Towson University Department of Economics **Working Paper Series**



Working Paper No. 2010-10

Emergent Pareto-Levy Distributed Returns to Research in a Multi-Agent Model of Endogenous Technical Change

By Michael D. Makowsky and David M. Levy

April, 2010

© 2010 by Authors. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Emergent Pareto-Levy Distributed Returns to Research in a Multi-Agent Model of Endogenous Technical Change

Michael D. Makowsky* Department of Economics Towson University

David M. Levy Department of Economics George Mason University

April 10, 2010

Abstract

We build a multi-agent model of endogenous technical change in which heterogeneous investments in patented knowledge generate Pareto-Levy and lognormal distributed returns to investment in research from very weak distributional assumptions. Firms produce a homogenous good and a public stock of knowledge accumulates from the expired patents of privately produced knowledge. Increasing returns to scale are derivative of endogenously produced technology, but the market remains competitive due to imperfect information and costly household search. The interaction of heterogeneous knowledge, research investment, revenues, and search outcomes across agents endogenously generates the empirically observed but seemingly idiosyncratic Pareto-Levy and lognormal mixture distribution of market returns. These distributional characteristics have ramifications for endogenous growth models given the importance of extreme values and market leaders in technological advancement. Average growth rates in the model have a global maximum at a finite, non-zero patent length. The distribution of growth rates increases with patent length.

JEL Codes: C63, L11, O33, D83

Keywords: patents, endogenous growth, increasing returns to scale, price dispersion, search, heterogeneous agents

*We would like to thank Robert Axtell and seminar participants in the Department of Social Complexity, George Mason University. We thank Omar Al-Ubaydli for a detailed list of comments. Makowsky thanks the Towson College of Business and Economics for summer financial support. Please send correspondence to <u>mikemakowsky@gmail.com</u>.

Theories of endogenous technical change built with knowledge serving as a non-rival input into productivity and, in turn, as a source of increasing returns to scale, have served to model exponential growth and offer a better understanding of disparate of rates of growth observed across countries (Grossman and Helpman 1994; Romer 1994). The capacity to cope with increasing returns to scale, however, motivated the abandonment of price taking perfect competition, and the allowance of market power within firms (Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1992). It should not be surprising that, given this reliance on knowledge inputs and market power, that intellectual property rights, or patents, have become a major topic of exploration in theories of endogenous growth (Horowitz and Lai 1996; Futagami and Iwaisako 2003; O'Donoghue and Zweimüller 2004; Iwaisako and Futagami 2007).¹ Models incorporating patents into theories of endogenous growth, however, have not accounted for the peculiar distributional properties of the returns to innovation. We offer an alternative modeling strategy that allows for endogenous technical change, is characterized by long run increasing returns to scale, and emerges a distribution of revenues across firms that is best characterized as a Pareto-Levy and lognormal mixture distribution, and is often dominated by a small number of extreme values. This peculiar mixture distribution is similar to those observed in patent revenue return research (Scherer, Harhoff et al. 2000; Silverberga and Verspage 2007).²

Patents bring the necessary market power to firms that seek to obtain monopoly rents from their excludable private knowledge. This excludable private knowledge, however, also engenders heterogeneity across firms that are all producing with differing knowledge inputs. Heterogeneous knowledge quickly leads to heterogeneity in productive capacity, marginal products of standard (rival) inputs, and prices. Such a world is considerably less tractable for traditional modeling, and is typically inhospitable to decentralized competition. The structural imposition of monopolistic competition in

¹ To varying degrees, the models proposed in this literature are built using the foundations laid out by Aghion and Howitt (1992), Grossman and Helpman (1991), and Judd (1985).

² Mixed distributions with extreme values have also been offered as a tractable representation of Knightian uncertainty and a challenging environment for policy (Epstein and Wang 1994).

the form of a continuum of goods produced by firms returns us to more tractable territory, but comes at a cost. With a continuum of goods in demand, and each firm producing a unique good that cannot be perfectly substituted for by goods produced by competing firms, our potentially Schumpeterian landscape looks considerably less destructive. Imperfect substitution, long thought to be necessary to allow many firms to exist in an industry with increasing returns, attenuates the consequence of discoveries which would be explosive in a world with perfect substitutes. The monopolistic competition model, governed by the Law of One Price, retains the representative firm by allowing for heterogeneous goods. We provide a model using the exact opposite: a set of heterogeneous firms competing to produce and sell a single homogenous good, each offering the good to consumers at their own unique price.

There is considerable evidence that the returns to research are highly skewed, with distributions dominated by extreme values. Research into these returns has used a variety of creative datasets, including citation records, initial public stock offerings (IPOs), and self-reported revenue returns to patents (Harhoff et al. 1998; Harhoff et al. 1999). The most appropriate statistical distribution for the characterization of the returns seems to be some combination of the lognormal and Pareto-Levy distributions (Scherer, Harhoff et al. 2000; Silverberga and Verspage 2007). The overall distribution within the empirical work is best characterized by a lognormal distribution with outliers in the upper tail. However, the upper tail of the distribution, particularly when looking at IPO data, is better characterized by the Pareto-Levy power law distribution. Such power law distributions are not unheard of in market competition and concentration data. Axtell (2001) finds that the size of firms, in terms of individuals employed, is Zipf distributed in the United States. Within power law distributions, the upper tail accounts for an extraordinary share of the distribution's value. Models that account for growth derivative of technical innovation that leverage some form of market power stand to benefit from either including such features of the returns to research or, preferably, generating them endogenously (Luttmer 2007). Concerns about the importance of the distribution of research outcomes, in particular of the upper tail and outliers have been recently

3

expressed (Silverberga and Verspage 2007). As they note such Pareto power law distributions might not even have first moments, something which has severe implications for risk analysis.

The aims of this paper are three-fold. First, we seek to build a model characterized by the long run increasing returns to scale and exponential growth properties of existing models of endogenous technical change and growth. Second, we abandon the traditional monopolistic competition model, and replace it with a model of competitive firms producing a homogenous good in a market characterized by price dispersion. Third, we simulate the model under a variety of parameterizations and examine the distributional properties of returns to investment in research. In doing so, we find that the distributions of returns to research in our model take on a mixture character, taking on a lognormal share in the lower quartiles while exhibiting Pareto-Levy power law properties in the upper quartiles. We also test the impact of the key parameters of the model, patent length and search costs, on average growth rates across large batches of simulation experiments.

