Los Alamos Unclassified Report LA-UR 98-3904 (This version: November 19, 1998)
Approved for public release; distribution is unlimited

A simple model for labor market dynamics (DRAFT, comments welcome)

Kai Nagel

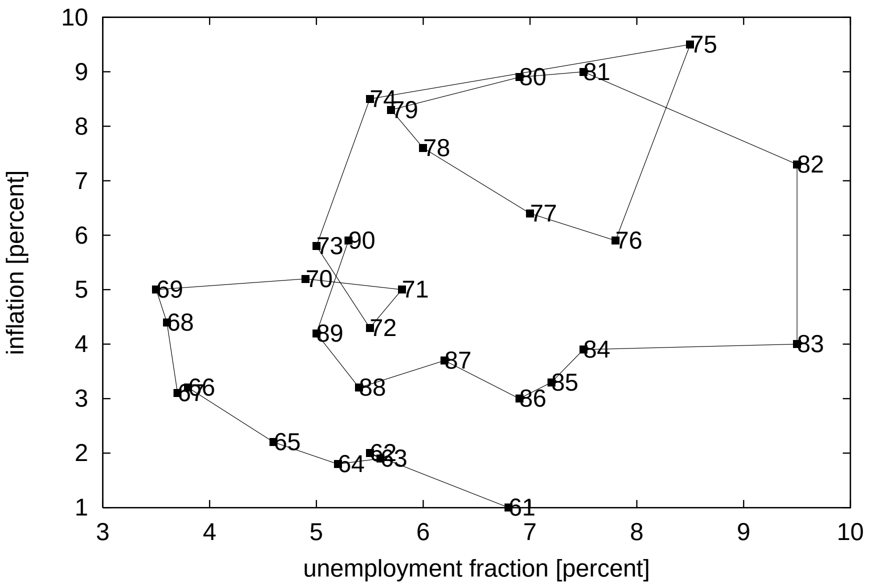

A simple model for labor market dynamics

Kai Nagel

Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe NM 87501,
U.S.A., kai@santafe.edu

Los Alamos National Laboratory,
TSA-Division/Simulation Applications,
Mail Stop M997, Los Alamos NM 87545, U.S.A.

Abstract

We present a simple dynamic model which can explain cycles in the labor market of a national economy. The main variable of the model is the quotient between wage, w, and the price p of the product that is produced for this wage. If this quotient, w/p, is larger than one, workers are inefficient for companies and they try to get rid of them, by letting them off and by reducing wages. If w/p is smaller than one, each additional worker makes additional money for a company and thus companies attempt to attrack workers by increasing their wages. This leads to oscillations around w/p which generate behavior in the relation between inflation and unemployment that resembles historical data. An agent-based implementation of the model is presented.

I. INTRODUCTION

Macroeconomics, i.e. the theory of how macroscopic variables of economics such as inflation, unemployment, interest rates, exchange rates, etc., relate to each other, is today an as interesting problem as ever. For example, for the European Community, people are attempting to predict outcomes for the common currency. In the United States, people are debating if it is possible to sustain low unemployment together with low inflation. Yet, journals that should be informed on such matters such as the *Economist* admit that economists are divided on these issues [1].

In this paper, we want to look at a simplified economy that still captures some of the above variables. In particular, we will look at the *labor market* side of the economy, i.e. at mechanisms that generate wage inflation and unemployment. We will start with some data which shows the historic relation between unemployment and inflation (Sec. II). This relation does not follow a simple curve, rather it moves in dynamic "cycles". Next, we will provide the economics textbook explanation (Sec. III). The textbook explanation of the phenomenon is quasi-static, that is, it assumes that an external shock brings the system out of equilibrium, and from there it relaxes back to a new equilibrium point *different from the old one*, and a succession of external shocks followed by relaxation generates the circle. We will then proceed to offer an alternative explanation of the data (Sec. IV). This alternative explanation will not be quasi-static, but entirely dynamic. This means that it will not need the existence of external shocks to explain the cycle. From there, we will proceed to show an agent-based computer implementation of our alternative explanation (Sec. V), and results of this implementation (Secs. VI to VIII). The paper is concluded by a discussion and a summary.

II. HISTORICAL DATA

Fig. 1 shows historical data of the relation between inflation and unemployment. The data shows cyclical behavior:

- Starting from high unemployment and low inflation, unemployment goes down while inflation remains low.
- Eventually, the data curves upward, that is, unemployment remains low while inflation goes up.
- With high inflation, unemployment increases.
- Finally, inflation decreases while unemployment remains high, which completes the circle.

III. TEXTBOOK BEHAVIOR AND TEXTBOOK EXPLANATION

The textbook explanation [2] for the above data is composed of two parts, which are labeled "traditional inflation theory" and "modern inflation theory".

"Traditional inflation theory" assumes that there is inertial inflation, that is, a certain level of inflation is built into everybody's expectations and behavior. There need to be external shocks to move inertial inflation from one value to another one. External shocks can for example be:

- Demand-pull inflation: Aggregate demand increases. As a result, prices go up more than the inertial rate because of higher demand, and wages go up as a consequence in order to attract more workers;
- Cost-push inflation: Costs increase, for example for raw materials or for labor, and prices go up in consequence.

"Modern inflation theory" starts out with the Phillips curve, see Fig. 2. The Phillips curve shows a trade-off between unemployment and inflation. When inflation is low, unemployment is high, and vice versa. Now, this is clearly inconsistent with the inertia assumption, which stated that the economy can operate at different levels of inflation without changes in any other variables, such as unemployment. It is also inconsistent with data (see Fig. 1), which moves in circles rather than following a Phillips-like curve.

The textbook explanation for this is that the Phillips curve is only valid in the short run; in the long run, there is a "natural level of unemployment" towards which each economy eventually relaxes. Thus, when increases of aggregate demand increase price inflation, then unemployment should decrease for a certain amount of time, but should eventually relax back to the original level while inflation remains at the higher level everybody got used to in the meantime. Under this scenario, the Phillips curve would shift upwards. And when there is a decrease of aggregate demand, then price inflation should go down, followed by higher unemployment, which eventually relaxes back to the "natural level", while inflation remains low.

