsilon}$, is given by

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

The right hand side is a closed-form expression, since the distribution of types θ (cdf F , pdf f) is exogenous and the elasticity ϵ is given. However, the left hand side is not a closed form expression for the income tax schedule $T(\cdot)$. Indeed, note that the left hand side gives the marginal tax rate faced by a type θ . But this tax rate is evaluated at the income $y(\theta)$ that the

agent earns given the (optimal) tax schedule. This variable is endogenous, and is given as the implicit solution to the first order condition of the agent, which obviously does not give $y(\theta)$ as a function of θ in closed form (even with the isoelastic functional form of the disutility of labor). In other words, the Diamond formula gives the tax on the type θ . We are interested in the tax schedule in the space of incomes. However, the relationship between the income and the type is unknown. In fact, it is the essence of the nonlinear income tax problem that the types are unobservable and the tax schedule is over incomes. The Diamond and the Saez formulas are essentially identical to each other and neither provides a full characterization of the optimum.

3 The optimal tax and the fairness principle

In this section, we provide a new property of the optimum – the fairness principle – viewing the static optimal taxation problem from a dynamical point of view.

First, we rewrite the optimal static tax formula in an operator notation. Consider a second order differential operator $L = \frac{\partial}{\partial y} \left(\sigma_*(y) \frac{\partial}{\partial y} \right)$, where $\sigma_*(y) = \varepsilon_*(y) \phi_*(y)$. The optimal tax in (7) is then given by

$$LT_* = -\phi_*.$$

It is known that operators of this type can be associated with the heat kernel (Grigor'yan 2009). We also show in Section 4 that the heat equation arises from studying the gradient flow of the tax reform, that is, the trajectory of the steepest increase in revenue. The static optimal tax is a stationary point of such a dynamical system.

For the analysis of the optimal tax we proceed as follows. We first introduce the heat kernel and show that the optimal tax satisfies a certain invariance property – the fairness principle – with respect to this object. Specifically, the optimal tax at a given income can be represented as the weighted (by the heat kernel) average of the optimal taxes at the other incomes. We then characterize the form of the heat kernel and show that it behaves as a Gaussian average that proportionally downweights more distant incomes. This weighting is indexed by time (or, maybe more intuitively, by scale) where averaging is done over a broader set of incomes. Importantly, this weighting is tightly connected to each other at every scale, thus providing one unifying weighting scheme at every scale (or set of incomes). Finally, we show a representation of the optimal marginal income tax as well as the higher derivatives of the optimal tax schedule in terms of the levels of the optimal taxes.

We first introduce the heat kernel. Let $q_t(x, y)$ be the heat kernel given by the solution to the Kolmogorov forward equation

$$\frac{\partial}{\partial t} q_t(x, y) = \frac{\partial}{\partial y} \left(\varepsilon_*(y) \phi_*(y) \frac{\partial}{\partial y} q_t(x, y) \right), \quad (8)$$

and $\lim_{t \rightarrow 0} q_t(x, y) = \delta(x - y)$, where δ is a Dirac delta function.⁸

An example of the heat kernel is a Gaussian in Figure 1, we show later in this section that the heat kernel satisfying equation (8) behaves similarly to the Gaussian. The heat kernel $q_t(x, y)$, for a given income x , is a function of two variables – income y and time t . One can think of time as a different scale over incomes. At time (scale) zero, the kernel puts the weight one on income x . The larger times (scales) average and encompass a wider set of incomes.

We now state and prove the new property of the optimal tax – the fairness principle.

Proposition 1. (*Fairness principle for the optimum*). *The optimal tax $T_*(y)$ is invariant under the heat kernel $q_t(x, y)$ given by (8), for any $x \in Y$ and any $t > 0$:*

$$T_*(x) = \int_0^t \int_{\mathbb{Y}} q_s(x, y) \phi_*(y) dy ds + \int_{\mathbb{Y}} q_t(x, y) T_*(y) dy. \quad (9)$$

Proof. Consider the derivative

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

by Kolmogorov forward equation

$$= \int_{\mathbb{Y}} \frac{\partial}{\partial y} \left(\varepsilon_*(y) \phi_*(y) \frac{\partial}{\partial y} q_t(x, y) \right) T_*(y) dy =$$

integrating twice by parts and rearranging

$$= \int_{\mathbb{Y}} q_t(x, y) \frac{\partial}{\partial y} \left(\varepsilon_*(y) \phi_*(y) \frac{\partial}{\partial y} T_*(y) \right) dy =$$

using equation (6)

$$= - \int_{\mathbb{Y}} q_t(x, y) \phi_*(y) dy.$$

Integrating the equation

$$\frac{\partial}{\partial t} \int_{\mathbb{Y}} q_t(x, y) T_*(y) dy = - \int_{\mathbb{Y}} q_t(x, y) \phi_*(y) dy,$$

we get

$$\int_{\mathbb{Y}} q_t(x, y) T_*(y) dy = T_*(x) - \int_0^t \int_{\mathbb{Y}} q_s(x, y) \phi_*(y) dy ds.$$

□

This proposition states that the optimal tax $T_*(x)$ is fair in the following sense – the optimal tax at a given income x is equal to the weighted (by the heat kernel) average of taxes at other

⁸The kernel $q_t(x, y)$ solves both the forward and the backward Kolmogorov equation $\partial_t q_t(x, y) = L_y q_t(x, y) = L_x q_t(x, y)$.

incomes and income densities. The level of the optimal tax at a given income x is determined as the average of taxes at all other incomes and at *all* times (scales) t . We start by focusing on the more interesting second integral, $\int q_t(x, y) T_*(y) dy$, in (9). Since this term arises from the behavioral effect of taxation it means that the planner smoothes the distortions or the deadweight loss of taxes at every scale t . The averaging is done with the 1-parameter family of local averaging functions $q_t(x, y)$, where each function corresponds to one particular instance of the underlying fairness principle. In other words, the optimal tax wants to ensure that an agent at a given income x is paying roughly the average of the amount of taxes paid by people working just a little less or just a little more (“little less” or “little more” is determined by the heat kernel that downweights the more distant incomes). This quantifies an underlying notion of fairness that is in no way built into the system. The first integral represents the mechanical effect of raising taxes and is a weighted average of all densities of incomes y at all times (scales) t .⁹ Of course, our term fairness has nothing to do with the notions of the social welfare function or the redistributive preferences for the government. In fact, the government here maximizes revenues. Yet, the most efficient way to raise revenue is to do it fairly in the sense of equating (with certain weights) the tax at any income to the taxes of incomes at all scales.