2 The Model

We construct a multi-agent model of endogenous growth that includes elements prominent in O'Donoghue and Zweimüller (2004) and Iwaisako and Futagami (2007). Within this model we create a market composed of heterogeneous, individually autonomous households and firms that make decisions in accordance with their type, unique information set and personal history, and the rules that govern their behavior. Like the model presented in O'Donoghue and Zweimüller (2004), our model is composed of two sectors, one in which technology investment and innovation are possible and one in which innovation is not possible, with inputs of only labor and capital. Individual, technology enabled, firms produce a homogenous quality primary good (q) while an aggregated non-technical sector (NTS) produces a secondary good (x). Households supply labor to both sectors, collect wages, earn uniform returns to shares of rents paid to capital, and maximize a universal utility function by purchasing a combination of x and q.

4

The model is always composed of a fixed quantity of households, and as such growth within the model is not dependent on the exogenous increase in labor.³

Time in the model occurs in discrete steps and substeps. Sets of agents (organized by type) are activated in a fixed schedule, but within each set, agents are activated in a randomized order. While firms are effectively acting simultaneously, households are not. A household may purchase the last of a firm's inventory or fill its final hiring slot. Potential order effects add to the complexity of model outcomes, but constant randomizing of activation order prevents model artifacts (Axtell 2001).

While agents, within their types, are homogenous in capacity, exogenous parameterization, and behavioral rules, they each face a world with costly, imperfect, and heterogeneous information. Households search for both lower prices and higher wages, seeking to maximize their consumptive bundle, while being constrained by a finite amount of time to be split between wage earning labor and search, and the ensuing time expenditures associated with searching the market. Firms, on the other hand, face the uncertainty of a research process that may or may not yield a competitive increase in excludable knowledge as well as a marketplace of consumers that may or may not discover them as a low price provider of goods. They respond to these uncertainties by making decisions regarding research investment predicated on simple heuristics and limited information. Given the complexity of the relationships between households and firms, the non-technology sector (NTS) is governed by a number of simplifying assumptions that grant the model additional tractability. The NTS operates as a single agent in the model, hires all who are willing to work for its offered wage, and always meets the sum of its market orders.

The labor supply in the model is fixed, but capital grows as a set fraction of the previous time steps's total productivity. Growth, nonetheless, is driven by technical innovation. As within a Schumpeterian model of creative destruction (Aghion and Howitt 1992), innovation is motivated by desire to both gain monopoly rents and avoid bankruptcy. In this manner both the carrot and the stick are applied every step of the

³ For a discussion of scale dependent vs. scale independent models of endogenous growth, see Eicher and Turnovsky (1999)

model: success in research and development leads to lower production costs, greater rents, and more customers, whereas failure leads to higher prices, fewer customers, and brings the firm one step closer to closing its doors. The prospect of permanent failure is one of the salient features of working with a competitive market for a homogenous good. In a monopolistic competition model, where goods exist along a continuum, there is no prospect for complete failure to attract customers. With agents searching over a set of producers offering a homogenous good, a firm with inferior productive technology will be unable to offer a competitive price and will be more likely to be passed over by potential customers. This market remains competitive⁴, as opposed to collapsing to monopoly, because of price dispersion and costly search, which allows second-best firms to attract sufficient customers to retain positive profits, or at the very least manageable losses that can endured in the short run (Levy and Makowsky 2010). Further, the expiring of patents and the subsequent sharing of previously private knowledge allows for turnover in who stands as the technology leader (Grossman and Helpman 1991). In reality, it is not just profit, but the prospect of losses and bankruptcy that motivates investment in research and development.

In contrast to traditional general equilibrium models, there is no social planner maximizing agent utility, nor a Walrasian auctioneer finding market clearing prices. Each agent, governed by type (firm, household) specific rules, is autonomous. From the thousands of interacting, decision-making agents emerge aggregate trends in research investment, technology, growth, wages, profits, and market concentration. Agents are myopic, backward looking, and absent any sophisticated strategy. They are governed by a strictly bounded rationality and costly information, but nonetheless manage to prosper in what are often rapidly growing economies.

2.1 The Multi-Agent Computational Model

⁴ In contrast to the bulk of the existing literature, Hellwig and Irme (2001) build a general equilibrium model of endogenous technical change that includes competitive markets, though their unique equilibrium is characterized by a low steady-state growth rate.

The model is composed of two vectors of agents, households

(1)
$$H = [1, 2, \dots n]$$
$$i \in H$$

and firms

(2)
$$F = [1, 2, \dots m]$$
$$j \in F$$

where each household (i) purchases q_i^j units from the firm, j_i^* , offering the lowest price known to her during time step *t*. All variables that are not exogenously set vary across time steps. For ease of explication, we will not include *t* as a subscript except when previous time steps (*t* - 1) are relevant.

Firms produce the primary good, Q_j , using inputs of labor, L_j , capital, K_j , and knowledge, A_j , where knowledge is composed of public, G, and private, R_j , knowledge:

(3)
$$Q_{j} = A_{j}^{\gamma} K_{j}^{\alpha} L_{j}^{\beta} \qquad \forall j$$
$$A_{j} = G + R_{j}$$

subject to the costs of production, C_j , including the wages, w_j , paid to employees; rent paid to capital, r, and the investment in research and development, S_j .