Thus, the explanation of the "Phillips curl" is as follows: When demand for products is high, companies need more people, unemployment goes down, wages go up. People then get used to the higher inflation, and so it remains high even when demand and unemployment go back to normal. Eventually, there will be a period of lower (external) demand, companies need less people, unemployment goes up, wages go down. People then get used to lower inflation, and so it remains low even when demand goes back to normal.

IV. AN ALTERNATIVE, *DYNAMIC* EXPLANATION

The textbook explanation is quasi-static. An external shock happens, and the system relaxes back to equilibrium. Another shock happens, and the system relaxes to another equilibrium point. Yet, Fig. 1 looks inherently dynamic, that is, if one looks at the process during time, one cannot say where the equilibrium points should be.

Also, external shocks need to be alternating between upward and downward in order to generate the circling motion. Yet, there is no obvious reason why this should be the case.

Last, the implication behind plotting wage inflation vs. unemployment is that it follows a similar relationship as price and supply: When supply of workers (i.e. unemployment) is low, the price for human labor (i.e. wage) should be high. Yet, note that this would imply a relation between unemployment and wage itself, not between unemployment and wage *inflation*.

All this implies that it might be worthwhile to look for alternative explanations, explanations which explicitely take the *dynamics* of the process into account. In the following, we will offer such an alternative, *dynamic* explanation.

Let us look at simplified economy where people work at factories and receive a certain amount of wage for their work. They produce a certain amount of a good; the value (i.e. the price) of their production is their productivity. From this alone, we can make the following general deductions:

(i) When wages are higher than productivity, companies lose money for each additional worker they employ.

(ii) When wages are lower than productivity, companies make more money for each additional worker they employ.

Let us in addition assume that workers move to the company that pays the highest wages. Now, the following happens:

- In situation (i), it is good for companies to decrease wages. Decreasing wages has two effects: (a) If the company pays lower wages than other companies, the workers will move away. This will cut losses for the company. (b) The company will eventually become profitable again.

- In situation (ii), companies have an interest to increase wages to attract more workers.

Next, we assume that there is *inertia* in the way companies set their wages, that is, it takes time to go from a period of high wage increases to a period of low wage increases.

Last, we assume that in times when wages are higher than productivity, there will be unemployment.

These four assumptions lead to a dynamic economic cycle as in Fig. 1 (see also Fig. 3 for results of a computer implementation of these rules, to be described later in this paper):

- When wages are higher than productivity, unemployment is high. At the same time, companies decrease their wage increases, leading to low wage inflation.

- Eventually, wages will be lower than productivity. In Fig. 3, this is represented by the wage level becoming lower than the price level. Unemployment goes down because companies are now making money. Wage inflation will still be low because of inertia.

- Since now companies want to attract workers, wage inflation goes up while unemployment remains low.

- Eventually, wages will again be higher than productivity. Unemployment will be high again while wage inflation remains high because of inertia. – And this closes the dynamic cycle.

V. AN AGENT-BASED IMPLEMENTATION

In this section, we will show a computer implementation which generates the dynamics outlined in the last section. The computer implementation will be agent-based, which means that we will have individual workers and individual companies that make decisions. The advantage of this is that we will be able to show that no "long-term planning" by any of the economic entities will be necessary in order to generate the dynamics – simple short-term heuristic rules like "go where the wage is highest" or "copy a successful company" will be sufficient. Agent-based models have been used for economics simulations before, see [3–8] for examples.

First, we want to give some intuitive insight into the model. A technically complete description will follow later.

The model consists of firms and agents. In general, during a period of the simulation, the following things happen:

- **Working.** Agents work at firms, i.e. they produce a certain amount of a good in return for a certain amount of wage.

- **Adaptation of agents.** Agents adapt their behavior. In this paper, we assume that their only choice is where they work. Agents will perform a variant of simple hill-climbing towards the company that pays the highest wage.

- **Adaptation of firms.** Firms adapt their behavior. In this paper, we assume that they can only adjust wage, i.e. we assume that prices are set by a central agency. Companies will perform a variant of simple hill-climbing towards being the company that has the highest gains. Gains are defined as "the value of all produced goods minus the wages paid".

  We also assume that companies that do not earn money go "bankrupt" with a certain probability, in which case they do not employ any people in the next period.

- **Bookkeeping.** We assume that goods are completely perishable, i.e. the amount of goods each company has in stock after one period is over is set to zero. Also, we reset all money accounts to zero at the end of each iteration.

All steps together represent one period of the "game"; we perform many iterations (periods) of the game.

Obviously, there is a wide range of possibilities on how to exactly specify all these rules. We expect the general dynamics to be fairly robust with respect to the specific implementation. Nevertheless, in the following we want to provide a precise description of the rules.

The model consists of N_f firms and N_a agents. As stated above, agents work, then agents adapt, then firms adapt, and then some general "bookkeeping" is done.

- **Working:** The "working" part of the rules is straightforward: Each agent, i, works at a firm, $j(i)$. Firms can reject applicants because of bankruptcy (see below). If the applicant is not rejected, she produces a certain amount, A, of a good, and receives a certain amount, $w(j)$, of wage for this. We assume A as a constant.

- **Agent adaptation:** Next, we look at adaptation. In each time step, agents adapt towards a higher wage. Agents remember the wage that they got; for example, if the company rejected the worker, this last wage is zero. In each time step, each agent picks a company at random, and it that company pays a higher wage than her last wage, the agent switches employers.

- **Firm adaptation:** Firms adapt towards higher gains. We define gains as the value of all goods that a firm produced minus the wages paid for the production. Every T_{fadapt}th time step, each firm picks another firm at random, and if, in the present period, that other firm has made higher gains than itself, it copies the wage parameter, w_j, and the firm-specific wage inflation, Δw_j (see below) from that other firm.

