Tax policy and entrepreneurial entry with information asymmetry and learning

Diego d'Andria¹

Published online: 20 April 2019
© The Author(s) 2019

Abstract

We study a market with entrepreneurial and worker entry where both entrepreneurs' abilities and workers' qualities are private information. We develop an agent-based computable model to mimic the mechanisms described in a previous analytical model (Boadway and Sato in *Int Tax Public Finance* 18(2):166–192, 2011). Then, we introduce the possibility that agents may learn over time about abilities and qualities of other agents, by means of Bayesian inference over informative signals. We show how such different assumptions affect the optimality of second-best tax and subsidy policies. While with no information, it is optimal to have a subsidy to labour and a simultaneous tax on entrepreneurs to curb excessive entry, with learning the detrimental effects of excessive entry are partly compensated by surplus-increasing faster learning.

Keywords Entrepreneurship · Taxation · Asymmetric information · Learning · Adverse selection · Agent-based computational model

JEL Classification D82 · D83 · G14 · H25

The author would like to thank Salvador Barrios, Ivan Savin, participants to the IIPF Congress 2017 in Tokyo and two anonymous referees of for useful comments and suggestions which greatly improved the paper. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author. They should not be attributed to the European Commission. Any mistake and all interpretations are his and his only.

✉ Diego d'Andria
diego.dandria@ec.europa.eu

¹ European Commission Joint Research Centre (JRC), Unit B2 Fiscal Policy Analysis, Calle Inca Garcilaso 3, 41092 Seville, Spain

1 Introduction

At least since Schumpeter (1934) entrepreneurial entry has been seen as a major driver for economic growth. Entrepreneurs challenge incumbent firms thus stimulating improvements and productivity gains and can themselves bring innovations into the market, in the form of new products or processes. A reflection of this view is the widespread existence of policies designed to foster entrepreneurship, to support small young firms and to ease credit constraints. At the heart of the market failure that these policies try to solve is asymmetric information. First, there is an intrinsic riskiness in starting a new venture which makes investors require a risk premium thus raising the cost of capital. Second, the chances of entrepreneurial success depend on a number of factors including the entrepreneur's human capital, which is not directly observable thus giving rise to adverse selection. Third, the entrepreneur himself faces asymmetric information at the time of acquiring input factors, notably when hiring personnel and in cases when third-party technologies have to be employed.

This paper focuses on simultaneous informational asymmetries in the credit market and in the labour market. We develop an agent-based computational model which is designed to mimic the mechanisms described in Boadway and Sato (2011). Our novel contribution lies in explicitly taking into account learning over time and its consequences on the effects of tax and subsidy policies. Simulation analysis is here used to test the effects of policies made of a tax on entrepreneurs and a subsidy to labour income, in a more realistic setting. We study a market that is not in an extreme condition of either full or no information, where excess entry favours learning and thus produces more information in less time.

The intuition behind our claim is rather straightforward. If agents over time learn about other agents' ability or quality, then having more entrants in earlier periods increases the amount of information available and leads the market closer to a condition of full information. Consequently, the optimal level of entry with learning is larger than in a first-best scenario with full information; as the loss in surplus suffered as more bad entrepreneurs enter, the market is (partly, or entirely) compensated by an increase in future surplus due to more efficient market conditions.

In our model entrepreneurs and workers decide whether to enter an entrepreneurial market, entrepreneurs hire workers and fund investment costs by means of external financing, and a government may levy taxes or subsidies on both entrepreneurial income and wages. Asymmetric information plagues both the credit market and the input (labour) market. The estimated entrepreneur's ability¹ affects the cost of financing she faces, while the estimated worker's quality affects his expected wage. After replicating the main features of a market as described in Boadway and Sato (2011) for the polar cases with full information and with no information available, we then introduce the possibility that agents may learn over time, thus better estimating abilities and qualities thanks to informative signals.

Our simulation results point to the fact that the way information is produced in the economy meaningfully affects the scoring of different policy options. In some settings

¹ We keep throughout the paper the convention to use male pronouns for workers and female pronouns for entrepreneurs. Also following the original convention in Boadway and Sato (2011), the term ability refers to entrepreneurs while quality refers to workers.

a tax on entrepreneurs, which would be optimal under a no-information regime, might be detrimental welfare-wise when learning occurs. Our results suggest that the optimal levels of a subsidy to workers and of a tax on entrepreneurs strongly depend upon the information regime observed in real markets and that learning always implies lower taxes on entrepreneurs at the optimum compared to the no-information regime.

In the following, Sect. 2 summarizes the relevant previous literature. Section 3 describes the model, and Sect. 4 discusses the simulation results. Section 5 concludes and points to future avenues for research.

2 Related literature

This paper draws from several previous research contributions. The literature on adverse selection in credit and entrepreneurial markets, particularly with regard to new ventures, poses the basis for the analysis of the asymmetric information on the side of entrepreneurs (Stiglitz and Weiss 1981). The main message from this literature often closely resembles Akerlof (1970) in predicting underinvestment and less than optimal (from a social planner's point of view) entry of new ventures, though subsequent research (De Meza and Webb 1987; Boadway and Keen 2006) finds conditions under which this outcome can be reversed to overinvestment.

Adverse selection on the side of labour inputs acquisition and its effects on new firms is relatively less studied. Weiss (1980) analyses a market with unobservable workers' quality and shows how adverse selection would draw more workers with low skills from the pool of candidates, with overall lower employment than optimal. Most other studies focused on moral hazard in the form of unobservable effort provision (Shapiro and Stiglitz 1984; Holmstrom and Milgrom 1991).