The key to our results is to 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, one 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. At time 0, the optimal tax is just equal to itself – thus being trivially fair. As time goes on, the heat kernel shows that the optimal tax is fair in the sense that it is equal to the average over an increasingly wider distribution $q_t(x, y)$. One can alternatively think of the variable t not as time but as a scale – and the tax being fair at each scale t encompassing the weight of more and more incomes. We expand on the dynamic system interpretation of the optimum in the next section where the optimum arises as a stationary point of a gradient flow.




Figure 1: The heat kernel $q_t(x, y)$ for various times t (smaller times correspond to larger maxima).

⁹One reason why this term is somewhat less central is that, for small t , it is generally of the smaller order than the second integral term (see Proposition 3 and footnote 16).

The precise form of the fairness principle depends on the precise form of q_t . We now turn to the more detailed characterization of the heat kernel.

Proposition 2. (1) *The heat kernel (8) satisfies a Gaussian upper bound, for any $t > 0$:*

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

for some positive constants c_1, c_2, c_3 ;

(2) *For $t \rightarrow 0$,*

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

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

(3) *For all point x, y and all times $t, s > 0$, the heat kernel satisfies the semigroup property:*

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

This proposition shows that the heat kernel essentially behaves similarly to the properly scaled Gaussian average.¹⁰ The first part of the proposition shows that the heat kernel satisfies a Gaussian upper bound for all times t . The second part of the proposition shows that for short t , the heat kernel is exactly the Gaussian. The third part of the proposition shows a fact of crucial importance – the heat kernels are tightly linked in at all time scales. We now provide a more extensive discussion of these results.

The heat kernel q_t is in general a very benign object that is fairly easy to compute to any desired degree of accuracy. Three basic properties are (1) $q_t(x, y) \geq 0$, (2) $q_t(x, y) = q_t(y, x)$ and (3) preservation of integral mass

$$\int q_t(x, y)dy = 1.$$

We can thus, for a given point x , understand $q_t(x, y)$ as a one-parameter family of probability distributions in the variable y . This motivates understanding them as averaging objects. A classical result of Aronson (1968) is that the heat kernel $q_t(x, y)$ on a general manifold M (satisfying very mild regularity assumptions) satisfies what is called a Gaussian upper bound

$$q_t(x, y) \leq \frac{c_1}{t^{n/2}} \exp\left(-\frac{d(x, y)^2}{c_2 t}\right), \quad \forall t > 0, x, y \in M,$$

¹⁰In one dimension that we have, the body of literature on parabolic PDEs shows that for almost any modification of the problem, the heat kernel looks and behaves exactly as a Gaussian – this is true for the heat kernels on arbitrary manifolds, for very wide classes of conductivities σ , and for a very broad range of spaces (see e.g. Grigor'yan (2009) or Bogachev, Krylov, Röckner, and Shaposhnikov (2015) for extensive reviews).

where $d(x, y)$ is the geodesic distance between x and y , and c_1 and c_2 are positive constants. In particular, while the heat kernel $q_t(x, y)$ may no longer look like a Gaussian centered at y having variance t , it certainly has the same decay behavior. That is, it acts as a local averaging operator at scale $d(x, y) \sim \sqrt{t}$. In other words, the fairness principle averages the nearby income, where the nearby is given by the scale \sqrt{t} .¹¹ In the proposition, we use a slightly more general result in Metafune, Ouhabaz, and Pallara (2011).

Finally, the classical results for the short-time asymptotics (see, e.g., Varadhan (1967), Molchanov (1975), and Grigor'yan (2009)) imply that for $t \rightarrow 0$, the heat kernel is the Gaussian with the scale determined by the conductivity parameter $\varepsilon_*(x) \phi_*(x)$. In the next section, we provide additional results and intuition for this small time asymptotics.

Returning to the interpretation of q_t as creating an averaging operator at scale $\sim \sqrt{t}$, the third part of the proposition shows the fairness principles for different scales are linked. A fact of crucial importance is that they are tightly linked via what is known as the semigroup property (see, e.g., Grigor'yan 2009). This shows that the behavior of q_t is tightly linked to both past and future behavior of the heat kernel and is thus far from arbitrary. That is, there is one unified weighting scheme at all income scales.

We now obtain from Proposition 1 the corresponding representation for the marginal tax.

Corollary 2. *The fairness principle is invariant under differentiation and implies that the marginal tax is given by*

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

for any $x \in Y$ and any $t > 0$

The relevant quantity, $\partial_x q_t(x, y)$, is quite simple to understand in one dimension: since $q_t(x, y)$ is a probability distribution, its derivative has total integral 0. This means that $\partial_x q_t(x, y)$ has a positive part and a negative part with the same total L^1 -mass and acts as a discrete differentiation operator. We plot it in Figure 2 (together with the higher derivatives of $q_t(x, y)$). In economic terms, the integral evaluates a weighted average of taxes paid by individuals with slightly higher incomes, subtracts a weighted average of taxes for individuals with slightly lower incomes, and this results in the quantity determining the size of the marginal tax $T'_*(x)$. Moreover, this fairness principle for the marginal tax holds, as the original fairness principle, for all $t > 0$.

It is an elementary mathematical fact that given any function $f \in C^1(\mathbb{R})$, it is possible to change it ever so slightly into a function $f_2 \in C^1(\mathbb{R})$ such that f_1 and f_2 give almost the same values everywhere

$$\max_{x \in \mathbb{R}} |f_1(x) - f_2(x)| \leq \varepsilon$$

¹¹If one is interested in the higher order expansions, those can be straightforwardly derived in closed form to any order using the parametrix method which represents the heat kernel as the sum of the Gaussian and the higher order corrections (see, e.g., Friedman (2008)).




Figure 2: q_t (dashed), its derivative q'_t and its second derivative (bold) q''_t .

but f_2 has a very different derivative

$$\max_{x \in \mathbb{R}} |f'_1(x) - f'_2(x)| \geq \frac{1}{\varepsilon}.$$

Put differently, even a very good understanding of the optimal tax code T_* need not a priori translate into a good understanding of the marginal tax T'_* . We show here that this is not the case, the marginal tax is uniquely determined by the global behavior of the optimal tax code.