- **Wage inflation:** A company j remembers two parameters about its wage: the wage itself, w_j, and the firm-specific wage inflation, Δw_j. As long as no copying happens as described in "firm adaptation", the wage inflation follows a random walk, i.e. in each time step companies change Δw_j randomly either to $\Delta w_j + \eta$ or to $\Delta w_j - \eta$. Wages w_j are then adapted to $w_j \cdot (1 + \Delta w_j)$.

  Note that, since the firm-specific wage inflation Δw_j follows is a random walk, it has inertia, which means that it cannot change by more then η per time step except by copying from another firm.

If not stated otherwise, for this paper the number of agents $N_f = N_a = 50$, the wage inflation noise $\eta = 10^{-11}$ and firm adaptation every $T_{fadapt} = 10^2$ periods. We also set price inflation to a constant 10^{-7} per period. The model does not pretend to be calibrated to reality and it may even impossible to calibrate it because agents and firms in the model are much dumber then in reality, thus needing more time to react. Yet, one can imagine that $Y = 10^6$ time steps correspond to one year. The reason for this is the following: Yearly (wage) inflation is defined as

$$I_w = \frac{\overline{w}(t+Y) - \overline{w}(t)}{w(t)} ,$$

where $\overline{w}(t)$ denotes the wage level (average wage) at time t. For $Y = 10^6$, we obtain realistic inflation values between zero and 20%.

All updates are done in parallel, although it really only matters for firm adaptation. Also, data points for the time-dependent plots are averaged over $T_{av} = 10\,000$ iterations, and results for the inflation-vs-unemployment plots are averaged over $T_{av} = 150\,000$ iterations. The inflation values as shown in the figures are calculated as

$$I_w = \frac{\overline{w}(t) - \overline{w}(t - T_{av})}{w(t - T_{av})} \cdot \frac{Y}{T_{av}} .$$

VI. RESULTS

Fig. 3 shows an example of the dynamics of the above model. As shown in Fig. 3 a, the average wage level tracks the average price level, which is set to a constant inflation here. Yet, the wage level oscillates around the price level. Fig. 3 b shows wage and price inflation, i.e. the first derivatives of Fig. 3 a. Fig. 3 c shows how unemployment behaves in the same simulation. Clearly, when the price level is higher than the wage level, we have full employment, whereas in the opposite case, we have unemployment. Fig. 4 shows what happens if one plots unemployment vs. inflation, as in the economics textbook. For clarity, only values for the first 2,350,000 time steps are plotted here. One gets clockwise cycling of the curve, with the following intuition:

- Workers in general move to the company which pays the highest wages.
- When wages are lower than prices, companies make additional money with each additional employee. Since workers move to the company which pays the highest wage, it is the company with the highest wages that attracts most of the workers (see Fig. 5 a). Since, in consequence, it makes the highest gains (see Fig. 5 b), the other companies attempt to copy it, leading to a "race" for the highest wage. – In this regime, wages go up faster than prices while unemployment is at zero.
- When wages become higher than prices (this is easiest to see in the "normalized" wage as in Fig. 3 d), companies can go "bankrupt" as described Sec. V. This leads to high unemployment while wage inflation is still high because of "inertia".

- When wages are higher than prices, companies with more employees have higher losses. This leads, in effect, to a competition for the lowest wages, trying to "push" people to other companies (see again Fig. 5 a and b). Thus, in this regime, wage inflation slowly decreases while unemployment stays high.

- Eventually, wages become lower than prices again. Unemployment goes to small values, while wage inflation is still low.

- From here, the cycle starts over again.

Clearly, the model generates a cycle similar to that observed in the real world data.

VII. MORE RESULTS

In the above Sec. VI, we have described how the model generates wage and unemployment oscillations similar to unemployment-inflation cycle observed in real world data. Closer inpection of Fig. 3 yields the following observations:

- Macroscopically, the mechanism is actually summarized quite easily. Let us define the average wage level as the average of all wages actually paid, i.e.

$$\overline{w} := \frac{1}{N_{empl}} \sum_i^N w(j(i)) \ ,$$

  where i is the index for the agents, $j(i)$ is the company where agent i works, $w(j)$ is the wage of company j, and N_{empl} is the number of agents who actually receive wages. Let us also define $\overline{p}$ as the average price level, which corresponds to the uniform price in our simplified model. Let us then denote by $W := \overline{w}/\overline{p}$ the normalized wage level. Let us denote by $\tilde{I}$ the *reduced* wage inflation, i.e. wage inflation minus price inflation; also $\dot{W} \approx \tilde{I}$, with corrections of order $w - p$. Clearly, what happens is roughly that

$$\frac{d^2 W}{dt} \approx \frac{d\tilde{I}}{dt} \approx \begin{cases} const1 & \text{if } W < 0 \\ -const2 & \text{if } W > 0 \end{cases} .$$

- On top of that, the amplitude of W decreases, probably linearly (see Fig. 6).

- The size of the largest firm follows the oscillations. Agents move to the company with the highest wage, although it takes some time until they have settled there.[1] As long

[1]Clearly, this is an exponential relaxation since the average number of agents per time step that switch to the company with the highest wage is proportional to the number of "remaining" agents, $N - n$, where N is the overall number of agents and n is the number of agents that already work for the company with the highest wage. Thus, $n = N - A\ \exp[-t/\tau]$, and we obtain $\tau \approx 35$ from a test simulation. That means that, for $N = 50$, after approx. 150 time steps all agents have moved to the company that pays the highest wages.

as $W < 1$, the company that pays the highest wages is happy about that large work force since it means high company gains, and so as a result, in this regime workers work for the same company. When $W > 1$, however, every additional employee means additional losses. A company that is paying high wages will thus quickly copy the lower wages of another company, thus "pushing" its employees away.