The taxation literature has analysed tax and subsidy policies that may be employed to obtain first- or second-best optimality. (Good summaries of results and open questions are found in Rosen 2005; Keuschnigg and Nielsen 2003; Henrekson and Sanandaji 2011.) Most of these, though, again focus on the problem of non-monitorable effort; for instance, Keuschnigg and Nielsen (2003) develop a model where venture capitalists and entrepreneurs jointly provide effort, and they find how taxes on labour, capital and capital gains income should be designed. The effects of specific types of taxes on bonus compensation with unobserved work effort have been studied with regard to managers (Radulescu 2012; Dietl 2013), and with regard to innovative employees when also corporate taxes and tax incentives are levied (d'Andria 2016).

The possible interactions between the two informational asymmetries, however, remain largely unexplored. One exception is of course Boadway and Sato (2011) which we are employing here as a basis to develop our own agent-based computational model. Boadway and Sato (2011) discuss a theoretical model where entrepreneurs with unknown ability decide whether to start a risky project. Those who decide to start such a project hire workers whose quality is, as well, unknown. This kind of market generates inefficient levels of entrepreneurial activity and adverse selection simultaneously on the side of entrepreneurial borrowing and on the side of labour supply. Policy-wise one of the main results is that in a scenario where double adverse selection induces excessive entry of low-ability entrepreneurs and the hiring of too many low-quality

workers, the second-best optimal policy is made of a subsidy to labour (which serves the purpose to attract high-quality workers) paired with a tax on entrepreneurs (which curbs the excessive entry of entrepreneurs, the latter being increased by the labour subsidy).

A related strand of literature models agents as facing the choice whether to become an entrepreneur or to work as an employee. Recent contributions have highlighted that the earnings from jobs alternative to entrepreneurship can either be randomly associated with the abilities as entrepreneurs (Scheuer 2013, where regressive profit taxation is found to be the optimal policy response) or correlated with abilities (d'Andria 2018a). In the latter case, taxes levied on entrepreneurial income can have very different effects on entry decisions based on the specific correlation assumed. In our simulation model, we assume that entrepreneurs only face the decision whether to become entrepreneurs or to invest in a safe asset, while workers decide whether to work for an entrepreneurial firm or for a non-entrepreneurial firm. Therefore, we rule out the possibility for an entrepreneurial-type agent to opt for employment. This choice stems from at least three reasons. First, we do so in order to allow for clearer comparisons with Boadway and Sato (2011) (and also with the previous literature, in particular, Stiglitz and Weiss 1981 and De Meza and Webb 1987) who opted for the same modelling choice. Second, as shown in d'Andria (2018a), different assumptions about how entrepreneurial ability and productivity as an employee correlate are not neutral and may completely change the way the model behaves. This would in turn multiply the number of scenarios and corresponding simulation runs needed to account for different assumptions, in a way that would be hard to manage in practice. Third, having an outside option that varies with own ability introduces another degree of heterogeneity across agents, which in turn would require larger populations and more replications for each simulation run, again posing technical issues with the optimization code used to find optimal policies.

Methodologically, our choice is agent-based computational economics (ACE) modelling, see Tesfatsion and Judd (2006) for an in-depth description and a number of examples taken from the literature. The rationale for using agent-based computational modelling techniques in economics has been widely discussed already, for example, in Tesfatsion (2003) and Judd (2005). In our case, as we aim at introducing learning over time—which is intrinsically a stochastic path-dependent process—together with heterogeneous agents, we believe this comes as a rather natural choice. We can then deal with multiple equilibria using a Monte Carlo approach and study the way the simulated model converges, on average, to different outcomes. We believe we avoid the critique, common to many agent-based models, of a lack of generality of the results stemming from simulated runs of the model with specific parametrizations thanks to the fact that we developed a computational optimizer to explore a wide range of policies and find the optimal mix of taxes and subsidies. Other parameters in the models are chosen following a minimal-assumptions approach to be as neutral as possible, often just taking the mean of the range of admissible values or deriving the parameters to meet a specific endogenous value for some variable (e.g. to obtain 50% entry rates under the baseline benchmark scenario with full information).

Agent-based models have been employed in the past (though only occasionally) to study labour markets (i.e. in Tesfatsion 2001 and Neugart 2008) and taxation (mainly

in relation to tax evasion, see, for instance, Bloomquist 2006; Hashimzade 2015 and Warner 2015). A relevant example is d'Andria and Savin (2018) where an agent-based model was developed to study a market for innovative workers with unobserved effort and workers' qualities, multiple job tasks and taxes both at the corporate and personal level. We follow here a philosophy similar to d'Andria and Savin (2018) in first developing a model that resembles as closely as possible the features and mechanisms of a corresponding analytical model (in the case of the present paper the reference analytical model is Boadway and Sato 2011), and then we change only one key assumption (by introducing learning) to see how the behaviour of the model changes. A similar approach has been used in the past in Yildizoglu (2002) where an agent-based model was developed that first follows an existing and well-known model (in that case, Nelson and Winter 1982), and then a different element is introduced (in Yildizoglu 2002 the new element consisted of a different set of assumptions about investment behaviour based on a learning mechanism); finally, the convergence value of the model is computed as a mean across several replications, and such convergence values computed with and without the new assumptions are compared.

3 The model

3.1 Overview of the model

The model² generates a population of entrepreneurs at the beginning of each simulation, drawing for each entrepreneur an ability value a from a uniformly distributed continuous interval $[0, 1]$. Similarly, a population of workers is generated with qualities q drawn from a uniformly distributed continuous interval $[0, 1]$.³ Each entrepreneur may hire a fixed number n of workers, and by assumption, the number of workers in the initial population is n times the number of entrepreneurs. Ability/quality is assigned to each agent and stays constant across the simulation. Ability represents the probability of success for an entrepreneur in sector E ,⁴ while quality is an index of productivity of a worker.