One way of interpreting this statement is as follows: for any function $f \in C^1$, f at a point is the local average of its neighboring values (indeed, this follows from continuity and boundedness and does not require differentiability)

$$\lim_{t \rightarrow 0^+} \int q_t(x, y) f(y) dy = f(x).$$

The fairness principle states that the optimal tax code satisfies a much stronger relationship

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

for *all* times $t > 0$. We can now see whether there is an analogous result for the derivative. Differentiating the fairness principle yields

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

and it is of interest to understand whether there is an analogous result for *all* functions f . We perform the relevant computations, for simplicity of exposition, for the heat kernel of the Laplacian on \mathbb{R} , i.e.

$$q_t(x, y) = \frac{1}{\sqrt{4\pi t}} e^{-\frac{|x-y|^2}{4t}}.$$

A simple computation shows that

$$\frac{\partial}{\partial x} q_t(x, y) = -\frac{(x-y)}{4\sqrt{\pi}t^{3/2}} e^{-\frac{|x-y|^2}{4t}}$$

and from this we see that

$$\lim_{t \rightarrow 0^+} \int \left(\frac{\partial}{\partial x} q_t(x, y) \right) f(y) dy = f'(x)$$

This follows easily from a Taylor expansion of f around x . While this relationship is true for all continuous $f \in C^1(\mathbb{R})$, it only holds in the limit $t \rightarrow 0$. In contrast, as shown by the derivative of the fairness principle, the optimal tax code T_* satisfies a relationship of this type for *all* $t > 0$. The same principle holds for higher derivatives, and these results can be obtained in the same manner.

Summarizing, in this section we derive a new characterization of the optimal tax. The fairness principle, while still of course not a closed-form solution, provides a new set of insights on the nature of the optimal tax.

4 Gradient flows of taxes

We now turn to the analysis of taxes from a different point of view. We construct a dynamic system, a gradient flow, which starts at any (optimal or suboptimal) tax function and then changes the tax system in the direction of the increased revenues. The optimal tax is a stationary point of this system.

In this section, we use the straightforward adaptation of notation in Section 1 to index the relevant variables by time.

We start by formally defining the gradient flow.

Definition 1. For all $t \geq 0$ and $y \in \mathbb{Y}$, the *gradient flow* of the government revenue functional $R(T_t)$ is defined as the dynamical system:

$$\frac{\partial T_t(y)}{\partial t} = \phi_t(y) + \frac{\partial}{\partial y} [T'_t(y) \varepsilon_t(y) \phi_t(y)], \quad (10)$$

where ϕ_t is governed by T_t according to the change of variables $\phi_t(y(\theta)) = (y'_t(\theta))^{-1} h(\theta)$ and by equation (2).¹²

¹²If the government is social welfare maximizing, it is immediate to show that the gradient flow has the form of $\frac{\partial T_t(y)}{\partial t} = \phi_t(y) (1 - \gamma_t(y)) + \frac{\partial}{\partial y} [T'_t(y) \varepsilon_t(y) \phi_t(y)]$, where $\gamma_t(y)$ is the social marginal utility of income (see Diamond (1975)).

4.1 Mathematical foundations

Finite-dimensional spaces. Gradient flows are natural mathematical objects attached to functions or functionals, mapping to real numbers. For simplicity, start with a differentiable function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ and define the gradient flow as a curve $x : [0, \infty] \rightarrow \mathbb{R}^n$ starting at some point $x_0 \in \mathbb{R}^n$ with the property that the curve always flows in the direction of steepest descent of V . Intuitively, this direction is determined by the gradient of V . Formally, we want to choose the vector $\hat{x} \in \mathbb{R}^n$ with $\|\hat{x}\|_{\ell^2} = 1$ that minimizes

$$\lim_{\mu \rightarrow 0} \frac{1}{\mu} [V(x + \mu\hat{x}) - V(x)] = \langle \nabla V, \hat{x} \rangle_{\ell^2},$$

where the equality follows from the definition of the gradient. This gives rise to an ordinary differential equation that describes the law of motion of $x_t \in \mathbb{R}^n$ for $t \geq 0$:

$$\frac{d}{dt} x_t = -\nabla V(x_t),$$

with the property that V is decreasing along the flow of x since

$$\frac{d}{dt} V(x_t) = \langle \nabla V, \frac{d}{dt} x_t \rangle = -\|\nabla V(x_t)\|^2 < 0.$$

While this model is rather classical and the existence and uniqueness properties of the solution are well known, understanding the actual dynamical behavior can pose considerable challenges (recent examples being given by Tao (2017), Steinerberger (2018)).

Infinite-dimensional spaces. The very same principle can be applied in settings where the underlying domain is not finite-dimensional but instead given by the space of functions. We illustrate this with a representative example. We may define a functional Ψ by assigning to any twice-differentiable *function* $f \in \mathcal{C}^2(\mathbb{R}, \mathbb{R})$, the number

$$V(f) = \frac{1}{2} \int_{\mathbb{R}} |f'(x)|^2 dx.$$

It is easy to show that the Gateaux differential of V in the direction \hat{f} is given by $\delta V(f, \hat{f}) = \int_{\mathbb{R}} f'(x) \hat{f}'(x) dx$. An integration by parts implies that $\delta V(f, \hat{f}) = -\int_{\mathbb{R}} f''(x) \hat{f}(x) dx$. More generally, for any function $f \in \mathcal{C}^2(\mathbb{R}^n, \mathbb{R})$, let $V(f) = \int_{\mathbb{R}^n} |\nabla f|^2$. We then have $\delta V(f, \hat{f}) = \int_{\mathbb{R}^n} \nabla f \cdot \nabla \hat{f}$. By Green's first identity, this can be represented as a functional $\hat{f} \mapsto \langle -\Delta f, \hat{f} \rangle_{L^2}$, where Δ denotes the Laplace operator, thus recovering the same structure as above. That is, in order to flow in the direction of steepest descent of the functional V , we must set

$$\frac{\partial}{\partial t} f_t = \Delta f_t.$$

This gives rise to a law of motion for the function f characterized by a parabolic PDE (namely, a heat equation). Needless to say, even showing that all of these operations remain valid for any time $t > 0$ is a difficult task, the theory of partial differential equations being substantially more challenging than that of ordinary differential equations.

4.2 The gradient flow of taxes

We now turn to the formal derivation of the the gradient flow (10). An integration by parts in the second integral of equation (5) implies that the impact of the tax reform \hat{T}_t on government revenue can equivalently be rewritten as

$$\delta R(T_t, \hat{T}_t) = \int_{\mathbb{Y}} \Lambda_t(y) \hat{T}_t(y) dy, \text{ with } \Lambda_t \equiv \phi_t + \frac{\partial}{\partial y} [T'_t \varepsilon_t \phi_t] \quad (11)$$

Assuming that the space of functions $\mathcal{C}^2(\mathbb{R}, \mathbb{R})$ is endowed with the L^2 norm $\|T\|^2 = \int (T(y))^2 dy$, this can be expressed as $\langle \Lambda_t, \hat{T}_t \rangle$.¹³ Therefore, the gradient flow we obtain in this case can be written as the dynamical system (10).