(4)
$$C_{j} = w_{j}L_{j} + rK_{j} + S_{j} \qquad \forall j$$
$$\pi_{j} = p_{j}Q_{j} - C_{j}$$

Profits, π , are a function of Q_j sold at price, p_j , and C_j . Firms post unique prices in the market equal to lagged AC, such that $p_{j,t} = AC_{j,t-1}$. Each firm also posts its own wage in the labor market in the hopes of attracting prospective employees. Firms set their wages equal to the monetized marginal product of labor from the previous turn (MPL $\cdot p_{j,t-1} l$):

(5)
$$w_{j,t} = \beta K_{j,t-1}^{\alpha} L_{j,t-1}^{\beta-1} \cdot p_{j,t-1} \qquad \forall j$$

Given this wage rate, firms establish a maximum number of employees they are willing to employ by engaging in standard cost minimization of the production function given⁵

$$C_{j}(K_{j}, L_{j} | Q_{j,t-1}, A_{j}, w_{j}, r, \alpha, \beta)$$

such that $L_{j}^{\max} = Q_{j,t-1} / A_{j}^{1/\alpha + \beta} r\beta / w\alpha^{\alpha/\alpha + \beta}$

During each step firms engage in research from which knowledge returns are uncertain, generating a quantity of private knowledge, or patent, $y_{j,t}$ that is temporarily excludable for Φ time steps, and contributes to a summed portfolio of private knowledge stocks, $R_{j,t} = \sum_{\phi=0}^{\Phi-1} y_{j,t-\phi}$. The process of research and development is modeled as an exponential probability function, dependent on the firm's investment, S_j , its current

portfolio of private knowledge, R_{j,t}, and the existing stock of public knowledge, G_t:

(6)
$$y_{j,t} = \left(\frac{-S_j}{G_t + R_{j,t-1}}\right) \log(Z) + R_{j,t-1} \qquad \forall j$$

where Z is a unit rectangular variate. The ratio of investment S_j to private knowledge created with each patent, $y_{j,t}$, is declining as the existing stock of knowledge, $G_t + R_{j,t-1}$, grows. This choice to model the costs of innovation as increasing with the existing stock of knowledge is based on the empirical observation that the costs of patents have been increasing over time (Kortum 1993). Firms choose unique research investments S_j equal to their investment from the previous turn adjusted by factor χ , where

(7)

$$S_{j,t} = TR_{j,t-1} \cdot \chi_{j,t} \quad \forall j$$

$$\chi_{j,t} = \chi_{j,t-1} + v_{j,t}$$
where
$$\begin{cases} v_{j,t} = v_{j,t-1} & \text{if } \pi_{t-1} \ge \pi_{t-2} \\ v_{j,t} = -1 \cdot v_{j,t-1} & \text{if } \pi_{t-1} < \pi_{t-2} \end{cases}$$

This research investment adjustment rule entails a simple profit seeking heuristic on behalf of the firm, with which each individual firm gropes towards an investment

⁵ If a firm fails to sell a single unit of q, but does not go out of business, they use the mean q_j among firms still in business and use that to establish an average cost and marginal product for setting their price and wage.

procedure that increases profits. The increment of change, $v_{t=0}$, is exogenously set parameter uniform across firms. Firms myopically grope towards greater profits, switching directions whenever their previous turn resulted in reduced profits.

Each firm's stock of private knowledge, $R_{j,t}$, is a rolling portfolio of patented knowledge. Each step, the oldest patent, $y_{j,t-\Phi}$, expires. The expired patent of greatest magnitude is added to the public knowledge stock, $G_t = G_{t-1} + \max y_{1,t-\Phi}, \dots y_{m,t-\Phi}$. Research results in more efficient production that is rewarded by greater profits and greater prospects for long run survival in the marketplace. This, in turn, incentivizes the long run contribution to the public stock of knowledge and ideas in the form of expired patents which lead to long run growth. At the same time, the rolling expiration of patents allows for turnover in private knowledge leadership at any given time step.

Once a firm has conducted its research, set its price and wage, and hired its employees, it can establish a profit maximizing quantity to produce and sell. Capital, K, in the model is available from exogenous pool at price $r(Q_{t-1}, \Psi)$, where Ψ is a fraction of total productivity in the model from the previous time step, and r() is the marginal product of the said fraction in the previous time step. The profit maximizing quantity to

be sold is
$$q_j^{\text{max}} = A_j L^{\beta} (\alpha L^{\beta} p_j A_j / r)^{\frac{1}{1/\alpha - 1}}$$

Agent search occurs within each time step t, in sub-steps $\tau=1...m$ where each increment of τ represents an act of search by the agent.⁶ Households first search over the set of wages offered by firms, then search over the set of prices posted for q. Their search activities are governed by simple income maximizing and cost minimizing search functions based on a desire to continue searching so long as the expected increase in the highest known wage, w_i^* , or decreases in the lowest known price, p_i^* , will result in a net increase in purchasing capacity given the cost of an additional sub-step of search, $w_i \varsigma$, where ς is the amount of an agent's time endowment expended by an act of search. In both wage and price search, the decision variable is the number of search actions, τ , that constitute the fixed sample size

⁶ There are *m* firms, and thus *m* prices over which to potentially search. If the cost of a unit of search, Δh , equaled zero, all agents would continue search until τ equaled *m*.

that households decide prior to the first discovered price. Both wage and price search result in a fixed sample size. Households assume a non-degenerate uniform distribution of wages F(w) on $[\underline{w}, \overline{w}]$ and maximize the expected total income (highest found wage earned over the time remaining after search).

(12)
$$E(M_{i,\tau+1}) = (1 - c\tau) \left[\underline{w} - \int_{\underline{w}}^{\overline{w}} F(w)^{\tau} dw \right] \qquad \forall i$$

Households similarly assume a non-degenerate distribution of prices F(p) on $[\underline{p}, \overline{p}]$ and minimize the expected total cost (cost of purchasing q_{t-1} plus cost of search).

(13)
$$E(C_{i,\tau+1}) = q_{i,t-1} \left\lfloor \underline{p} - \int_{\underline{p}}^{\overline{p}} F(p)^{\tau} dp \right\rfloor \qquad \forall i$$

For the sake of simplicity, households will assume the center of the price and wage distributions are simply the wage or price they chose the previous time step (

 $\underline{a}_i = \frac{1}{2}a_{i,t-1}^*; \overline{a}_i = \frac{3}{2}a_{i,t-1}^* \quad a \in w, p$). Additionally, the wage paid by the NTS is known to

each household without cost.