Each simulation runs for a fixed number of periods. In each period, each worker decides to enter the entrepreneurial sector E if, and only if, he expects his wage to be not lower than an exit option wage he can obtain with certainty from working in a

² The model code was written in Matlab/Octave. It was tested in Matlab (R2015a) and GNU Octave (4.0.3). To improve reproducibility of results, we used Octave to produce the figures reported in the following text, running on an Intel Xeon E5-2690 2.60GHz CPU with 14 cores, 128 Gb of RAM memory and under Microsoft Windows Server 2008 R2 Enterprise. All source code is made available for research purposes upon request to the author.

³ In the actual code, the values of abilities and qualities are rounded down to the nearest second decimal for coding convenience. For the same reason, extracted abilities and qualities equal to zero are converted to 0.01. This way of coding introduces a very minor bias in the learning algorithm, which we tested to be of trivial importance and in no way affecting our arguments.

⁴ An alternative interpretation of the values a assigned to entrepreneurs, which we will not pursue here, would be as probabilities of success for an underlying entrepreneur-specific project that requires repeated financing over time to be completed, and which discloses information about such probability over time as it gets developed.

traditional sector T . The wage in sector T is made equal to the worker's quality and is known to him. The expected wage from working in sector E depends on the scenario that is being simulated. If full information is assumed, then the true quality of each worker is common knowledge, and each worker knows that if hired in E , the wage he will be paid is equal to his true quality. If the scenario is of no information, all workers are always assigned the average quality calculated across the worker population employed in E (as assumed in the no-information case in Boadway and Sato 2011). If the scenario allows for learning, wage is equal to the estimated quality for each worker. (The learning algorithm that generates such estimated qualities is described below.) Therefore, only in the zero-information scenario are all workers paid the same wage. Wages are assumed to be paid upfront and therefore do not depend upon the probability of entrepreneurial success. (Again the latter assumption is taken to follow Boadway and Sato 2011 closely.)

In each period, each entrepreneur is assigned a gross interest rate $r \geq 1$ that incorporates a risk premium. Entrepreneurs face a probability of success equal to their ability a . If they fail, they are assumed to go bankrupt and repay nothing to the bank. If they succeed, they repay to the bank the cost of labour times r . As the risk-free interest rate is assumed constant and equal to a value ρ , the interest rate for an entrepreneur is simply ρ divided by the entrepreneur's estimated ability, $\frac{\rho}{a_e}$, so that higher-ability entrepreneurs face lower interest rates.⁵ Similarly to workers' qualities, the estimated ability a_e for entrepreneurs depends upon the scenario of choice. With full information, it is equal to the true ability value. With no information, it is equal to the average across the population of entrepreneurs. The latter is equivalent to assuming that the distribution of abilities and qualities in sector E is common knowledge. Finally, with learning, estimated ability a_e changes over time based on a learning algorithm. Consequently, only under zero-information will all entrepreneurs be pooled together and pay the same interest rate r , while they will face different rates in the other two scenarios.

Given a true ability a and an estimated ability a_e , the expected profit of an entrepreneurial project from the point of view of banks is:

$$E_{\text{banks}}(\pi) = a_e(R\bar{q}_e^\alpha - r(a_e)n\bar{q}_e) \quad (1)$$

where $\bar{q}$ is the mean estimated quality of hired workers. Produced value is here given by a standard production function $y = R\bar{q}^\alpha$ where R is total factor productivity (assumed the same for all entrepreneurs) and $0 < \alpha < 1$ such that the quality of the hired workforce increases production with marginally decreasing returns. As said before, the value of the interest rate r depends on the estimated ability a_e that other agents believe the entrepreneur possesses. From the point of view of an entrepreneur, her expected profit is:

$$E(\pi) = a(R\bar{q}_e^\alpha - r(a_e)n\bar{q}_e) \quad (2)$$

the difference with Eq. (1) being that the entrepreneur is assumed to know his own ability value a . Entrepreneurs have too an external option of value π_0 which is inter-

⁵ In the simulation code, a cap of 1,000 is put on the maximum value of r to avoid infinite values.

puted as a risk-free investment opportunity, and they decide to enter sector E if, and only if, $E(\pi) \geq \pi_0$.

After entry is determined for all agents, a matching algorithm assigns the n workers having the best-looking estimated qualities q_e to the entrepreneur having the best-looking estimated abilities a_e , and then the next n best-looking workers to the second best-looking entrepreneur, and so on. The reason for assuming a ranked-matching stems from the fact that better-looking entrepreneurs can always offer a slightly larger wage to a better-looking worker. (This point was already discussed in Boadway and Sato 2011, for the cases with full and no information.)

Finally, each enterprise enters the production phase and can either be successful (with probability a) or fail (with probability $1-a$). The cost of production is determined by the sum of the estimated qualities of employees, times r . (There are no capital costs other than what constitutes the net interest paid to banks in case of success.) Revenues are determined by the true qualities of employees, meaning that better quality implies larger productivity.

The model is parametrized such that it replicates Boadway and Sato (2011),⁶ and in the scenario with full information such parametrization makes 50% of entrepreneurs enter sector E and causes the same share of workers to be hired there. The total surplus of this market is maximized without policy intervention.

In the scenario with no information, on the contrary, a “market for lemons” dynamics occurs and less than 50% of workers are employed in E while entrepreneurial entry is also below 50%, resulting in lower total surplus. Under zero-information, low-ability and low-quality agents have larger incentives to enter sector E (the opposite holds true for higher-ability and higher-quality agents) because of the pooling mechanism that provides them with lower interest rates and higher wages, respectively, compared to full information. Moreover, as an additional source of inefficiency, better entrepreneurs and better workers are not matched. These effects further decrease surplus under the zero-information scenario.

3.2 Learning algorithm

The model under scenarios with full and no information is not dynamic, in the sense that each period independently produces a market equilibrium. Thus, the number of periods assumed in a simulation does not affect the derivation of an optimal policy. In scenarios with learning, the model becomes truly dynamic as estimated ability, or quality, is assigned to an agent based on past behaviour and outcomes. Thus, the equilibrium in each period is affected in a path-dependent way by all previous periods.