This can be seen in Fig. 5 a, which shows the share of the workforce working at the largest company, i.e.

$$S := \frac{1}{T} \sum_{t=1}^{T} \frac{\max_j N_{empl}(j,t)}{N} , \tag{1}$$

where $T = 10\,000$ is the averaging time, N is the number of workers, $N_{empl}(j,t)$ is the number of workers at company j at time t. Thus, $\max_j N_{empl}(j,t)$ is just the number of employees working at the company that is largest at iteration t. In words: In each time step, we find the workforce share of the largest company, and average over those values over T iterations. – Clearly, the largest firm employs about 80% of the workforce during low wage periods, and about 22% of the workforce during high wage periods.

The gains, in Fig. 5 b, vary accordingly. Gains (or losses), as defined earlier in this text, equal the value of a day's production minus the wages paid for this. The average gains as in the figure is simply the average of these values over all companies and all T iterations of a measurement period. The average gains *of the largest company*, G_l, is obtained in correspondance to Eq. (1), that is, in each time step, the largest company is found, its gains are added to the sum, and after T iterations, the result is divided by T. The result is also divided by the price level, $\overline{p}$. That is

$$G_l := \frac{1}{T \cdot \overline{p}} \sum^{T} g(j_{max}(t), t) ,$$

where $j_{max}(t)$ is the largest company at iteration t, and $g(j,t)$ are the gains of company j at iteration t. – Clearly, companies earn money during low wage periods and lose money during high wage periods. During low wage periods, the largest company collects all the work force and thus has high gains, whereas during high wage periods, the largest company is much smaller and thus losses are distributed across more firms and are thus lower for the largest firm.

VIII. UNSTABLE FIXPOINT

The decreasing amplitude of the oscillations indicates that the cycle will eventually settle into a stable fix point. This means that the dynamics would need to be permanently driven by external shocks, otherwise the cycle would die out. Although it is quite possible that this is indeed what happens in reality, it makes sense to look for circumstances under which the fix point becomes unstable, giving way to a periodic or possibly chaotic attractor. A way to achieve this is to *only* copy the wage *adjustment*, Δw_j, in the firms adaptation step, i.e. to not copy the wage itself. A possible justification for this would be that the model situation is actually very artificial. In reality, different firms have different productivities, and it is therefore very difficult to compare wage levels directly. For example, one company may try to produce a product with many low-skilled workers, whereas another one may try to produce the same product with fewer but high-skilled workers. It is, however, comparatively straightforward to compare wage *increases* between companies.

A possible result looks like in Fig. 7. In order to show the instability, the simulation here is started with wages equal prices, i.e. $w_j = p$, and and wage inflation equal price inflation, i.e. $\Delta w_j = I_p$. Clearly, the oscillations keep getting bigger, and the above mechanism gets started all by itself. Yet, the wage adaption mechanism is considerably less stable than it was before. This is clear: For example, during periods when it is good for companies to pay high wages, just copying the wage inflation parameter from your competitor is not enough to catch up if you have a lower wage than she has.

IX. DISCUSSION

First, let us mention that price inflation is not necessary as an element of the described dynamics. Setting the price inflation to zero would simply mean to reduce the m_i by a constant amount; with respect to the ratio $\overline{w}/\overline{p}$ nothing would change.

Second, we need to point out that in our results we look at wage inflation, while price inflation remains fixed during the simulations. Yet, in the real world data (Fig. 1), price inflation is used. Does that not render our model invalid? We think not: It seems plausible to us that prices (and therefore price inflation) follow wages (and therefore wage inflation) with some delay. And as soon as $\overline{p}$ lags behind $\overline{w}$, the same oscillating behavior as before would occur. Such an effect could be added to the model, but would make it more complicated. Also, the principal dynamic variable of our model is not wage or price, but their quotient w/p. And we suggest that if one could measure w/p in reality, one would see a curve similar to Fig. 1. The problem is thus that we need to make the model more complex not in order to explain the dynamics, but in order to relate the dynamics back to what can be measured.

Next, note that unemployment does not have a strong influence on the dynamics. In the version of the model presented, the only influence is that people who become unemployed have a higher incentive to switch to a new employer (since their last wage in that situation is zero). This means that in this model unemployment effectivly "paints" some aspects of the dynamics, but is not involved directly. This in turn means that any description of the inflation-unemployment curve in terms of unemployment is misleading – their coupling is much more indirect.

We tried to simulate the textbook explanation directly, without success. Also, the explanation our model offers is dynamically considerably simpler since "external" shocks are not necessary. Yet clearly, simulations in general cannot settle the question *which* microscopic rule-set is the correct one since simulations can only tell us that certain models generate a certain dynamic behavior, not if they are the *only* models generating this behavior.

Yet, we find that our model is conceptually clearer than the textbook explanation. To recast, our model just states that, (i) when wages are lower than productivity, companies make more money for each additional worker they empley and thus there is a competition for the highest wage, (ii) when wages are higher than productivity, companies lose money for each additional worker they employ and thus they decrease (real) wages, (iii) inflation possesses inertia, and (iv) there is unemployment when companies lose money. Ingredients (i) – (iii) together lead to the cycle; (iv) just "paints" the dynamics.

One has to note, though, that the quotient between wage level and price level is a quantity that will be difficult to observe in practice. Companies can judge if they do something wrong because then they do not make money. But in a highly specialized economy, workers have no way to determine if the buying power of their wage is higher or lower than their productivity. It is one of the advantages of simulations that such "hidden" quantities can be extracted easily.

This also has a political implication. For workers, the only way to know if their wages correspond to their productivity is to know if their company makes money or not. Thus, during periods of high company earnings, unions have a vested interest in knowing the "true" gains of a company. Short of this, the only thing they can do is to keep fighting for higher wages, until the company reacts by deliberately shrinking its output. As long as that does not happen, they can probably assume that, for each added worker, the company still makes more money than it pays for her.