The gradient flow (10) can be equivalently derived as the solution to the problem of choosing the trajectory of the tax schedule $t \mapsto T_t$ that maximizes at each instant t the increase in government revenue:

$$\max_{T_t} \frac{\partial}{\partial t} R(T_t),$$

in the L^2 norm. Specifically, the evolution of government revenue $R(T_t)$ over time for a given trajectory T_t is given by

$$\frac{\partial}{\partial t} R(T_t) = \frac{\partial}{\partial t} \int T_t(y_t(\theta)) dH(\theta) = \int \left[\frac{\partial T_t}{\partial t}(y_t(\theta)) + T'_t(y_t(\theta)) \frac{\partial y_t}{\partial t}(\theta) \right] dH(\theta).$$

Imposing that the individual's first-order condition remains satisfied over time requires (by differentiation of (2)) that

$$\frac{\partial y_t}{\partial t}(\theta) = -\varepsilon_t(\theta) \frac{\partial T'_t(y_t(\theta))}{\partial t}.$$

That is, at each instant, individual θ adjusts his income in the opposite direction and proportionally to the change in the marginal tax rate that he faces.

Plugging this equation back into the law of motion of government revenue and integrating

¹³We could have considered the weighted- L^2 norm $\|T\|^2 = \int \kappa_t(y) (T(y))^2 dy$, for some weights $\kappa_t(y)$ and this expression can be represented as $\langle \kappa_t^{-1} \Lambda_t, \hat{T}_t \rangle$ with the resulting gradient flow $\frac{\partial T_t(y)}{\partial t} = (\kappa_t(y))^{-1} \phi(y) + (\kappa_t(y))^{-1} \frac{\partial}{\partial y} [T'_t(y) \varepsilon_t(y) \phi_t(y)]$. The analysis for this case is identical. There is a re-interpretation of such a weight as simply changing the metric of the underlying manifold \mathbb{R} : put differently, one can interpret everything as an equal-weight problem on a curved geometry; heat and associated processes are not very sensitive to “curving” (heat propagates on a plane and on a sphere in roughly the same sense). See, e.g., Taylor (1996), Hörmander (2003), and Grigor'yan (2009).

the second term by parts leads to

$$\frac{\partial}{\partial t} R(T_t) = \int \frac{\partial T_t}{\partial t} \{ \phi_t(y) + \frac{\partial}{\partial y} [T_t'(y) \varepsilon_t(y) \phi_t(y)] \} dy = \left\langle \Lambda_t(y), \frac{\partial T_t(y)}{\partial t} \right\rangle.$$

This expression is maximized when $\frac{\partial T_t}{\partial t}(y) = \phi_t(y) + \frac{\partial}{\partial y} [T_t'(y) \varepsilon_t(y) \phi_t(y)]$, thus leading to (10).

For another way to understand the economic meaning of (10), consider the problem of choosing the tax reform \hat{T} that maximizes the increase in government revenue, subject to the following quadratic cost of reforming the tax payments:¹⁴

$$\max_{\hat{T}} \delta R(T_t, \hat{T}) - \frac{1}{2} \int_{\mathbb{Y}} (\hat{T}(y))^2 dy.$$

Using the representation (11), the solution is given by

$$\hat{T}_t(y) = \Lambda_t(y).$$

Now, the law of motion of the tax schedule in the small interval of time $\delta t \rightarrow 0$ is given by $T_{t+\delta t}(y) = T_t(y) + \hat{T}_t(y) \delta t$, or $\hat{T}_t(y) = \frac{\partial T_t(y)}{\partial t}$. We therefore obtain the gradient flow (10).

4.3 Short-term evolution and fairness property

In this section, we describe the short-term evolution of the tax schedule $T_{\tilde{t}}$ under the gradient flow by solving the heat equation (10) over a short time interval $[t, \tilde{t}]$. We derive a version of the fairness principle that now applies to any point on the path of evolution of the tax system.

Proposition 3. (*Fairness principle for the gradient flow*). *Consider any initial time t with the corresponding tax profile $T_t(y)$, density of incomes $\phi_t(y)$, elasticity $\varepsilon_t(y)$, and conductivity $\sigma_t(y) = \phi_t(y) \varepsilon_t(y)$. Then, for small $\tilde{t} - t$, the tax $T_{\tilde{t}}(y)$, generated by the gradient flow (10), is given by a weighted Gaussian average of the incomes:*

$$T_{\tilde{t}}(x) \sim (\tilde{t} - t) \phi_t(x) + \int_{\underline{y}}^{\bar{y}} q_{t,\tilde{t}}(x, y) T_t(y) dy,$$

$$\text{where } q_{t,\tilde{t}}(x, y) = \frac{1}{\sqrt{4\pi\sigma_t(x)(\tilde{t}-t)}} \exp\left(-\frac{(y-x-\sigma_t'(x)(\tilde{t}-t))^2}{4\sigma_t(x)(\tilde{t}-t)}\right).$$

Proof. The proof uses the Feynman-Kac formula for path integrals (Lorinczi, Hiroshima, and Betz 2011) to study the short-time behavior of solutions of equations of the type (10). Let $B(s)$ denote the diffusion process that satisfies the SDE.

$$dB_s = \sigma'_s(B_s) ds + \sqrt{2\sigma_s(B_s)} dW_s,$$

¹⁴Again, we can use weights $\kappa(y) > 0$ in the cost function (e.g., they can be equal to the density function $\phi_t(y)$) and derive essentially the same results.

where W is a Brownian motion. We then have

$$T_{\tilde{t}}(x) = \mathbb{E} \left[\int_t^{\tilde{t}} \phi_s(B(s)) ds \right] + \mathbb{E}[T_t(B_{\tilde{t}})],$$

where the expectation runs over the diffusion process B_s , started in y and running up to \tilde{t} .¹⁵ We can now perform a Taylor expansion of these quantities. Up to the first order, diffusivity is constant at a certain scale. The short-time asymptotics for Brownian motion is then given by a Gaussian distribution $B_{\tilde{t}} \sim x + \sigma'_t(x) (\tilde{t} - t) + \sqrt{2\sigma_t(x)} W_t \sim N(x + \sigma'_t(x) (\tilde{t} - t), 2\sigma_t(x) (\tilde{t} - t))$ and

$$\text{distribution of } B_{\tilde{t}} \sim \frac{1}{\sqrt{4\pi\sigma_t(x) (\tilde{t} - t)}} \exp \left(-\frac{(y - x - \sigma'_t(x) (\tilde{t} - t))^2}{4\sigma_t(x) (\tilde{t} - t)} \right).$$