The scenarios with learning feature a learning algorithm. At the beginning of a simulation, all agents are assigned the mean value from their respective population (as per the scenario with no information). Then, in subsequent periods a noisy informational signal is observed by all agents and used to infer the true underlying value of ability or quality. Both parameters σ_a and σ_q , which represent the standard deviations

⁶ We only consider the case in Boadway and Sato (2011) where the employment rate of workers who opted for entry in sector E is not 100%, and therefore, some of the workers will not be hired in E and will work in sector T instead.

of the noise factors, are assumed constant over time and equal for all entrepreneurs and workers, respectively.

In each simulated period, each entrepreneur who entered sector E is observed to be either successful or not successful. The number of successes in past periods divided by the number of total entrepreneurial projects started, plus a noise $\sim \mathcal{N}(0, \sigma_a)$, is used as an estimation for the true underlying probability of success (which is just equal to the ability value a). The noise parameter is such that a high-ability entrepreneur may look, after some periods of learning, like an average ability one but very unlikely like a low-ability one. (The opposite holds for low-ability entrepreneurs.)

For workers, an informative signal of their quality is observed in each period, regardless of whether they entered or not sector E . This signal is equal to their true quality q plus a noise $\sim \mathcal{N}(0, \sigma_q)$. The vector of such signals obtained in the past, $\bar{q}_e$, is used within a Bayesian inference algorithm. The value for q corresponding to the largest estimated probability $P(q|\bar{q}_e)$ is picked as best-guess and assigned as the estimated quality q_e to the worker. (If multiple q values are associated with an equal probability, their mean is taken instead.)

As stated, we employ Bayesian learning. This is interpreted as an as-if representation of rational expectations (see Feldman 1987) formed over a worker's quality, given a set of observed signals. The prior belief on (conditional) quality distribution is assumed to be correct and common knowledge among all agents, that is, the probability $P(\bar{q}_e|q)$ for each possible q is obtained from the normal probability density function of $\mathcal{N}(q, \sigma_q)$. The learning algorithm is "non-adaptive" in the sense explained in Marimon (1996) and disregards the possibility that Bayesian predictions are somewhat affected by reinforcement-based predictions (Charness and Levin 2005).

With learning, therefore, the estimated values for abilities and qualities on average converge towards their true values. Also the longer the simulation, the more observation points are available to the agents so that the closer estimated values will be to the true ones. As entrepreneurs gain more observations, if they enter sector E , increasing entry may make the market converge faster to a higher informational level. For workers on the contrary, we assume that they produce signals about their quality even when working in the traditional sector T : this assumption represents the idea that in most cases workers can build up their curriculum vitae regardless of being hired by a new entrepreneurial firm. While debatable in principle, the latter assumption runs against our claim that the optimal entry level is larger with learning, so it does not really bear any implication for optimal policy (other than, possibly, strengthening our point).

3.3 Policy

As in Boadway and Sato (2011), the model allows for two types of policy instruments. The first is a tax or subsidy σ on labour income. The other instrument is a tax or subsidy τ on entrepreneurs. Both instruments are taxes if negative and are subsidies if positive. As there is no explicit bargaining process in the model to endogenously determine the split of profit between entrepreneurs and workers, we assume that any tax or subsidy σ on labour is split between the two parties in fixed shares. (This share in our calibration is 50%.)

A tax on entrepreneurs $\tau < 0$ is meant to represent, in a very synthetic way, the combined effects of the personal and corporate tax system on the individual choice to enter entrepreneurship. It can be viewed as a “success tax” as explained in Gentry and Hubbard (2000) exceeding personal taxation (if one assumes that the risk-free investment is taxed under personal taxation), or as a tax design for corporate taxation that penalizes new entrepreneurial firms (for example, by having an imperfect loss offset). Agents in the model are assumed risk-neutral, so the tax on entrepreneurs should not be interpreted as a device affecting risk-taking choices akin to Domar and Musgrave (1944).

The effects of a subsidy $\sigma > 0$ to labour are to induce more workers’ entry (by raising the expected wage in sector E), and also to induce more entrepreneurial entry (because of lower investment costs). A tax on labour $\sigma < 0$ would bear the opposite effects. Differently from σ , a tax or subsidy to entrepreneurs only affects their entry decision. Boadway and Sato (2011) (see their Proposition 3.i) argue therefore for a policy that is made of a subsidy to labour, and a tax on entrepreneurs meant to reduce excessive entrepreneurial entry stemming from the combined effects of no information and labour subsidy.

The expected profit of an entrepreneur, including the tax on profit and the subsidy to labour, can be therefore written as:

$$E(\pi) = a[R\bar{q}_e^\alpha - r(a_e)n(\bar{q}_e - \sigma\sigma_{\text{incidence}}) + \tau] \quad (3)$$

meaning that both the tax τ and the costs associated with the repayment of the debt are born by the entrepreneur only in case of successful outcome of the entrepreneurial project. The amount $(\bar{q}_e - \sigma\sigma_{\text{incidence}})$ is gross wage less the share of the worker subsidy σ falling onto firms (as stated, the incidence parameter $\sigma_{\text{incidence}}$ is set to 50%).

4 Simulation results

The model is run each time using the parameters listed in Table 1. It is just the case to stress that these parameters are not meant to resemble any level of realism nor are they calibrated based on real-world figures: they are meant to make the model behave in a certain way in order to support a theoretical argument. Table 1 summarizes the set of parameters employed.