Each household *i* searches over the wage set Θ , where $\Theta_{i,\tau}$ is the subset of wages known to household *i* after τ search efforts.

(8)

$$\Theta \equiv w^{1}...w^{m}, w^{NTS}$$

$$\Theta_{i,\tau} \subset \Theta$$

$$w_{i}^{j^{*}} = \max w_{i}^{j} | w_{i}^{j} \in \Theta_{i,\tau} \qquad \forall i$$

In addition to their wage, each household receives a uniform dividend, *d*, of the rent outlaid by firms and the NTS to capital inputs and any positive profits accumulated by firms.

After searching for a wage, each household *i* then searches over the price set Ω , where $\Omega_{i,\tau}$ is the subset of prices known to household *i* after τ search efforts.

(9)

$$\Omega \equiv p^{1}...p^{m}$$

$$\Omega_{i,\tau} \subset \Omega$$

$$p_{i}^{j^{*}} = \min p_{i}^{j} | p_{i}^{j} \in \Omega_{i,\tau} \qquad \forall i$$

Once households have executed their searches and found a lowest known price and highest known wage, they maximize a constant elasticity of substitution (CES) utility function,

(10)
$$U_i = (q_i^{\lambda} + x_i^{\lambda})^{1/\lambda}$$

For a given wage rate and price, the optimal quantities of q and x are

(11)
$$q_i^* = 1/p^{1/1-\lambda} \cdot M / p^{-\lambda/1-\lambda} + \eta^{-\lambda/1-\lambda}$$
$$x_i^* = 1/\eta^{1/1-\lambda} \cdot M / p^{-\lambda/1-\lambda} + \eta^{-\lambda/1-\lambda}$$

Where the total income of the household, M_i , is a function of the household's wage, the number of sub steps spent searching, and the costs of search, ς , and its dividend from capital rents and firm profits, d, such that $M_i = w_i(1 - \tau_i \varsigma) + d$

The non-technical sector (NTS) acts as a single agent. It sets the price for x, η , price based on the average cost of production from the previous time step

$$\eta = C_{t-1}^{NTS} / \left(\sum_{h=1}^{n} x_h + \sum_{b=1}^{z} x_z \right)$$
. The NTS pays a wage to its employees equal to the marginal

product of labor from the previous step, $W_{NTS,t} = \beta K_{NTS,t-1}^{\alpha} L_{NTS,t-1}^{\beta-1} \cdot p_{NTS,t-1}$.

At the end of each step, all firms are evaluated for potential bankruptcy. All firms for which costs exceed revenues ($\pi_j < 0$) must borrow funds to remain solvent. This debt accumulates across steps. Bankruptcy occurs when accumulated debt exceeds the limit of *B*,

(14)
$$B = \Gamma \cdot \max(\pi_{i,t=0} \dots \pi_{i,t})$$

B is a function of the greatest profits previously realized by any firm in a single step, adjusted by an exogenous multiplier, Γ .

2.2 Simulation Steps and Sub-step Ordering

Our model is characterized by a schedule of agent decisions and model events. This schedule plays out in a series of steps and sub-steps.⁷ A run of the model is constituted by an initialization (t = 0) followed by a set number of model steps (t=1...T), during which every agent is activated in random order, as arranged by the model sub-steps. The sub-steps are ordered as follows:

- Each firm, j=1...m, sets its offered wage (see Equation 5) and its offered price for primary goods.
- 2) All expired patents are made public; the largest patent value is added to the cumulative stock of public knowledge. All sub-superior knowledge disappears.
- 3) Each firm conducts research (see Equation 6).
- 4) The NTS sets both its offered wage and the price for secondary goods.
- 5) Households, i = 1...n, are activated in random order and execute τ searches over the set of all available wages. Households are always aware (without cost) of the NTS wage. Once they have decided on their fixed sample size, the first wage in their discovery set is their employer from the previous time step (see Equation 8).
- 6) Given the fruits of their research investment, their posted price and wages, the price of capital, and the number of employees they were able to hire, firms establish a profit-maximizing limit to the amount of the primary good they will produce.
- 7) Households are activated in random order and execute τ searches over the set of all available prices. Once they have decided on their fixed sample size, the first price in their discovery set is their seller from the previous time step (see Equation 9). Once search is concluded, the household maximizes its utility function, choosing an optimal bundle of *q* and *x*. If the firm offering the lowest known price to the household is unable to fulfill the entire desired quantity of *q*, the household purchases the remaining amount from the firm with the second lowest price firm. For tractability, the household will not seek out a third firm if the quantity desired

⁷ The model is written in Java using the MASON agent modeling library (Luke et al. 2005). The step/substep construct is built into the MASON model scheduling system.

is still not met. Once a firm has orders for $q_j = q_j^{\text{max}}$, it is withdrawn from the set of unknown prices Ω .

- 8) Having received all of their market orders, firms will acquire the amount of capital necessary to produce q_i and fulfill all existing market orders.
- 9) If a firm is unable to procure any market orders, it may go bankrupt. Bankruptcy results when a firm's outstanding debt is greater than the quantity that is available in the commercial loan market (see equation 14). In the model simulations executed in this paper, firms were exempt from bankruptcy rules during the first ten simulation steps, allowing firms to adapt to initialized conditions.

3 Simulation Results

We ran the model under a variety of patent length and search cost parameterizations, with 400 time steps constituting a run. In experiments where we simulate the model for a single run, we ran it with 4000 households and 200 firms. For larger batches where we made comparisons across runs, we ran it with 2000 households and 100 firms. The key exogenously set parameters are summarized in Table 1.⁸

Our emphasis, in this paper, is on the distributional properties that are observable across firms. These properties, however, are of limited interest if they do not occur in a model of endogenous, exponential growth. Figure 1 plots log Q, where $Q = \sum q_j$, over time in a single run of the model. All firms produce the same good, price differences at any time tick the model are simply the consequences of positive search costs and technology differences, so there is no particular merit to working with "real" output.