Then, up to a first order for \tilde{t} small,

$$\int_t^{\tilde{t}} \phi_t(B(s)) ds \sim (\tilde{t} - t) \phi_t(x).$$

This implies that

$$\begin{aligned} T_{\tilde{t}}(x) &\sim (\tilde{t} - t) \phi_t(x) + \mathbb{E}[T_t(B_{\tilde{t}})] \\ &\sim (\tilde{t} - t) \phi_t(x) + \int_{\underline{y}}^{\bar{y}} \frac{1}{\sqrt{4\pi\sigma_t(x) (\tilde{t} - t)}} \exp \left(-\frac{(y - x - \sigma'_t(x) (\tilde{t} - t))^2}{4\sigma_t(x) (\tilde{t} - t)} \right) T_t(y) dy. \end{aligned} \quad (12)$$

□

This equation extends the notion of fairness that we derived for the optimal tax to that of the trajectory of the gradient flow of taxes. The government that considers a tax reform in the direction of maximizing revenues changes the tax such that the new, evolved tax $T_{\tilde{t}}(x)$ is equal to the weighted average of the initial taxes $T_t(y)$ plus another source term $\phi_t(x)$ that evaluates the mechanical effect of the revenues collected. In other words, the gradient flow wants to ensure that an agent at a given income x is paying roughly the average of the amount of taxes paid by people working just a little less or just a little more. The conductivity $\varepsilon_t(x) \phi_t(x)$ determines the scale of the Gaussian and the drift correction and the function $\phi_t(x)$ determines the asymptotic deviation.¹⁶

Note, that the property (12) holds for any starting time t on the gradient flow trajectory

¹⁵Technically, we need to specify boundary conditions for the Brownian motion $B(s)$, however, since we are only using Brownian motion for very small times \tilde{t} , it does not matter very much whether we prescribe absorbing boundary conditions corresponding to Dirichlet conditions or reflecting boundary conditions corresponding to Neumann conditions.

¹⁶Of course, this equation also holds for the optimal tax, as it is a stationary point of the gradient flow in which both the taxes and agents' behavior no longer change: $T_*(x) = t\phi_*(x) + \int_{\underline{y}}^{\bar{y}} \frac{1}{\sqrt{4\pi\sigma_*(x)t}} \exp \left(-\frac{(y - x - \sigma'_*(x)t)^2}{4\sigma_*(x)t} \right) T_*(y) dy + o(t)$.

and that we average over the known and given parameters $\varepsilon_t(x)$ and $\phi_t(x)$ evaluated at the time t . That is, it is a closed-form expression. This characterization is valid for the short time \tilde{t} as these parameters are essentially frozen over that short time interval. For the large time, it is also a fairly good approximation of the heat kernel, but with the conductivity and the source needing to be adjusted as the agents change their behavior in response to the evolution of the tax function. In contrast, the results for the optimal tax are derived for any time (scale) t as well as for the short term asymptotics but include the elasticities and the densities evaluated at the optimum.

There is also a sense in which the gradient flow acts on the points that deviate most from this principle. It is encoded in the equation

$$T_{\tilde{t}}(x) \sim (\tilde{t} - t) \phi_t(x) + \mathbb{E}[T_t(B_{\tilde{t}})].$$

Say the source is $\phi_t \equiv 0$ and suppose that $T_t = 1$ everywhere except in $x = 0$. Then $T_{\tilde{t}}(0) = 1$ as the averaging corrects this deviation. We elaborate the discussion of the smoothing properties of the gradient flow in the next section.

4.4 Gradient flow as a heat equation

In this section, we propose a construction of a trajectory of the tax reform and show how the gradient flow of the revenue functional gives rise to the heat equation.

Equation (10) changes the tax schedule in favor of increasing government revenue, letting ϕ_t and ε_t be endogenously driven by T_t – that is, taking into account the fact that the density and the elasticity change in response to the evolution of taxes.

We propose to evolve the system separately (this underlying idea is a straightforward application of “operator splitting”). The simplest instance of this idea is as follows. Suppose we are given a system of ordinary differential equations given as

$$\frac{d}{dt}u(t) = (A + B)u(t),$$

then the solution is given by the matrix exponential $u(t) = e^{t(A+B)}u(0)$. A formal Taylor series expansion suggests that

$$\begin{aligned} e^{t(A+B)} &= \text{Id} + t(A + B) + \mathcal{O}(t^2) \\ &= (\text{Id} + tA)(\text{Id} + tB) + \mathcal{O}(t^2) \\ &= e^{tA}e^{tB}u(0) + \mathcal{O}(t^2). \end{aligned}$$

These computations are purely formal but they do suggest that, at least for small values of t , we may solve the system by first evolving along the simpler system $\dot{u}(t) = Au$ and then along the system $\dot{u} = Bu$ and alternate in this manner (Varga (1962), Glowinski and Osher (2016)).

We apply the very same method in our problem: more precisely, we fix the distribution of incomes ϕ_t and the elasticity ε_t for a short period of time δt , evolve T_t , and then re-compute $\phi_{t+\delta t}$ and $\varepsilon_{t+\delta t}$ based on the new tax function $T_{t+\delta t}$.¹⁷ In standard situations, this procedure will converge to a solution path of the dynamical system as $\delta t \rightarrow 0$ (Glowinski, Osher, and Yin (2016)). The operator splitting technique also has a natural economic meaning. The government evolves taxes in the direction of increased revenues, keeping the density of agents' incomes and elasticities fixed at their value observed in the current economy. That is, the government evaluates the changes in revenues under the current information given by the exogenous sufficient statistics evaluated at a given initial time.

This implies equation (10) is a heat equation (with source term ϕ and local conductivity $\sigma = \varepsilon\phi$), i.e. a PDE of the form

$$\frac{\partial T}{\partial t} = \phi(y) + \frac{\partial}{\partial y}[\sigma(y) \frac{\partial T}{\partial y}] \quad (13)$$

and guarantees in particular that the problem always has a solution (the heat equation being well-posed). We also note that our assumption that the density tends to 0 at the boundary of the interval implies that no boundary conditions need be imposed. Since heat equations are among the most well-known and well-behaved partial differential equations, we can apply standard mathematical results to obtain theoretical properties of the evolution of the tax schedule over time.

Fixing ϕ and ε , letting T evolve for a short amount of time, then unfreezing ϕ and ε and recomputing it can be regarded as a classical example of operator splitting. While the analysis of convergence of this dynamic system is outside the scope of the paper, one can expect that for sufficiently short time steps, the solution converges to the global optimum at a great level of generality. For example, the review of Glowinski, Osher, and Yin (2016, p.13) concludes: “Last but not least, operator splitting algorithms are theoretically attractive because they converge under very few assumptions.” More broadly, the splitting procedure we use is similar in spirit to the ones used in physical sciences where the split terms correspond to different physical processes – for example, splitting convection from diffusion (see, e.g. MacNamara and Strang (2016)) or splitting fast from slow variables.