Also, in view of the high unemployment in Europe vis-a-vis the United States and Britan, one could speculate that this model supports the view that Europe is in the regime where the general wage level is above the general price level, and thus unemployment is high. A closer look might suggest that this is true only for part of the labor market: The fact that unemployment is high *and* company earnings are high suggests that in some segment of the labor market, the value of the productivity of employees is lower than their wages, and so companies push these employees out, whereas in other segments, employees have higher productivity than they cost, thus leading to company earnings. It is easy to see that a high minimum wage (or in general a wage structure that does not reflect productivity) can exactly lead to this: People with low qualifications are too expensive and get pushed out of the market, whereas people with high qualifications are relatively cheap and lead to high company earnings. – It is hopefully clear that in view of the simplicity of our model, remarks such as the last paragraphs are highly speculative.

X. SUMMARY

We have presented a simple dynamic model which can explain cycles in the labour market of a national economy. The model represents the labor market side of a very simplified economy. Workers produce a certain amount of a product with value/price p and receive a certain wage w in return. Workers continuously search for companies that pay higher wages; companies copy other companies that are more successful; and randomness is introduced by making each companies individual wage inflation follow a random walk.

The main dynamic variable of this model is the quotient of wage and price, w/p. If $w/p > 1$, wages are higher than the value of the product that the wages pay for, and companies loose money. For that reason, companies decrease (real) wages in this situation. If $w/p < 1$, wages are lower than the value of the product that they pay for, and thus companies make more money for each additional worker that they hire. In consequence, companies increase (real) wages in this situation in order to attract more workers. These elements together lead to a dynamic cycle, where w/p oscillates around one.

If one in addition assumes that companies with $w/p > 1$ actively get rid of workers and thus generate unemployment, then this model generates a cycle in the inflation-unemployment plane that resembles historical data.

ACKNOWLEDGMENTS

The original, very early motivation to work in this area came from Nick Vriend, at that time postdoc at the Santa Fe Institute. Many thanks go to Alan Kirman, who allowed me to present a (too?) early version of this work at a workshop at the European University Institute. Maya Paczuski was interested in this work from near the beginning and considerably helped it along the way; she and Jörg Esser gave invaluable suggestions in structuring the material. This work has been done in part at Los Alamos National Laboratory, which is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36.

REFERENCES

[1] An awfully big adventure. EMU Survey, The Economist, April 1998.

[2] Paul A. Samuelson and William D. Nordhaus. *Economics.* McGraw-Hill, Inc., New York, 14th edition, 1992.

[3] R.G. Palmer, W. Brian Arthur, John H. Holland, Blake LeBaron, and Paul Tayler. Artificial economic life: a simple model of a stockmarket. *Physica D*, 75:264–274, 1994.

[4] B. Arthur. Inductive reasoning, bounded rationality, and the bar problem. *American Economic Review (Papers and Proceedings)*, 84:406–411, 1994.

[5] R. Axelrod. Advancing the art of simulation in the social sciences. Santa Fe Institute Working Paper 97-05-048, `www.santafe.edu`.

[6] A. Leijonhufvud. Towards a not-too-rational macroeconomics. *Southern Economic Journal*, 50(1), 1993.

[7] TRANSIMS, TRansportation ANalysis and SIMulation System, Los Alamos National Laboratory, Los Alamos, U.S.A. See `transims.tsasa.lanl.gov`.

[8] Ken Steiglitz, Michael L. Honig, and Leonard M. Cohen. A computational market model based on individual action. In Scott Clearwater, editor, *Market-based control: A paradigm for distributed resource allocation*, chapter 1. World Scientific, Hong Kong, 1996.

FIGURES

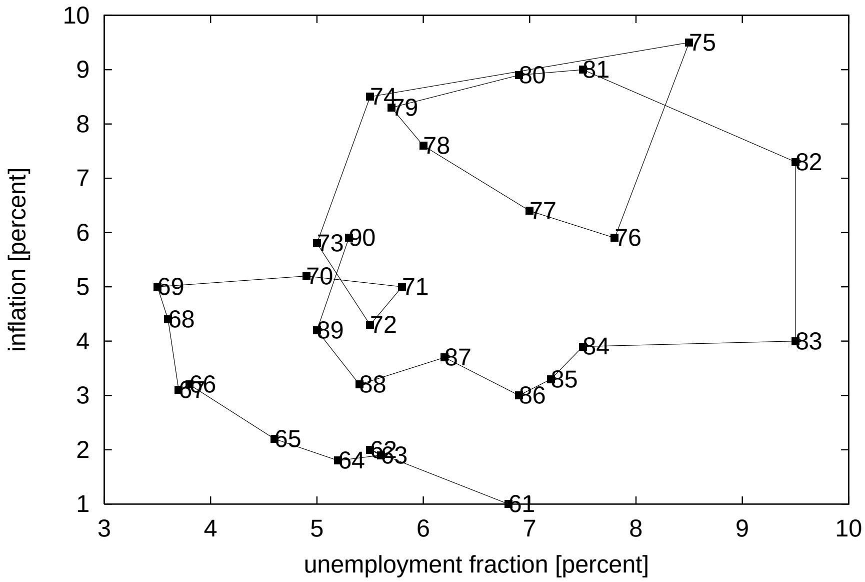

FIG. 1. Relation between inflation and unemployment between 1961 and 1990. Data copied from Ref. [2].

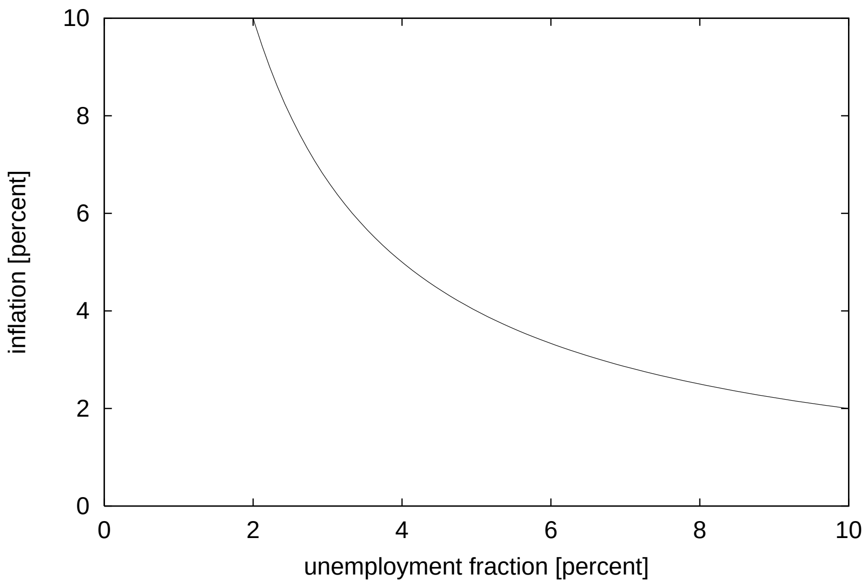

FIG. 2. Schematic Phillips curve. For the purposes of this paper, we set increases in productivity and thus the difference between price inflation and wage inflation to zero.