The main output of a simulation is surplus, which (following Boadway and Sato 2011 and most of the literature on optimal taxation in partial equilibrium settings) is computed summing up all value generated by entrepreneurs (production minus labour costs times gross interest rates), less their opportunity costs (the alternative safe investments producing value π_0), plus wages earned by workers less their opportunity costs (the wage they could have earned in sector T). Profit and wages entering S are gross of taxes and subsidies, meaning that S represents value generated (or destroyed) after compensating for any subsidy paid, or tax collected, by the government. For this reason, the following simulation results include a measure of the net financial position of the government, in order to provide information about the budgetary cost of the simulated policy. Dividing S by this net position provides a measure of the

Table 1 Parameters used

Label	Description	Base value	Sensitivities
Periods	No. of periods	10	from 5 to 40
No. of workers	Initial population of workers	1000	
No. of entrepreneurs	Initial population of entrepreneurs	100	
R	Revenues multiplier	31	
α	Revenues exponent	0.5	0.3 and 0.7
ρ	Risk-free interest rate	2	
π_0	Profit from risk-free investment	1	
n	No. of employees per firm	10	
Learning (starts at)	First periods before learning applies	3	
σ_a	Std. dev. of abilities	0.25	0.35 and 0.15
σ_q	Std. dev. of qualities	0.15	0.25 and 0.05
Interest cap	Maximum interest rate allowed	1000	
$\sigma_{\text{incidence}}$	How σ is distributed	0.50	

additionality of the policy, that is, of how much value is created per net unit of public money spent. We label a policy “optimal” when it is found to maximize S , regardless of the corresponding net financial position of the government.

Surplus is here computed at the beginning of each period; thus, as workers are paid upfront while success for entrepreneurial projects is determined at the end of a period, any surplus coming from entrepreneurs is discounted using the same safe rate ρ assumed in the credit market. All surplus figures provided in the following tables have to be interpreted as measuring the net present value of surplus at the beginning of each period. In symbols, surplus in a given period is:

$$S = \sum (E(\pi)\rho^{-1} - \pi_0) + \sum (q_e - q) \quad (4)$$

Surplus is computed in each period based on the expected profit of entrepreneurial projects rather than actual realizations. The reason for this choice is purely technical and meant to reduce randomness in order not to have to run several replications of the same simulation type, thus drastically reducing computation time for the optimization algorithm used to find optimal policies.

For the scenarios with either full or no information, the use of expected surplus allows us to run one single simulation for each policy to be evaluated. For the scenarios with learning, given the path-dependencies implied by model design, multiple replications are run (10, or 5 when using the optimization algorithm and many periods). When replications are performed, average values are taken across periods and replications (this is the case for entry and employment figures) or they are derived by summing up across periods and then averaging across replications. (This is the case for surplus.) These average values are what we refer to as results in the following text. When no replications are performed, the values are averaged or summed up only across periods.

We first run the model assuming full information to calibrate it. We chose parameters that make the model obtain an entry of 50% for both types of agents under full information and without any tax or subsidy. (We refer to the latter as a no-policy scenario.) We then looked for a policy mix (made of a subsidy to labour and a tax on entrepreneurs)⁷ that maximizes surplus under no-information first, and then under learning, using our custom optimization algorithm. Because of the stochastic behaviour of the simulation model under the scenario with learning, some of the most common optimizers (gradient-based, directed search) are either unusable or do not perform well, the problem being akin to the optimization of a “rugged landscape”. A genetic algorithm could be used but it would take too long to perform the search. We therefore wrote a grid-based optimizer that computes the outcome from several policy mixes by changing taxes and subsidies, one at a time and by a fixed value, and then in a second step performs the same search locally, centering a smaller grid around the maximum value found in the first step.

Table 2 summarizes the outcome of the main simulations. Surplus is reported both as non-discounted sum across all periods (as stated, this is net present value of surplus evaluated at the beginning of each period in which it is generated), and as net present value in the first period of the simulation (NPS in the table), where the discount rate used is the safe interest rate ρ (thus, a very large discount rate as $\rho = 2$ in our parametrization). Table 2 reports entry rates and mean abilities and qualities of entered agents. Also, in the last column, the net financial commitment of the government is shown as total subsidies paid less total taxes collected (summed across all periods).

The model is such that with no information there is sub-optimal entry of entrepreneurs and workers, with worse average workers’ quality than the case with full information. Note that the average entrepreneurial ability in the market is close (even larger) than with full information: because the average worker quality deteriorates so does expected profit, consequently the few entrepreneurs entering the market have to be of high-ability type in order to have positive expected profit. This occurrence changes with a subsidy to workers: workers’ entry and quality both increase, and thus, more entrepreneurs are attracted into sector *E* but, because of the information asymmetry in the credit market, the pooling of abilities resulting in a common interest rate attracts lower-ability types. Thus, although entrepreneurial entry is close to the first-best level of entry, the mean ability is much lower than at the first-best (0.51 in the no-information scenario with subsidy, versus 0.75 in the full-information scenario).

As in Boadway and Sato (2011), total surplus is found to be greatest in the no information scenario, with a tax and subsidy policy mix. The optimal values found by the optimization algorithm are -75.7 for the tax on entrepreneurs, and 3.85 for the subsidy to workers. Table 2 reports also the results for a subsidy-only policy (equal to 3.85) to allow for comparisons and to appreciate the effects of the subsidy in isolation. The subsidy-only policy even reduces surplus (total and NPS) below what was obtained without policy intervention.

We then allowed for agents to learn. We run simulations first without any policy, and then with the optimal policy found for the no-information scenario ($\sigma = 3.85$

⁷ We also experimented by allowing the optimization to test for subsidies to entrepreneurs and/or taxes on workers. These were never found to be optimal; therefore, we omitted them to focus the time-consuming optimization algorithm only on those policies that are actual candidates for an optimal policy.