4.5 Smoothing properties of the heat equation

We now present various smoothing properties of the gradient flow arising within a step of the operator splitting; this gradient flow is realized by a heat equation. Most of the results in this section are straightforward adaptations of very classical results for the heat equation. Since the heat equation can alternatively be realized as a gradient flow in a certain Sobolev-type space, it has a very nice underlying smoothing structure which is reflected in a large number of

¹⁷This also can be regarded as a classical numerical technique for systems of this type.

beneficial mathematical properties. Moreover, since the heat equation is one of the most studied (and well-behaved) objects in mathematics, the list of the useful properties is very large. We therefore focus only on some of them in this section that show how the heat equation smoothes the underlying functions. Various other properties could be of interest but are outside the scope of this paper. For the rest of the section to ease the notational burden we suppress the indexing of taxes with t whenever it does not cause confusion.

Since the initial tax system can be arbitrary, we need to introduce the proper reference point for the analysis of smoothing of a given tax system. We proceed as follows. Since $\phi(y)$ and $\sigma(y)$ are fixed, the solution will converge to a fixed point as $t \rightarrow \infty$ and the unique fixed point is given by setting $\partial T / \partial t = 0$ resulting in

$$0 = \phi(y) + \frac{\partial}{\partial y} \left[\sigma(y) \frac{\partial T(y)}{\partial y} \right]. \quad (14)$$

We define the solution to this equation as $\tau(y)$. This stationary point is the solution to the problem of maximizing tax revenue, conditional on keeping the density of agents' incomes and elasticities fixed. In other words, this is the tax schedule under exogenous sufficient statistics that would be optimal if these sufficient-statistic variables were fixed and equal to their value observed in the current economy, i.e., given the current tax T_t .

It is important to note, however, that we are not interested in letting the gradient flow for all $t \rightarrow \infty$ and thus finding τ but only evolving it for a very small time t (so as to respect the operator splitting). In this sense, τ serves as a proper reference point to describe various smoothing properties of the gradient flow.

We first relate the gradient flow arising from revenue maximization to another gradient flow with strong smoothing properties.

Proposition 4. *Equation (14) coincides with the gradient flow of the functional \mathcal{J} given by the weighted Sobolev-type seminorm H^1*

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

Proof. Let us compute the directional derivative in the direction of a function w evaluated at T

$$\delta \mathcal{J}(w) = \lim_{\varepsilon \rightarrow 0} \frac{\mathcal{J}(T + \varepsilon w) - \mathcal{J}(T)}{\varepsilon}.$$

We see that

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \frac{\mathcal{J}(T + \varepsilon w) - \mathcal{J}(T)}{\varepsilon} &= \int \sigma(y) (T'(y) - \tau'(y)) w'(y) dy \\ &= \int w(y) \frac{\partial}{\partial y} (\sigma(y) (\tau'(y) - T'(y))) dy. \end{aligned}$$

This shows that the negative gradient flow is given by

$$\frac{\partial T}{\partial t} = -\frac{\partial}{\partial y} (\sigma(y) (\tau'(y) - T'(y))) = -\frac{\partial}{\partial y} (\sigma(y)\tau'(y)) + \frac{\partial}{\partial y} (\sigma(y)T'(y)).$$

Using the equation for τ , we see that $-\frac{\partial}{\partial y} (\sigma(y)\tau'(y)) = \phi(y)$ and we have established the desired claim. \square

This proposition has an interesting economic meaning. This gradient flow has the effect of trying to smooth out rough irregularities in the difference between T and τ – a large value of $|T'(y) - \tau'(y)|$ implies the existence of a large value of $(T'(y) - \tau'(y))^2$ and the flow is trying to decrease this as quickly as possible. We note that the quantity $\sigma(y) \geq 0$ serves as a natural weighting measure: if $\sigma(y)$ is large, then irregularities in that region count even more severely and are dampened quicker than in regions where $\sigma(y)$ is very small.

Moreover, if $T(y)$ has large amounts of strong oscillations or maybe even discontinuous jumps, then the gradient flow acts strongest on those parts first. This leads to the following proposition which follows from the classical result on parabolic equations.

Proposition 5. *Let $T(t, y)$ denote the solution of (13). If $\sigma(y)$ is bounded away from 0 and $T(0, y)$ is bounded, then $T(t, y)$ is infinitely differentiable for any $t > 0$.*

This implies that the 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.

The next proposition shows that 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. As we discussed above, 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. As the initial tax function $T(0, y)$ and hence the associated agents' behavior (that determine $\phi(y)$ and $\varepsilon(y)$) can be arbitrary, the correct reference point for this smoothing behavior is the corresponding stationary solution $\tau(y)$.

Proposition 6. *Let $T(t, y)$ be the solution to (13), $T(0, y)$ be an arbitrary initial tax schedule, and $\tau(y)$ be the solution to the stationary problem (14), and λ_1 be the first eigenvalue of the associated Sturm-Liouville operator:*

$$\frac{\partial T}{\partial t} = \phi(y) + \frac{\partial}{\partial y} [\sigma(y) \frac{\partial T}{\partial y}]. \quad (15)$$

Then, $\forall t > 0$:

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

Proof. We use the standard Sturm-Liouville theory (see Zettl (2010), Teschl (2012), Titchmarsh (1962)) to prove this result. This equation can be studied by first solving for the stationary

problem

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

Subtracting both equations leads to an equation for $z(t, x) = T(t, x) - \tau(x)$ given by

$$\frac{\partial z}{\partial t} = \frac{\partial}{\partial y}[\sigma(y) \frac{\partial z}{\partial y}].$$

It remains to study problems of this type. We will do so by studying the spectrum of the associated differential operator H given by

$$Hz = -\frac{\partial}{\partial y}[\sigma(y) \frac{\partial z}{\partial y}]$$

or, in other words, we study the problem $Hz = \lambda z$. This eigenvalue problem has a discrete sequence of admissible values λ for which the equation has a solution: these values $0 < \lambda_1 < \lambda_2 < \dots$ are the eigenvalues of this operator of Sturm-Liouville type, the corresponding solutions will be denoted by η_1, η_2, \dots and are assumed to be L^2 -normalized, i.e. $\|\eta_n\|_{L^2} = 1$. We note that $\lambda_0 = 0$ is a special value and $\eta_0 = \text{const}$. We see that these eigenfunctions are necessarily orthogonal in $L^2(a, b)$ since, again by integration by parts,

$$\begin{aligned} \int_a^b \eta_k(y) \eta_\ell(y) dy &= \frac{1}{\lambda_k} \int_a^b -(\sigma(y) \eta'_k(y))' \eta_\ell(y) dy \\ &= -\frac{1}{\lambda_k} \int_a^b -\sigma(y) \eta'_k(y) \eta'_\ell(y) dy \\ &= \frac{1}{\lambda_k} \int_a^b (-\sigma(y) \eta'_\ell(y))' \eta_k(y) dy \\ &= \frac{\lambda_\ell}{\lambda_k} \int_a^b \eta_k(y) \eta_\ell(y) dy \end{aligned}$$