⁸Model results are, unsurprisingly, sensitive to the specification of γ (the output elasticity of A_j) and Γ (the maximum debt firms can take on). This sensitivity is the result of their influence on market concentration. Specifically, large values of γ and small values of Γ result in faster rates of attrition, driving the model towards monopoly.

 Table 1 Model Parameters

Parameter	Context/Related Function	Value	
M	Starting number of firms	100, 200	
N	Number of customers	2000, 4000	
α, β	$Q_j = (G + R_j)^{\gamma} K_j^{\alpha} L_j^{\beta}$	0.5	
γ	$Q_j = (G + R_j)^{\gamma} K_j^{\alpha} L_j^{\beta}$	0.15	
λ	$U_i = (q_i^{\lambda} + x_i^{\lambda})^{1/\lambda}$	-0.1	
$G_{t=0}$	Initial public stock	1	
Г	Loanable funds multiplier	5	
Ψ	$r(\mathbf{Q}_{t-1}, \Psi)$	0.05	
$V_{j,t=0}$	$\chi_{j,t} = \chi_{j,t-1} + \nu_{j,t}$	0.002	
ς^{\dagger}	Search cost	[0.00001, 0.0001]	

[†] The total sub step time endowment for an agent is 1. As such when search costs, ς , equal 0.00005, that is equivalent of 0.005% of their time endowment, meaning it takes 0.005% of an agent's sub step time endowment to engage in another act of search.

Tracking the growth of log Q over time, we observe two distinct periods. In the early time steps of the run, we see an "organizational" period in the model, which typically (but not always) concludes within the first 50 steps, within which growth is erratic, often characterized by large swings up and down, as firms grope towards profitable strategies and unprofitable firms go bankrupt and exit the model. Eventually the model settles into steady growth trend, which is ostensibly a random walk, but does sometimes exhibit small, semi-regular cycles. Growth observed in the model is exponential and consistent, and largely parameter insensitive (given minimal returns to research and elasticities of output). Given this type of growth, we can proceed to focus on the distributional properties observed within the set of active firms.



Figure 1 Log Q over Time. Patent length = 12, Search Costs = 0.01%. The model was initiated with 4000 households and 200 firms.

3.2 Outcomes Across Firms Within a Single Run

Firms in our model are homogeneous ex ante and heterogeneous ex post. Given the randomness of research outcomes and cost constraints faced by searching households, each firm experiences its own unique history of research outcomes, sales, and profitability. Firms are confronted with two levels of uncertainty: research uncertainty and commercial uncertainty.⁹ They do not know the stock of excludable knowledge that their research investment will bear, nor do they know whether customers will successfully find them *even if* they are able offer a relatively low price. These

⁹ This is not unlike the three types of uncertainty (technical, commercial, and financial) laid out in Scherer et al. (2000)

uncertainties result in differing research investments, private knowledge stocks, posted prices, q_i sold, and revenues generated.

Research investment (as a percentage of revenues), χ_j , is the manner in which the individual firms most directly respond to their own unique history. In Figure 2 we present the histograms of the distributions of χ_j across all firms at time step 400 in several runs of the model, parameterized with different patent lengths and search costs. As is expected, firms are on average investing more when patents last longer. Perhaps more interesting, however, is that we observe a far greater variety of investment rates, including much larger ones, under higher search costs. When search costs are very low, households sample a larger fraction of the offered prices, and their purchasing of goods will more directly follow the distribution of knowledge and prices. When search costs are higher, however, there is greater randomness in the model. The correlation between sales and prices is far murkier when search costs are high, and in turn firms have vastly different experiences in their personally observed connection between research investment and profit. Greater variety of experienced histories leads to a greater variety of firm behavior.

Figure 3 presents histograms of logged total revenue, TR_{j} , across all firms at time step 400 in several runs of the model, parameterized with different patent lengths and search costs. The symmetry of log TR_{j} in many of the histograms in Figure 3 gives the appearance of a distribution that is potentially log normal. In fact, the distributions of log TR_{j} passes the Shapiro-Wilk normality test at the better than 1% level (p<0.01 for all fifteen runs, see Table 2). While the bulk of the distributions are likely lognormal, there are several visual characteristics that warrant interest. Many of the distributions have prominent outliers, especially in the upper tail, and exhibit fat tails more generally. Twelve of the fifteen histograms qualify as leptokurtotic, six of which have a raw kurtosis greater than 6.0.

16



Figure 2. Histograms of research investment percentage across firms, organized by Patent Length (4 to 20, vertically) and Search Cost (0.001%, 0.005%, and 0.01%, horizontally), at t = 400. Each subfigure is from a single run of the model.



Figure 3 Distribution of Total Revenue, organized by Patent Length (4 to 20, vertically) and Search Cost (0.001%, 0.005%, and 0.01%, horizontally), at t = 400. Each subfigure is from a single run of the model.

In Figure 4 we chart the rank, N, of each observation, where the rank can be interpreted as number of other observations within the same model simulation run that are of equal or greater value than TR_j at step t = 400, using data from the single simulation run with patent length of 16 and search costs of 0.005%. The shape of the results in Figure 4 bears a strong resemblance to what was found by Scherer et al. (2000) in their study of the value (in Deutschmarks) of German patents from 1977 to 1995, duplicated here in Figure 5. In both the simulation and patent data, the lower observations within the distribution are concave to the origin, but the higher value observations take a more linear relationship between log value and log rank. A log-linear relationship between rank and value is indicative of a potential power law nature of the distribution of values. Both the simulation results and German patent return observations also include outliers significantly beyond the rest of the distribution.



Fig 4 Log Rank over Log TR, Single Run of the Model. Patent length = 16, Search costs = 0.001%



Figure 5 Rank over Estimated Value (Log-Log scale), from Scherer, Harhoff, et al. (2000)

Again, we simulated the model under a range of parameterizations. In Figure 6 we cluster similar charts of the Log N_j over Log TR_j , with each graph charting the results at time step 400 of a single run of the model, under different combinations of patent length and search cost parameterizations (charts of simulations with longer patent lengths can be found in the appendix). The concave to the origin shape and increased linearity in the upper tail are fairly consistent. Further, the top ranked observation, and sometimes several more, is frequently a significant outlier from the rest of the observations.