Table 2 Summary of simulation results

Scenario	Policy	S	NPS	Entr. entry (%)	Work. entry (%)	Avg. a	Avg. q	Gov. Net
Full info	None	9930	550	50	50	.75	.75	0
No info	None	3047	127	24.50	25.21	.87	.45	0
	$\sigma = 3.85$	2820	117	49.50	49.80	.51	.50	15,361
Learning	$\sigma = 3.85, \tau = -75.7$	3360	140	34.90	38.40	.65	.50	-1766
	None	6960	302	44.20	43.20	.70	.50	0
	$\sigma = 3.85, \tau = -75.7$	-3248	108	28.90	28.90	.68	.72	-2027
	$\sigma = 0.76$	8711	405	81.26	79.29	.53	.56	5354
	$\sigma = 0.76, \tau = -10.0$	9252	465	69.30	69.30	.57	.60	637

and $\tau = -75.7$). The previously found optimal policy reduced surplus to negative values in the scenario with learning. The optimal policy for the learning scenario was found to be made of a subsidy $\sigma = 0.76$ and a tax $\tau = -10.0$, thus both smaller in absolute values than for the no-information scenario. Table 2 summarizes these results. With 10 replications⁸ the surplus obtained from the learning scenario has a standard deviation of less than 2.5%, which means that the difference from the no-policy case is statistically significant at p -value < 0.0001 .

With learning and without policy actions, the market produces sub-optimal entry for both types of agents. The reason lies in the noisiness of ability and quality estimates, such that there is always a share of agents who are undervalued and thus do not enter sector E , and this share decreases in time as better estimates of abilities and qualities become available. Note that the optimal policy raises entry well above 50%, the reason being that with learning it is best to elicit larger entry in the first periods to improve the production of information over time. The latter result remains true even when comparing net present surplus which heavily discounts future surplus gains (but, see Sect. 4.1 for a sensitivity analysis w.r.t. the number of periods assumed). Table 2 also reports results for a scenario with learning and only a subsidy ($\sigma = 0.76$), again to allow for comparisons with the other policies.

Contrary to the no-information scenario, having learning agents implies that a policy made only of a labour subsidy is surplus-improving over the tax and subsidy policy that was optimal under no information. This remains true even if the subsidy-only policy produces some level of apparently excessive entrepreneurial entry: as explained, the efficiency gains from accelerated learning more than compensating the efficiency losses due to more low-ability entrepreneurs entering sector E .

An interesting property of the scenario with learning is that subsidies to workers elicit much more entry compared to no information. After in-depth examination of the microdata stemming from our simulations, our conclusion is that this larger reactivity is mostly due to the fact that even a small additional incentive to entry enhances the matching of workers and entrepreneurs. Improved matching greatly increases expected profit for the high-ability entrepreneurs, thus prompting their entry, which in turn allows for even larger entry of workers. This mechanism is reflected, in Table 2, in the relatively large average ability of entrepreneurs in the market ("large" relative to the fact that the entry rate is close to 70%).

Turning to the government position, one can see from Table 2 (last column) that optimal policies may produce both positive or negative net positions. Interestingly, our simulations suggest that the optimal policy under the scenario with learning produces very large additionality. On average, value generated in the latter scenario exceeds spent net public money by more than 14 times (obtained dividing S by the Government Net position from Table 2). Moreover, looking at individual simulations, after 100 replications of the scenario with learning and optimal policy we found that this ratio never fell below 3.5. While only suggestive as our model is not calibrated using real data, these results point to the possibility that taking into account dynamic (learn-

⁸ This is the number of replications used for all simulations involving learning reported in Table 2. As explained already, simulations with full- or no-information scenarios only need one single replication.

ing) effects the cost–benefit analysis of such policies generally might produce more favourable assessments.

4.1 Sensitivity analyses

In this section, some sensitivity analyses are presented and discussed.

The parameters characterizing production and outside earnings (these are R , α , n , ρ , π_0) cannot be changed individually without requiring adjustments to the other parameters in order to keep the calibration such that, with full information and no policy, entry rates are 50%. Consequently, we refrain from reporting a full set of sensitivity tests on these and just discuss here some noteworthy general effects. In particular, the parameter α affects the reactivity of produced value to the cost of labour and, consequently, affects expected profit and entry rates of entrepreneurs. Deviations from $\alpha = 0.5$ imply that entrepreneurs under scenarios with no information or with learning react more, or less, in terms of entry decisions to a change in the expected average quality of workers they can hire. Thus, values for α larger than 0.5 imply that a subsidy to workers increases entrepreneurial entry more than what is shown in Table 2, and vice versa if α is smaller than 0.5. We experimented with $\alpha = 0.7$ and $\alpha = 0.3$ which affected the results quantitatively but never qualitatively. In particular, optimal taxes and subsidies under no information are always much larger than in learning scenarios, and the subsidy-only policy that is optimal under no information is superior welfare-wise to one also including the optimal profit tax (again, we refer here to the tax that was found optimal under no information), when learning is included.

One seemingly important key parameter is the number of periods included in the simulations. A longer time span implies that the benefits from early discovery of abilities and qualities will have a long-lasting effect. But as the improvements in the estimation of abilities and qualities become marginally smaller in later periods, a longer time span also means that the positive effects of learning on total surplus will decrease relative to the total effect from having larger entrepreneurial activity. (The latter depresses surplus as the policy attracts more low-ability entrepreneurs into sector E .)

We thus changed the number of periods, halving it to 5 (from our central case with 10 periods) and then increasing it stepwise to study the effects of different parametrizations,⁹ and simulated both a no-policy scenario and one using the optimal policy found for 10 periods. The aim is to see how the number of periods affects the ability of the policy to improve market outcomes. Table 3 reports the resulting surplus for each simulation run, while Table 4 reports surplus as net present value. The last column in both tables reports the increase in surplus, or in net-present-value surplus, caused by the policy.

⁹ Our model in Octave takes about three days to perform 50 replications with 30 periods each. The optimization algorithm for the learning scenario usually needs some more days with 10 periods. As the amount of calculations needed for each additional period increases more than proportionally because of the Bayesian learning algorithm, we decided to stop the sensitivity analysis at 40 periods. For simulations with more than 10 periods, we also halved the number of agents (50 entrepreneurs and 500 workers, instead of 100 and 1000, respectively).