If $k \neq \ell$, then the factor in front of the integral is different from 1 and the integral is therefore 0. This together with the completeness of the system of eigenfunctions in L^2 allows us to expand an arbitrary initial function $T(0, x)$ into a series

$$T(0, x) = \tau(x) + \sum_{k=1}^{\infty} \langle z(0, x), \eta_k(x) \rangle \eta_k(x).$$

We will abbreviate $a_k = \langle z(0, x), \eta_k(x) \rangle$ for simplicity of exposition. We then claim that

$$T(t, x) = \tau(x) + \sum_{k=1}^{\infty} a_k e^{-\lambda_k t} \eta_k(x)$$

is a solution of the partial differential equation (15). This can be verified by computing

$$\left(\frac{d}{dt} - \frac{d}{dx}\sigma(x)\frac{d}{dx}\right)\sum_{k=1}^{\infty} a_k e^{-\lambda_k t} \eta_k(x) = \sum_{k=1}^{\infty} a_k \left(\frac{d}{dt} - \frac{d}{dx}\sigma(x)\frac{d}{dx}\right) e^{-\lambda_k t} \eta_k(x).$$

The separation of variables implies that

$$\left(\frac{d}{dt} - \frac{d}{dx}\sigma(x)\frac{d}{dx}\right) e^{-\lambda_k t} \eta_k(x) = e^{-\lambda_k t} \left(-\lambda_k \eta_k(x) - \frac{d}{dx}\sigma(x)\frac{d}{dx}\eta_k(x)\right) = 0$$

as desired. Since we now have a complete description of a solution, we can analyze the convergence to the limiting function arising for $t \rightarrow \infty$ at a greater level of detail: we have

$$\begin{aligned} \int_a^b (T(t, y) - \tau(y))^2 dy &= \int_a^b \left(\sum_{k=1}^{\infty} a_k e^{-\lambda_k t} \eta_k(y)\right)^2 dy = \int_a^b \sum_{k, \ell=1}^{\infty} a_k e^{-\lambda_k t} \eta_k(y) a_{\ell} e^{\lambda_{\ell} t} \eta_{\ell}(y) dy \\ &= \int_a^b \sum_{\ell=1}^{\infty} a_{\ell}^2 e^{-2\lambda_{\ell} t} \eta_{\ell}(y)^2 dy = \sum_{\ell=1}^{\infty} a_{\ell}^2 e^{-2\lambda_{\ell} t} \leq e^{-2\lambda_1 t} \sum_{\ell=1}^{\infty} a_{\ell}^2. \end{aligned}$$

We note that

$$\sum_{\ell=1}^{\infty} a_{\ell}^2 = \int_a^b (T(0, y) - \tau(y))^2 dy$$

and that we have therefore shown that

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

Appealing to the classical Rayleigh-Ritz formula, we see that

$$\lambda_1 = \inf_{\int_a^b f(y) dy = 0} \frac{\int_a^b \sigma(y) f'(y)^2 dy}{\int_a^b f(y)^2 dy} \quad (16)$$

where the last step follows from the classical Neumann eigenvalue computation for the homogeneous rod (see Courant and Hilbert (1989)). This shows that for sufficiently regular values of $\phi(y)$, we can expect $\lambda_1 > 0$ and therefore the distance to $\tau(y)$ undergoes exponential decay. \square

Finally, note that the short term asymptotics results in Section 4.3 do not require the use of operator splitting methods.

5 Conclusion

We have shown that dynamical methods can provide new insights on the analysis of static optimal taxation problems. We show that the heat kernel and the heat equation are intimately connected with the analysis of this classic problem. Since the heat equation is one of the

most basic mathematical objects, it possesses a variety of useful properties that can enrich our understanding of the mechanism design problems. One such new characteristic is the fairness property where the taxation system implies that a tax on a given income is a proper average of taxes at other incomes and income densities. The fairness principle does not stem from any desire of the planner to be fair but rather shows that the most efficient way to raise revenues is to equalize a properly weighted average of taxes at any given income. We have shown that this principle holds for both the optimum at any scale and for the gradient flow for the short time approximation. The derivation of the fairness principle fundamentally relies on the dynamical view of the static problem.

We now briefly outline some extensions. We considered the simplest taxation environment to starkly highlight the main contributions of the paper. Several extensions, some of which we already discussed in the body of the paper are immediate. Consider a model where utility is not quasi-linear, the government maximizes social welfare function, agents have multidimensional types (not incomes). This model delivers a gradient flow that is very similar to the one that we constructed, with the properly modified elasticities and the social welfare weights. More broadly, one can consider a variety of other mechanism design problems – nonlinear pricing, matching, etc., – where gradient flows, or more broadly viewing a static problem from the point of view of a dynamical system, may be useful in characterizing the static optimum and the evolution of the locally improving suboptimal policies.

References

Aronson, Don. Non-negative solutions of linear parabolic equations, *Ann. Sci. Norm. Sup.* 22, p. 607–694. 1968.

Artzner, Philippe, Carl P. Simon, and Hugo Sonnenschein. "Convergence of Myopic Firms to Long-Run Equilibrium via the Method of Characteristics." *Models of Economic Dynamics*. Springer, Berlin, Heidelberg, 157-183. 1986.

Bierbrauer, Felix, and Pierre Boyer. "Politically feasible reforms of non-linear tax systems." 2018.

Blanchet, Adrien, and Guillaume Carlier. "Optimal transport and Cournot-Nash equilibria." *Mathematics of Operations Research* 41, no. 1: 125-145. 2015.

Bogachev, Vladimir I., Nicolai V. Krylov, Michael Röckner, and Stanislav V. Shaposhnikov. *Fokker-Planck-Kolmogorov Equations*. Vol. 207. American Mathematical Soc., 2015.

Bolton, Patrick, and Christopher Harris. "The dynamics of optimal risk sharing." No. w16094. National Bureau of Economic Research. 2010.

Brewer, Mike, Emmanuel Saez, and Andrew Shephard. "Means-testing and tax rates on earnings." *Dimensions of Tax Design: the Mirrlees Review*. 2010.

Chiappori, Pierre-André, Robert J. McCann, and Lars P. Nesheim. "Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness." *Economic Theory* 42, no. 2: 317-354. 2010.