Figure 6 Log Rank over Log TR, organized by Patent Length (4 to 20, vertically) and Search Cost (0.001%, 0.005%, and 0.01%, horizontally), at t = 400. See Appendix Figure A for longer patent lengths.

Scherer, et al. (2000) analyzes the Pareto-Levy, or power law, distribution parameters of the patent data with simple regression analysis using the log-linear modeling function

(15)
$$N = k(\operatorname{TR}_{j})^{-\varpi} \log N = \log k - \varpi \log \operatorname{TR}_{j}$$

where N is the rank of the return on Investment observation, TR_j is total revenue, and k and ω are parameters. Absolute values of ω greater than or equal to one are indicative of a Pareto-Levy distribution.

In most of their patent data, Scherer et al do not find a log linear fit between TR_i and rank across the full body of observations, instead finding that the bulk of the distribution is better characterized as lognormal. However, they do observe a much closer to log-linear fit in the upper tail of the data. We find similar results in our simulation data across a variety of parameterizations. The Shapiro-Wilk test results suggest that the full distributions appear to be lognormal. Simple ordinary least squares analysis of Pareto-Levy model parameters, when regressed over the full distribution, offers further support, with values less than one. However, if we isolate the third and fourth quartiles of the distributions, the resulting coefficients correspond to a Pareto-Levy power law distribution. The results of the ω slope parameter are included in Table 2 for the overall distribution (column 4), the third quartile (column 5) and in the fourth quartile (column 6) for each combination of patent length and search cost parameter combination tested. In all 15 model simulations, the regression of Pareto-Levy parameters on the full distribution resulted in slope coefficients less than one. Regression on the third quartile observations resulted in much larger slope coefficients, with $|\omega\rangle > 1$ in four of the specifications. Regression on the fourth quartile observations produced slope coefficients $|\omega\rangle > 1|$ in six of the specifications. Of the fifteen specifications, eight produced $|\omega\rangle > 1|$ in at least one of the upper two quartiles.

Examined in tandem, the plots from Figure 6 and the Pareto-Levy slope coefficients for the upper quartiles of the simulation distribution paint a telling picture. The plots can be visually broken down into three components. The first is the lowermiddle portion with a shallow slope and weakly concave shape. The second, is a steep, flat region usually in the upper-middle portion of the distribution. Third, but not always, is a small number of extreme values that represent significant outliers from the rest of the distribution. Even in the simulations whose upper quartiles did not have $|\omega| > 1|$, we can often visually identify a portion of the distribution characterized by significant steepness, with the strong possibility of Pareto-Levy characteristics. It is our view that considered together, the Shapiro-Wilks tests, Pareto-Levy coefficients, and general visual shape of the data plots, are evidence of a distributional pattern emergent from the model that is strikingly similar in character to that observed in the Scherer, et al. German patent data – specifically a mixture of lognormal and Pareto-Levy distributions of revenues across firms accruing rents from temporarily excludable knowledge stocks.

		(1)	(2)	(3)	(4)	(5)
				OLS	OLS	OLS
Patent	Search		S-W		$ \sigma $	$ \sigma $
	Costs	Kurtosis*	(p	$ \sigma $	3rd	4th
Length	COStS		value)		quartile	quartile
4	0.01%	4.05	< 0.01	0.399	0.788	0.468
4	0.05%	4.28	< 0.01	0.228	1.011	1.591
4	0.1%	3.46	< 0.01	0.256	0.565	1.863
8	0.01%	6.90	< 0.01	0.497	1.158	0.699
8	0.05%	3.67	< 0.01	0.281	0.616	1.465
8	0.1%	2.06	< 0.01	0.077	0.061	0.200
12	0.01%	7.99	< 0.01	0.349	0.617	1.041
12	0.05%	4.83	< 0.01	0.287	0.555	1.866
12	0.1%	6.35	< 0.01	0.08	0.081	0.067
16	0.01%	6.81	< 0.01	0.281	1.351	1.222
16	0.05%	2.21	< 0.01	0.07	0.036	0.197
16	0.1%	9.02	< 0.01	0.114	0.262	0.081
20	0.01%	7.24	< 0.01	0.519	1.074	0.634
20	0.05%	2.30	< 0.01	0.068	0.077	0.135
20	0.1%	2.06	< 0.01	0.069	0.037	0.298

Table 2 Log TR: Kurtosis, Shapiro-Wilk Tests, and Pareto-Levy RegressionCoefficients at t=400

*Raw Kurtosis (kurtosis = 3.0 at normality).

3.3 Growth Rates

In the previous section and simulation experiment, we inspected the distributions of results across firms within single runs of the model, each with differing patent length and search cost parameterizations. In this section, we take a different approach, simulating the model thousands of times, and inspecting the how various outcome properties change. In turn, we are not looking at individual firms, but rather outcomes that are aggregated across all firms from each instantiation of the model. In this simulation experiment, we ran the model 4000 times, with 2000 households and 100 firms in each run, for 400 time steps in each run. All parameters besides patent length [1,2,...40] and search costs [0.001%, 0.002%,... 0.01%] are held constant and are identical to those reported in Table 1. Our outcome of concern is the average growth of aggregate Q from step 200 to step

400 in the model. We are purposely allowing the model to settle into the "steady state" portion of its simulation history before tracking growth rates. Figure 7 shows the quadratic fits of the observations – average growth rates, mapped over patent length, color/shape coded by search cost. All quadratic fits are statistically significant at the 1% level (p < 0.01 for each regressor), which is not surprising given that we are modeling data from a simulation whose only varying parameter and its square are being used as regressors. In the lower four search cost parameterizations, average output is increasing, though at a decreasing rate, with patent length. At the highest search cost level ($\varsigma = 0.01\%$), however, the impact of patent length has a global maximum in the interior of the parameter domain tested, declining at a significant rate beyond it. Futagami and Iwaisako (2003) arrived at a similar result in their dynamic analysis of patents in their endogenous growth model, within which social welfare was maximized by patents with finite length (see also Horowitz and Lai 1996).