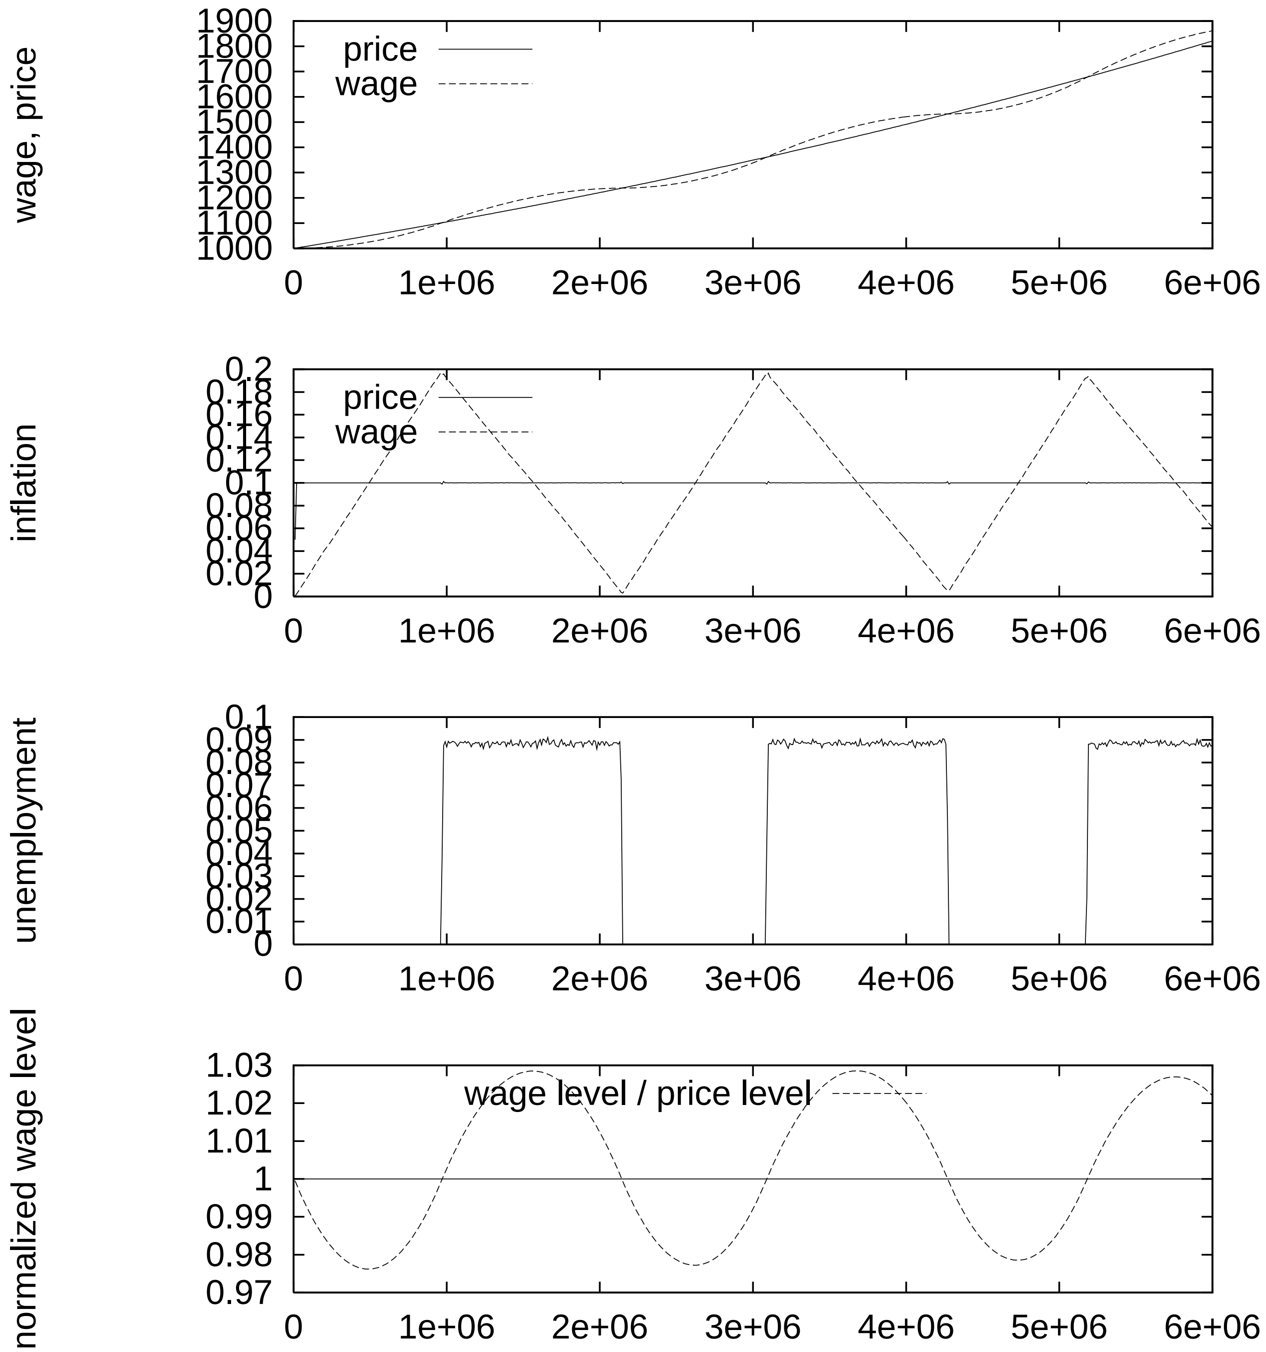

FIG. 3. First 6 000 000 iterations of a model as described in the text. (a) Wage level and price level as function of the iteration. (b) Inflation as function of the iteration. (c) Unemployment as function of the iteration. (d) Normalized wage level, i.e. wage level divided by price level.