Table 3 Sensitivity on the number of periods in the scenario with learning

	A: S no policy	B: S optimal policy	Change (A to B) (%)
5 periods	2806	3924	+ 39.8
10 periods	6960	9252	+ 32.9
20 periods	15,739	20,049	+ 27.4
30 periods*	12,282*	15,352*	+ 25.0
40 periods*	16,859*	20,480*	+ 21.5

Figures are for total surplus. The optimal policy tested is $\sigma = 0.76$ and $\tau = -10$. An asterisk in columns 1, 2, 3 means that the simulated population was halved to reduce computation time

Table 4 Sensitivity on the number of periods in the scenario with learning

	C: NPS no policy	D: NPS optimal policy	Change (C to D) (%)
5 periods	38	39	+ 2.6
10 periods	302	465	+ 54.0
20 periods	303	467	+ 54.1
30 periods*	146*	226*	+ 54.8
40 periods*	143*	229*	+ 60.1

Figures are for surplus in net present value in period 1. The optimal policy tested is $\sigma = 0.76$ and $\tau = -10$. An asterisk in columns 1, 2, 3 means that the simulated population was halved to reduce computation time

Reading Tables 3 and 4, one can see how net present value surplus increases sharply with the number of periods up to a point where the additional value generated in future times becomes too small in discounted terms to make any visible difference (considering that the standard deviation of total surplus across 10 replications in the learning scenario is on average rather small, about 2.5% of the mean surplus). Total surplus monotonically decreases (at decreasing rates) with the number of periods as the efficiency gains obtained thanks to accelerated learning weight less, relative to all surplus generated in subsequent periods when estimated abilities and qualities would anyway have converged close to the true underlying values. Together these results suggest that for a policymaker committing to a given policy in period 1 and evaluating it based on net present values, such policy will look more beneficial welfare-wise when the time horizon used for the cost–benefit analysis is long enough to account for the full effects of learning on the economy.

A question that remains open at this point is whether the optimal policy found for 10 periods is still optimal when assuming a different number of periods. Table 5 shows results for our central case with 10 periods and from simulations with 5 and 13 periods. For each parametrization, we used the optimizer algorithm to find the corresponding optimal policy. Coherently with the arguments presented in previous sections, when reducing the number of periods the optimal policy features larger tax and subsidy rates (and therefore gets closer to the one under no information), the reason being that the simulated economy with fewer periods produces less information compared to simulations with 10 periods. As the number of periods grows, the model produces more information and thus approaches a condition (after the initial periods when the

Table 5 Sensitivity on the number of periods in the scenario with learning

	Optimal σ	Optimal τ	Entr. entry (%)	Work. entry (%)	Avg. a	Avg. q
5 Periods*	0.89	− 12.2	82.88	82.88	.57	.51
10 Periods	0.76	− 10.0	69.30	69.30	.57	.60
13 Periods*	0.70	− 9.0	69.01	68.87	.58	.62

An asterisk in column 1 means that the simulated population was halved (from 100 entrepreneurs to 50; from 1000 workers to 500) to reduce computation time

majority of learning is occurring) that is closer to full information. This is the reason why optimal policies feature smaller taxes and subsidies when assuming more periods.

The results in Table 5 therefore reassure that choosing a parametrization with a number of periods different from 10 would not change the main points stressed throughout previous sections. With learning, optimal taxes and subsidies still need to be significantly smaller (in absolute values) than with no information. This is true because the number of periods does not affect the optimality of a policy for the no-information case and therefore, for any number of periods considered, the optimal policy under no information always remains the one found using 10 periods ($\sigma = 3.85$ and $\tau = -75.7$). Some further testing (not reported for brevity) performed with 5, 8, 13 and 15 periods and assuming either $\sigma = 3.85$ and $\tau = -75.7$, or $\sigma = 3.85$ and $\tau = 0$, shows that it is always true that with learning the policy made only of a subsidy $\sigma = 3.85$ outperforms the policy also including the tax $\tau = -75.7$, in terms of total surplus produced, similarly to what is reported in Table 2 for 10 periods.

Table 5 highlights another aspect of introducing learning: the entry rate under optimal policies decreases with the number of periods. This might be counter-intuitive as one would expect early discovery of information produces more benefits over a longer time span. However, as the simulation keeps the tax and subsidy rates fixed in all periods, this means that larger entry in later periods also implies a loss in surplus as more low-ability and low-quality agents stay in the market. (Recall that the model is calibrated in a way that makes 50% entry optimal under full information.) An optimal policy has to strike a balance between having more information produced and having excessive entry, and Table 5 shows that such trade-off is (also) a function of the time horizon used for evaluating the effects of the policy.

Another set of sensitivity analyses relates to the variance of the information signals about abilities and qualities. We built two series of simulations that, in comparison with our central scenarios, only differ because the parameters σ_a and σ_q are either both increased by +0.10, or decreased by −0.10. We found that changing the standard deviation of the distributions from which the noise values for estimated ability and quality are drawn does not meaningfully affect the results. Also the relative magnitude of the welfare gains from the two alternative policies (subsidy only, or tax and subsidy) are very similar.

5 Conclusions

We developed a simulation exercise to mimic in an agent-based fashion a market with adverse selection both on the side of entrepreneurial entry and on the input market for workers, in this following the theoretical work of Boadway and Sato (2011). Our simulation results show that the introduction of learning over time can significantly affect the optimality of tax and subsidy policies. A policy made of a subsidy to labour and a simultaneous tax on entrepreneurs, which would be socially second-best under a no-information scenario, was shown to be inferior welfare-wise to a policy made only of a labour subsidy (where the level of such subsidy was optimal under no information). An optimal policy under learning was found to require much smaller subsidy to workers (and consequently, also a much smaller tax on profit) than under no information, as the entry rate of agents tends to be more reactive to the policy due to the dynamic effects of learning and better matching of high-quality workers with high-ability entrepreneurs.