Figure 7 Average Growth rate from t = 200 to t = 400 over Patent Length, organized by Search Cost. Search costs = 0.002% (red), 0.004% (orange), 0.006% (green), 0.008% (blue), and 0.0010% (black). Lines are quadratic regression fits over each search cost subset of observations. N = 4000.

We are also interested in the distribution of outcomes from the model across multiple runs of the model. Specifically, we were interested in the consistency of growth outcomes from simulation to simulation under identical model parameterizations. To test this, we simulated model across different patent lengths, 1 to 40, and under two search cost paramterizations: low search costs (0.002%) and high search costs (0.01%). Each possible combination of patent length and search cost was simulated fifty times, with identical simulation parameters to the previous experiment (400 times steps, 100 firms, 2000 households), resulting again in 4000 simulation runs. Each of the 80 sets of 50 simulations resulted in a distribution of average growth rates from step 200 to step 400 in each simulation.¹⁰

In Figure 8 we chart the eighty different observed mean average growth rates, as well as the variance and kurtosis of the distributions. In the first graph of Figure 8, we see that the mean of average growth rates follows patterns at the low and high search costs similar to what we observe in Figure 7.¹¹ In the low search cost set of observations, the mean of average growth rates are increasing over the entire range of patent lengths. In the high search cost set of observations, there is again a global maximum in the interior of the patent lengths tested.

The variance of average growth rates across simulations is increasing with patent length in both search cost parameterizations but is increasing at a much higher rate under low search costs. Even at the highest levels, however, the variance remains small relative to the means of the distributions. The raw kurtosis of the individual distributions of average growth rates may in fact be the most interesting. The median kurtosis of the distributions was 3.19, almost exactly that of a normal distribution. However, in simulations with high search costs, the median kurtosis of these distributions was 10.55,

¹⁰ Average growth rate =
$$\left(\sum_{j=1}^{m} q_{j,t=400} / \sum_{j=1}^{m} q_{j,t=201}\right)^{1/200} - 1$$

¹¹ We can also see in Figure 8 the often truly prodigious rates of growth in the model. It has been our experience that the model results remain salient and clear, despite our relatively limited size of 100 firms and 2000 customers, when growth is rapid. When using single runs, such as in earlier sections of the paper, we have the luxury of a larger agent sets. Running the model thousands of times per experiment is less practical with larger numbers of agents, however. Larger scale investigation remains for future research.

and often exceeded 40. In the high search cost runs, the correlation to patent length was U-shaped; the largest kurtosis distributions occurred under the highest and lowest patent lengths, while the lowest kurtosis distributions occurred when patents were of an intermediate length (Figure 8, lower right). Under low search costs, the effect of patent length on the kurtosis rates is greatly muted, and may even be the inverse. Longer patents in the model increase the mean growth rate, at least early on and up to a global maximum. Perhaps more importantly, however, both the highest and lowest patent lengths can increase the possibility of extreme outcomes, dependent on search costs. This would appear to suggest the merits of patents of an intermediate (finite) length, which are not just increasing the mean rate of growth, but also its predictability. The extreme kurtosis of growth rates in the shortest parameter and longest patent length regimes suggests an extremely peaked distribution, with the majority of observations clustered around the middle, but also outliers exceptional in number and deviation from the mean. Growth rates within each individual simulated "history" in model are pulled from these fat-tailed distributions. In deference to Nicholas Taleb, "black swans" abound, particularly when search costs are high and secrets are either fleeting or forever.



Figure 8. Distribution of average growth rates across

runs from Steps 200 through 400, over Patent length. Clockwise from upper left: a) Mean of Growth Rates

b) Variance of Growth Rates c) Kurtosis of Growth Rates. Ouadratic fits are included in all subfigures.

graphs are the results from 4000 model simulations.

Two search costs parameterizations are charted: 0.0002% (black/dashed/circles) and 0.001% (red/solid/triangles). Each observation is the mean/variance/kurtosis for a set of 50 runs. The



4 Discussion

The distribution of returns to research investment, based on our analysis, appears to be a mixture of a lognormal and the Pareto-Levy power law distribution. This especially curious mixture distribution is nowhere assumed in the model. The economic consequences of perfect substitution combined with positive search costs allow an occasional innovation to revolutionize the industry without creating a single firm. We have not attenuated the consequences of such an innovation by the ad hoc assumption of imperfect substitution. The only distributional shapes assumed in the model are an exponential distribution of returns to dollars invested in research and the uniform price and wage distributions assumed by searching households. It is this conspicuous absence from the model structure that makes the emergence of this mechanically idiosyncratic and

empirically observed distributional shape so compelling. Such a distribution is characterized both by its skewness and proclivity towards producing outliers that dominate the rest of the distribution. William Feller (1950) begins the elementary half of his celebrated pair of volumes on probability theory with distributions without moments. He demonstrates how an innocent looking process is likely to have a single event which is as large as the all the other events combined. He expresses his concern that we have to believe that something can happen before we see it:

In practice such a phenomenon would be attributed to an "experimental error" or be discarded as "outlier." It is difficult to see what one does not expect to see. (Feller 1950, p. 91).

In our account of endogenous technical change distributions without moments are not censored by our intuition. What might such an outlying event look like in technical change? Perhaps the electronic computer, but economic intuition is uncomfortable with the naïve question of asking what would 2010 GDP look like if we were to evaluate computational expenditures in 1950 prices (Nordhaus 2001). We find it more tractable to simply trim out 1950 computational expenditures by putting them in 2010 prices.

If we take the approach to Knightian uncertainty suggested by Epstein and Wang (1994) in which familiar distributions are mixed with something strange, then trimming "outliers" can take Knightian uncertainty and transform it into seemingly well-behaved risk. Regardless of approach, the point remains that "outliers" are not observations that can be dismissed. Outliers are driving observed growth and are, in turn, the source of "fat tailed" distributions of growth observed in Figure 8. When an outlier event occurs, it changes the entire trajectory of the simulated economy. It is our opinion that models of endogenous growth would do well to account for such distributional properties, either in their assumptions or their outcomes (Silverberga and Verspage 2007).