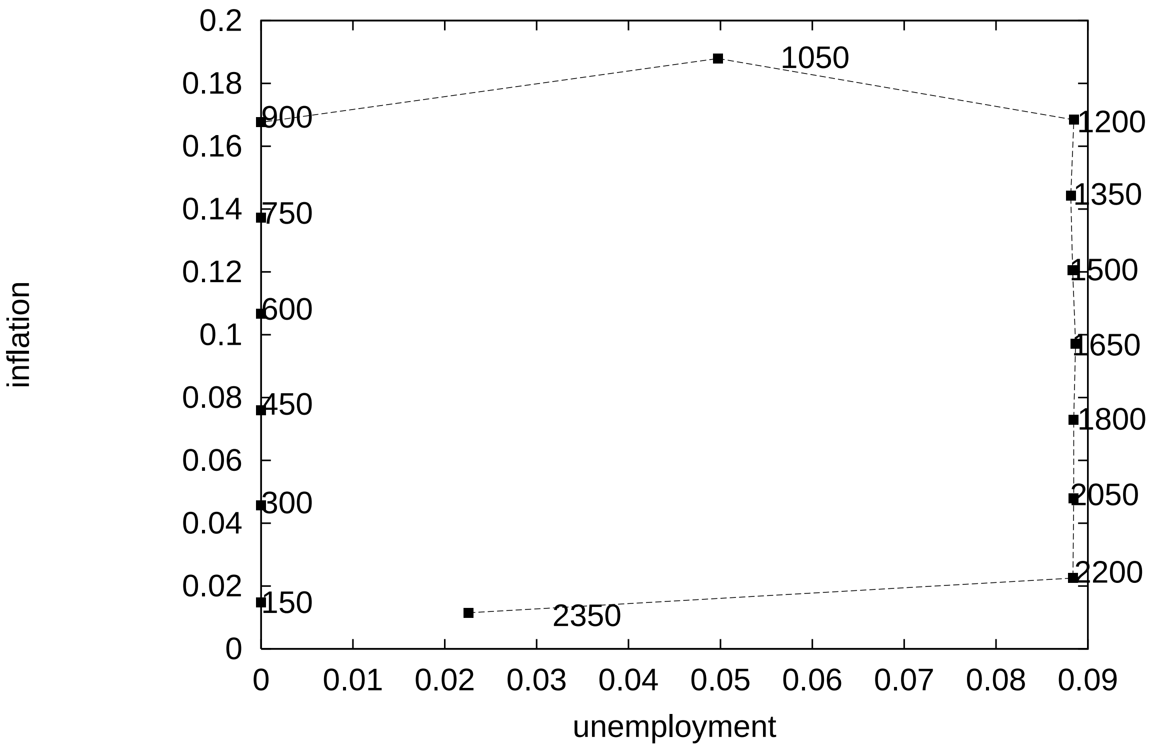

FIG. 4. Same simulation as in Fig. 3. Inflation vs. unemployment for the first 2 350 000 iterations. Numbers near data points denote iterations in thousand.

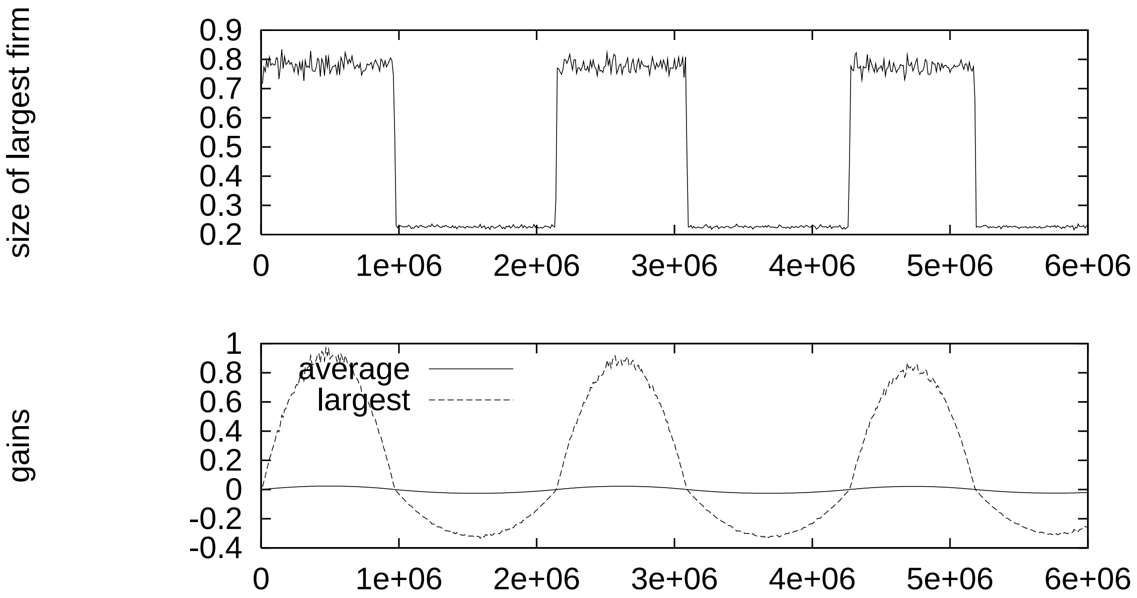

FIG. 5. Same run as in Fig. 3. (a) Share of the labor market of the largest company in the system. (b) Gains of the largest company according to (a), and average gains of all companies.