The intensity of information asymmetries probably varies across industries, countries and times. Highly innovative sectors are likely to be more plagued as it takes time to evaluate the capacity of a technical employee to innovate, an example of this being the highly skewed distribution of patent applications among the population of inventors where most employees never patent anything or at most manage to patent once in a lifetime. In such cases, the economy is likely to be represented quite closely as a polar scenario with no information, with the exception of “star” scientists and entrepreneurs for whom a sizable number of observations may exist about their performance so that other agents may evaluate them accordingly. The latter observation points to a candidate extension of this research by assuming a segmented market with learning over ability/quality happening only above a threshold number of observed successes or failures.

The simulation model here presented is in many respects a simplified representation of reality, and multiple improvements could be devised to increase its external validity. We would like to mention one, in particular, which would represent the entrepreneurial market as an overlapping-generations framework with several cohorts of entrepreneurs and workers (similar, for example, to a model with symmetric uncertainty as in Ter-vio 2009 and d’Andria 2018b). Here the learning dynamics of our simulation would apply to each cohort individually, but all cohorts would simultaneously interact on the markets for credit and jobs. In such context, and contrary to the present model, cohort-specific policies (such as age-conditioned subsidies) could also be studied.

Acknowledgements This funding was supported by Joint Research Centre.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. In: *The quarterly journal of economics* (pp. 488–500).
- Bloomquist, K. M. (2006). A comparison of agent-based models of income tax evasion. *Social Science Computer Review*, 24(4), 411–425.
- Boadway, R., & Keen, M. (2006). Financing and taxing new firms under asymmetric information. *FinanzArchiv: Public Finance Analysis*, 62(4), 471–502.
- Boadway, R., & Sato, M. (2011). Entrepreneurship and asymmetric information in input markets. *International Tax and Public Finance*, 18(2), 166–192.
- Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. *The American Economic Review*, 95(4), 1300–1309.
- d'Andria, D. (2016). Taxation and Incentives to innovate: A principal-agent approach. *FinanzArchiv: Public Finance Analysis*, 72(1), 96–123.
- d'Andria, D. (2018a). Start-ups, venture capital financing, and tax policy under adverse selection. *FinanzArchiv: Public Finance Analysis*, 74, 462–480.
- d'Andria, D. (2018b). Superstars and mediocrities: A solution based on personal income taxation. *Journal of Economic Behavior and Organization*, 149, 459–463.
- d'Andria, D., & Savin, I. (2018). A Win–Win–Win? Motivating innovation in a knowledge economy with tax incentives. *Technological Forecasting and Social Change*, 127, 38–56.
- De Meza, D., & Webb, D. C. (1987). Too much investment: A problem of asymmetric information. *The Quarterly Journal of Economics*, 102(2), 281–292.
- Dietl, H. M., et al. (2013). Incentive effects of bonus taxes in a principal-agent model. *Journal of Economic Behavior and Organization*, 89, 93–104.
- Domar, E., & Musgrave, R. (1944). Effects of proportional taxes on risk-taking. *Quarterly Journal of Economics*, 59, 338–422.
- Feldman, M. (1987). Bayesian learning and convergence to rational expectations. *Journal of Mathematical Economics*, 16(3), 297–313.
- Gentry, W. M., & Hubbard, R. G. (2000). Tax policy and entrepreneurial entry. *The American economic review*, 90(2), 283287.
- Hashimzade, N., et al. (2015). The use of agent-based modelling to investigate tax compliance. *Economics of Governance*, 16(2), 143164.
- Henrekson, M., & Sanandaji, T. (2011). Entrepreneurship and the theory of taxation. *Small Business Economics*, 37(2), 167–185.
- Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization*, 7, 24–52.
- Judd, K. L. (2005). Computationally intensive analyses in economics.
- Keuschnigg, C., & Nielsen, S. B. (2003). Tax policy, venture capital, and entrepreneurship. *Journal of Public Economics*, 87(1), 175–203.
- Marimon, R. (1996). *Learning from learning in economics*. Fiesole: European University Institute Florence.
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of technical change*. Cambridge: Harvard University Press.
- Neugart, M. (2008). Labor market policy evaluation with ACE. *Journal of Economic Behavior and Organization*, 67(2), 418–430.
- Radulescu, D. (2012). The effects of a bonus tax on manager compensation and welfare. *FinanzArchiv: Public Finance Analysis*, 68(1), 116.
- Rosen, H. S. (2005). Entrepreneurship and taxation: Empirical evidence. In C. Keuschnigg (Ed.), *Venture capital, entrepreneurship and public policy* (pp. 251–279). Cambridge: MIT Press.
- Scheuer, F. (2013). Adverse selection in credit markets and regressive profit taxation. *Journal of Economic Theory*, 148(4), 1333–1360.
- Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle* (Vol. 55). Piscataway: Transaction publishers.
- Shapiro, C., & Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *The American Economic Review*, 74(3), 433–444.
- Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3), 393–410.

- Tervio, M. (2009). Superstars and mediocrities: Market failure in the discovery of talent. *The Review of Economic Studies*, 76(2), 829–850.
- Tesfatsion, L. (2001). Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, 25(3), 419–457.
- Tesfatsion, L. (2003). Agent-based computational economics: Modeling economies as complex adaptive systems. *Information Sciences*, 149(4), 262–268.
- Tesfatsion, L., & Judd, K. L. (2006). *Handbook of computational economics: agent-based computational economics* (Vol. 2). Amsterdam: Elsevier.
- Warner, G., et al. (2015). Modeling tax evasion with genetic algorithms. *Economics of Governance*, 16(2), 165–178.
- Weiss, A. (1980). Job queues and layoffs in labor markets with flexible wages. *Journal of Political Economy*, 88(3), 526–538.
- Yildizoglu, M. (2002). Competing R&D strategies in an evolutionary industry model. *Computational Economics*, 19(1), 51–65.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.