Thinking about models of industrial organization, in 1949 George Stigler saw something very remarkable in Edward Chamberlin's theory of monopolistic competition.

28

With the abolition of an "industry" there was no way to keep the "group" from being the whole of the economy:

It is perfectly possible, on Chamberlin's picture of economic life, that the group contain only one firm, or, on the contrary, that it include all of the firms in the economy. This latter possibility can readily follow from the asymmetry of substitution relationships among firms: taking any one product as our point of departure, each substitute has in turn its substitutes, so that the adjacent cross-elasticities may not diminish, and even increase, as we move farther away from the "base" firm in some technological or geographical sense. (Stigler 1949, p. 15)

This property would suggest that a technological development in one firm could disrupt firms arbitrarily distant. This explosive general equilibrium property, however, was long seen as a defect in monopolistic competition which would be later tamed with a "preference for diversity" which kept firms safely in their niches. It was the tamed monopolistic competition models which would become the basis for models of increasing returns,

5 Concluding Remarks

It would be useful to extend and test the model in a computing environment that allowed for larger scale simulations. Given the importance of highly skewed distributions, larger agent pools could have important ramifications for growth rates and the incentive to participate in innovation races. Patents in our model are greatly simplified. Future work would benefit from introducing more sophisticated intellectual property rights, including both "breadth" and length, as well as a continuum of imitation and obsolescence. This paper is largely concerned with the unusual distributional characteristics of the returns to research and the need for their realization in endogenous technical change models. More generally, the nature of the "optimal" patent length is given only cursory attention here, and would benefit from finer grain analysis in future work.

29

References

- Aghion, P. and P. Howitt (1992). "A Model of Growth Through Creative Destruction." <u>Econometrica</u> **60**(2): 323-351.
- Axtell, R. (2001). Effects of Interaction Topology and Activation Regime in Several Multi-Agent Systems. <u>Multi-Agent-Based Simulation</u>. G.Goos, J. Hartmanis and J. v. Leeuwen. Berlin / Heidelberg, Springer: 33-48.
- Axtell, R. (2001). "Zipf Distribution of U.S. Firm Sizes." Science.
- Eicher, T. S. and S. J. Turnovsky (1999). "Non-Scale Models of Economic Growth." <u>The Economic</u> Journal **109**(457): 394-415.
- Epstein, L. G. and T. Wang (1994). "Asset Pricing under Knightian Uncertainty." <u>Econometrica</u> **62**: 283-322.
- Feller, W. (1950). *An introduction to probability and its applications*. New York.
- Futagami, K. and T. Iwaisako (2003). "Patent Policy in an Endogenous Growth Model." <u>Journal of</u> <u>Economics</u> **78**(3): 239-258.
- Grossman, G. M. and E. Helpman (1991). "Quality Ladders in the Theory of Growth." <u>Review of</u> <u>Economic Studies</u> **58**(1): 43-61.
- Grossman, G. M. and E. Helpman (1994). "Endogenous Innovation in the Theory of Growth." Journal of Economic Perspectives 8(1): 23-44.
- Harhoff, D., F. Narin, et al. (1999). "Citation Frequency and the Value of Patented Inventions." <u>The Review of Economics and Statistics</u> **81**(3): 511-515.
- Harhoff, D., F. M. Scherer, et al. (1998). "Exploring the tail of patent value distributions." Working Paper, Center for European Economic Research.
- Hellwig, M. and A. Irme (2001). "Endogenous Technical Change in a Competitive Economy." Journal of Economic Theory 101, 1 39 (2001) **101**: 1-39.
- Horowitz, A. W. and E. L. C. Lai (1996). "Patent Length and the Rate of Innovation." <u>International</u> <u>Economic Review</u> **37**(4): 785-801.
- Iwaisako, T. and K. Futagami (2007). "Dynamic analysis of patent policy in an endogenous growth model." Journal of Economic Theory **132**(3): 306-334.
- Judd, K. L. (1985). "On the Performance of Patents." Econometrica **53**(3): 567-586.
- Kortum, S. (1993). "Equilibrium R&D and the Patent-R&D Ratio: U.S. Evidence." <u>American</u> <u>Economic Review</u> **83**(2): 450-57.
- Levy, D. and M. D. Makowsky (2010). "Price Dispersion and Increasing Returns to Scale." <u>Journal</u> of Economic Behavior and Organization **73**: 406-417.
- Luke, S., C. Cioffi-Revilla, et al. (2005). "MASON: A Multiagent Simulation Environment." <u>SIMULATION</u> **81**(7): 517.
- Luttmer, E. G. J. (2007). "Selection, Growth, and the Size Distribution of Firms." <u>Quarterly</u> Journal of Economics **122**(3): 1103 -1144.
- Nordhaus, W. D. (2001). "The Progress of Computing." <u>Yale Cowles Foundation Discussion Paper</u> **No. 1324**.
- O'Donoghue, T. and J. Zweimüller (2004). "Patents in a Model of Endogenous Growth." <u>Journal</u> of Economic Growth **9**(1): 81-123.

- Romer, P. M. (1990). "Endogenous Technological Change." <u>The Journal of Political Economy</u> 98(5, Part 2: The Problem of Development: A Conference of the Institute for the Study of Free Enterprise Systems): S71-S102.
- Romer, P. M. (1994). "The Origins of Endogenous Growth." <u>The Journal of Economic</u> <u>Perspectives</u> **8**(1): 3-22.
- Scherer, F. M., D. Harhoff, et al. (2000). "Uncertainty and the size distribution of rewards from innovation." Journal of Evolutionary Economics **10**(1): 175-200.
- Silverberga, G. and B. Verspage (2007). "The size distribution of innovations revisited: An application of extreme value statistics to citation and value measures of patent significance." Journal of Econometrics **139**(2): 318-339
- Stigler, G. J. (1949). <u>Five lectures on economic problems</u>. London, New York, London School of Economics and Political Science; Longmans Green.

Appendix



Figure A1. Log Rank over Log TR, organized by Patent Length (24 to 40, vertically) and Search Cost (0.1%, 0.05%, and 0.1%, horizontally), at t = 400.