Chiappori, Pierre-André, and Bernard Salanié. "The econometrics of matching models." *Journal of Economic Literature* 54, no. 3 : 832-61, 2016.

Courant, Richard and David Hilbert, *Methods of mathematical physics*. Vol. I. Interscience Publishers, Inc., New York, N.Y., 1953.

Daskalakis, Constantinos, Alan Deckelbaum, and Christos Tzamos. "Strong Duality for a Multiple-Good Monopolist." *Econometrica* 85, no. 3: 735-767. 2017.

Diamond, Peter A. "A many-person Ramsey tax rule." *Journal of Public Economics* 4.4: 335-342. 1975.

Diamond, Peter A. "Optimal income taxation: an example with a U-shaped pattern of optimal marginal tax rates." *American Economic Review*: 83-95. 1998.

Figalli, Alessio, Young-Heon Kim, and Robert J. McCann. "When is multidimensional screening a convex program?." *Journal of Economic Theory* 146, no. 2: 454-478. 2010.

Friedman, Avner. *Partial differential equations of parabolic type*. Courier Dover Publications, 2008.

Galichon, Alfred. *Optimal transport methods in economics*. Princeton University Press. 2016.

Glowinski, Roland, Stanley J. Osher, and Wotao Yin, eds. *Splitting Methods in Communication, Imaging, Science, and Engineering*. Springer, 2017.

- Grigor'yan, Alexander. Heat Kernel and Analysis on Manifolds, AMS/IP Studies in Advanced Mathematics. 2009.
- Golosov, Mikhail, Aleh Tsyvinski, and Nicolas Werquin. A variational approach to the analysis of tax systems. No. w20780. National Bureau of Economic Research, 2014.
- Hellwig, Martin F. "Incentive problems with unidimensional hidden characteristics: A unified approach." *Econometrica* 78, no. 4: 1201-1237. 2010.
- Hörmander, Lars. The Analysis of Linear Partial Differential Operators I – IV, Springer, 2003.
- Jordan, Richard, David Kinderlehrer, and Felix Otto. "The variational formulation of the Fokker-Planck equation." *SIAM Journal on Mathematical Analysis* 29.1: 1-17. 1998.
- Jacquet, Laurence, and Etienne Lehmann. "Optimal Income Taxation when Skills and Behavioral Elasticities are Heterogeneous." 2015.
- Kleven, Henrik Jacobsen. "Sufficient Statistics Revisited." 2018.
- Kleven, Henrik Jacobsen, and Claus Thustrup Kreiner. "The marginal cost of public funds: Hours of work versus labor force participation." *Journal of Public Economics* 90, no. 10-11: 1955-1973. 2006.
- Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez. "The optimal income taxation of couples." *Econometrica* 77, no. 2: 537-560. 2009.
- Lőrinczi, Jozsef, Fumio Hiroshima, and Volker Betz. Feynman-Kac-type theorems and Gibbs measures on path space: with applications to rigorous quantum field theory. Vol. 34. Walter de Gruyter, 2011.
- MacNamara, Shev, and Gilbert Strang. In Glowinski, Roland, Stanley J. Osher, and Wotao Yin. Operator splitting. In *Splitting Methods in Communication, Imaging, Science, and Engineering* (pp. 95-114). Springer, Cham. 2016.
- Metafune, Giorgio, El Maati Ouhabaz, and Diego Pallara. "Long time behavior of heat kernels of operators with unbounded drift terms." *Journal of Mathematical Analysis and Applications* 377, no. 1: 170-179. 2011.
- McCann, Robert J. "Academic wages, singularities, phase transitions and pyramid schemes." *Proceedings of the International Congress of Mathematicians (Seoul 2014)*. Vol. 3. 2014.
- Mirrlees, James A. "An exploration in the theory of optimum income taxation." *The Review of Economic Studies* 38.2: 175-208. 1971.
- Molchanov, Stanislav A. "Diffusion processes and Riemannian geometry." *Russian Mathematical Surveys* 30, no. 1: 1-63. 1975.
- Musgrave, Richard A., and Tun Thin. "Income tax progression, 1929-48." *Journal of Political Economy* 56.6: 498-514. 1948.
- Sachs, Dominik, Aleh Tsyvinski, and Nicolas Werquin. Nonlinear tax incidence and optimal taxation in general equilibrium. No. w22646. National Bureau of Economic Research. 2016.
- Saez, Emmanuel. Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, vol. 68, no 1, p. 205-229. 2001.

- Saez, Emmanuel, and Stefanie Stantcheva. "Generalized social marginal welfare weights for optimal tax theory." *American Economic Review* 106, no. 1: 24-45. 2016.
- Saez, Emmanuel, and Stefanie Stantcheva. "A simpler theory of optimal capital taxation." *Journal of Public Economics* 162: 120-142. 2018.
- Salanie, Bernard. *The economics of taxation*. MIT press, 2011.
- Scheuer, Florian, and Iván Werning. *Mirrlees meets diamond-mirrlees*. No. w22076. National Bureau of Economic Research. 2016.
- Sonnenschein, Hugo. "Price dynamics and the disappearance of short-run profits: An example." *Journal of Mathematical Economics* 8.2: 201-204. 1981.
- Sonnenschein, Hugo. "Price dynamics based on the adjustment of firms." *The American Economic Review* 72.5: 1088-1096. 1982.
- Steinerberger, Stefan. "Fast escape in incompressible vector fields." *Monatshefte für Mathematik* 186, no. 3: 525-537. 2018.
- Tao, Terence. "On the universality of potential well dynamics." arXiv preprint arXiv:1707.02389. 2017.
- Taylor, Michael. *Partial Differential Equations I-III*, Springer. 1996.
- Teschl, Gerald. *Ordinary differential equations and dynamical systems*. Vol. 140. American Mathematical Soc., 2012.
- Tirole, Jean, and Roger Guesnerie. "Tax reform from the gradient projection viewpoint." *Journal of Public Economics* 15.3: 275-293. 1981.
- Titchmarsh, Edward Charles. *Eigenfunction expansions associated with second-order differential equations. Part I*. Second Edition Clarendon Press, Oxford. 1962.
- Varadhan, Sathamangalam R. Srinivasa. "On the behavior of the fundamental solution of the heat equation with variable coefficients." *Communications on Pure and Applied Mathematics* 20, no. 2: 431-455. 1967.
- Varga, Richard S. *Matrix Iterative Analysis*, New Jersey: Prentice-Hall, 1962
- Villani, Cedric. *Topics in optimal transportation*. No. 58. American Mathematical Soc., 2003.
- Zettl, Anton. *Sturm–Liouville Theory*. Providence: American Mathematical Society, 2005.