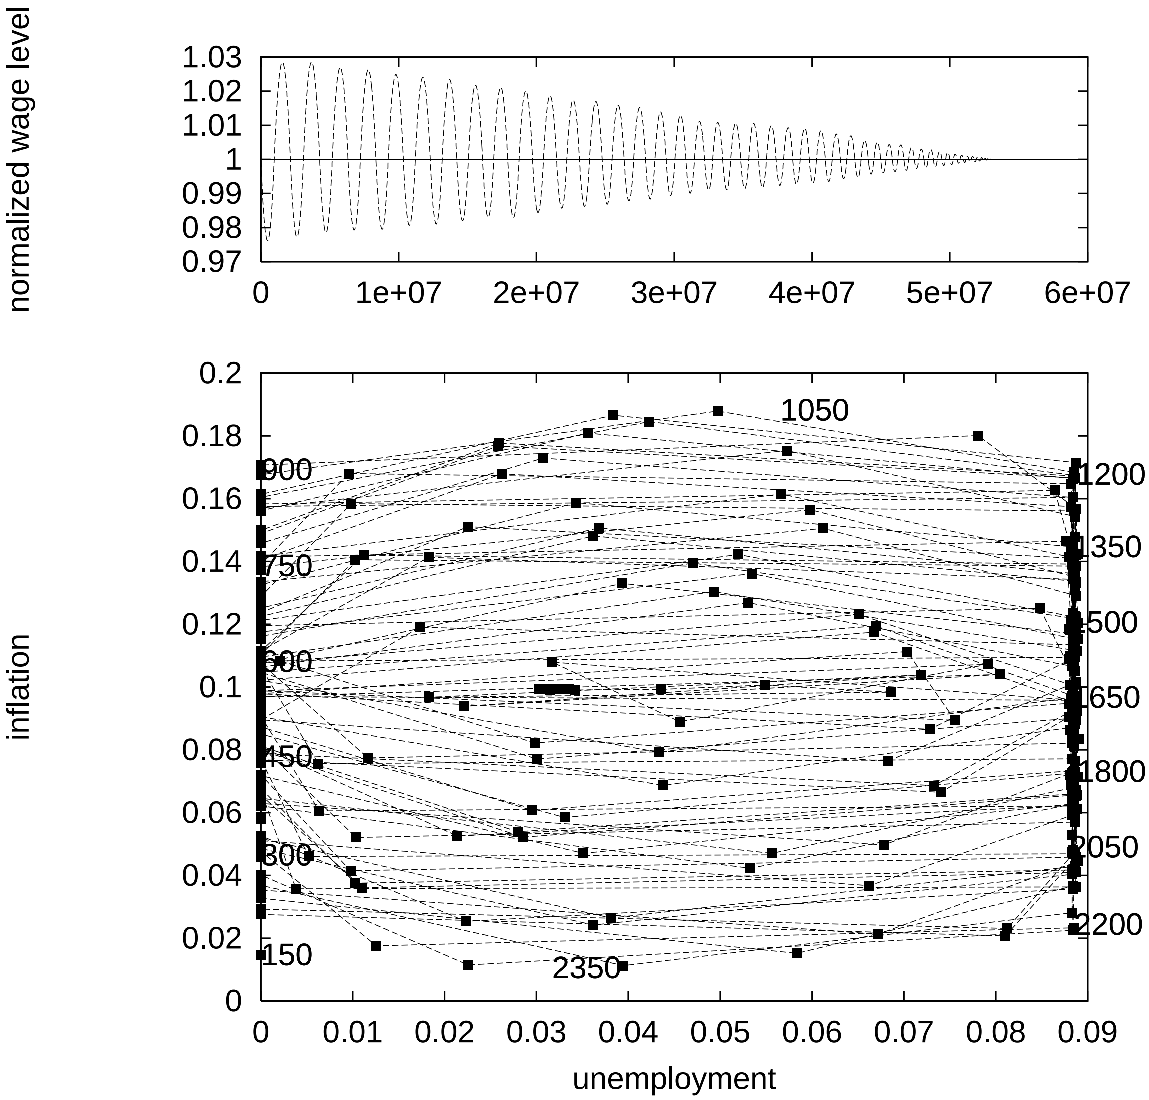

FIG. 6. Plotting a longer time period for the same simulation as in Figs. 3 and 5. (a) Normalized wage level vs. time; (b) inflation vs. unemployment.

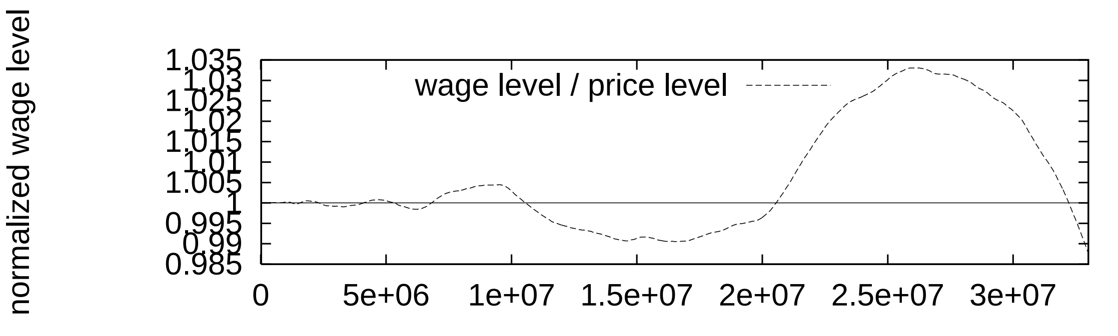

FIG. 7. Demonstrating instability. As described in the text, firms copy other firms by only copying the wage *inflation* parameter from other firms, not the wage itself. (a) Normalized wage over time. (b) Inflation vs. unemployment.