

Pallage, Stéphane; Scruggs, Lyle; Zimmermann, Christian

**Working Paper**

## Unemployment insurance generosity: a trans-atlantic comparison

IZA Discussion Papers, No. 3869

**Provided in Cooperation with:**

IZA – Institute of Labor Economics

*Suggested Citation:* Pallage, Stéphane; Scruggs, Lyle; Zimmermann, Christian (2008) : Unemployment insurance generosity: a trans-atlantic comparison, IZA Discussion Papers, No. 3869, Institute for the Study of Labor (IZA), Bonn,  
<https://nbn-resolving.de/urn:nbn:de:101:1-20081216232>

This Version is available at:

<https://hdl.handle.net/10419/35625>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

IZA DP No. 3869

## Unemployment Insurance Generosity: A Trans-Atlantic Comparison

Stéphane Pallage  
Lyle Scruggs  
Christian Zimmermann

December 2008

# Unemployment Insurance Generosity: A Trans-Atlantic Comparison

**Stéphane Pallage**

*Université du Québec à Montréal*

**Lyle Scruggs**

*University of Connecticut*

**Christian Zimmermann**

*University of Connecticut  
and IZA*

Discussion Paper No. 3869  
December 2008

IZA

P.O. Box 7240  
53072 Bonn  
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: [iza@iza.org](mailto:iza@iza.org)

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

## ABSTRACT

### **Unemployment Insurance Generosity: A Trans-Atlantic Comparison<sup>\*</sup>**

The goal of this paper is to establish if unemployment insurance policies are more generous in Europe than in the United States, and by how much. We take the examples of France and one particular American state, Ohio, and use the methodology of Pallage, Scruggs and Zimmermann (2008) to find a unique parameter value for each region that fully characterizes the generosity of the system. These two values can then be used in structural models that compare the regions, for example to explain the differences in unemployment rates.

JEL Classification: E24, J65

Keywords: unemployment insurance, labor market policy evaluation

Corresponding author:

Christian Zimmermann  
Department of Economics  
University of Connecticut  
341 Mansfield Road, Unit-1063  
Storrs, CT 06269-1063  
USA  
E-mail: [christian.zimmermann@uconn.edu](mailto:christian.zimmermann@uconn.edu)

---

<sup>\*</sup> We thank seminar participants at the OECD and the University of California Santa Barbara, and of course at the EDHEC conference for their comments on this work.

# 1 Introduction

It is often argued that Europe is more generous than the United States when it comes to protecting unemployed workers. Yet, it is difficult to establish the truth behind that conclusion and it is quite a challenge to quantify any difference in generosity.

Unemployment insurance [UI] policies have many different dimensions and thus are extremely difficult to compare through time and space. Income replacement ratios, the proportions of past income that constitute unemployment benefits, are the typical metric used in comparisons. Yet, replacement ratios are meaningless in terms of relative generosity if we do not account for differences in eligibility criteria, duration of benefits, etc. Comparing durations of benefits across countries is no more meaningful if we do not acknowledge differences in labor market conditions. For example, whether the reduction of the eligibility period for UI benefits matters depends on local unemployment duration. Thus while duration of benefits is much shorter in the United States than in many European countries, this does not necessarily mean that the US unemployment insurance is less generous, as unemployment duration is also much shorter, and the US program may be more generous in other dimensions that matter more for its labor market.

In this paper, we want to contribute to a better understanding of how generous, in an aggregate sense, unemployment insurance agencies are relative to one another. We use a methodology developed in Pallage, Scruggs and Zimmermann (2008) to summarize all policy dimensions into a single parameter and compare two very distinct UI agencies, that of France and that of an American state, Ohio.

The idea of measuring the generosity of UI systems is not new. It has generated an important literature. The OECD, for instance, has a research program that makes international comparisons of UI coverage for very specific types of workers. Martin (1996) summarizes these results. More recently, Scruggs (2006) compiles various measures of social programs for a specific type of household and looks at how they compare, one dimension at a time, through space and time. These works, however, ignore how the local labor market conditions may matter.

The approach we take here follows Pallage, Scruggs and Zimmermann (2008). In a simulation, we compare an economy having the complete characteristics of the actual UI program to an economy with a one-dimensional UI program. This single dimension is the level of UI benefits with no time limit. We measure the overall generosity of an unemployment insurance program as the level of benefits in the one-dimensional UI program that makes agents indifferent between that

and the actual programs. The base model we use is one of households facing repeated employment lotteries. They are liquidity constrained and they can try to self-insure against employment shocks if the UI program is not generous enough. This economy also exhibits moral hazard, which influences the optimal generosity, as seen in Hansen and İmrohorođlu (1992) and Pallage and Zimmermann (2001) in a similar set-up.

Our original exercise in Pallage, Scruggs and Zimmermann (2008) focused on how the unemployment insurance reforms in the United Kingdom had affected the system's generosity. In the current paper, instead of studying the temporal dimension, we make a spatial analysis by comparing a European country, France, to the US state of Ohio. This is of particular interest given the growing literature that tries to understand the gap between unemployment rates in continental Europe and North America. For instance, Ljungqvist and Sargent (1998) show that the unemployment insurance's income replacement ratio is critical in understanding differences in unemployment rates, but they offer little guidance on what true values should be used and what their impact would be. Den Haan, Haefke and Ramey (2001) find less of an impact of UI generosity, but again provide little in the way of measuring regional differences. The present paper offers a means to properly calibrate the difference in UI generosity in such cross-country analyses.

Our approach has several advantages. First it takes into account the changing labor market conditions. Second, it considers how agents' behavior may be modified by changing policies or the changing environment. Third, it allows us to consider what facet of the policy or environment accounts most for changes in UI generosity. Finally, as just mentioned, we can consider the effects of moral hazard on the design of an optimal UI policy.

In the following sections, we first detail the modeling approach, then discuss the parametrization of the households, the labor market and the UI policies. This calibration procedure is crucial, as we want to obtain quantitative answers. We then provide results and conclude.

## **2 Modeling Approach**

We use two models, the first features a complex unemployment insurance program, while the second has a simple one. For exposition purposes, we want to start by describing the common parts, i.e., the problem of typical households.

## 2.1 The problem of the households

Households care about consumption  $c$  and leisure  $l$ . They maximize an infinite stream of expected, discounted utilities. They can accumulate assets  $m$ , but are not allowed to borrow. Every period, they draw an employment opportunity or not. The likelihood of this event depends on whether they had an opportunity the period before. Job opportunities are drawn from Markovian lotteries. Agents may choose to turn down a job opportunity. An unemployment insurance system is in place, which allows households to obtain some benefits under some conditions. The UI agency balances its budget every period by collecting a tax.

Let us be more precise: The preferences of each household can be represented by the following function:

$$\max E_0 \sum_{t=1}^{\infty} \beta^t u(c_t, l_t)$$

where  $u(\cdot)$  is a utility function with the usual properties, i.e. increasing in each argument and concave;  $l_t = 1$  for someone who does not work,  $l_t = 1 - \hat{h}$  for someone who works, with  $\hat{h} < 1$ , a constant;  $\beta \in (0, 1)$  is the factor by which the household discounts time.

Asset holdings of the households evolve according to the following rule:

$$m_{t+1} = m_t + y_t^d - c_t, \quad m_t > 0 \quad \forall t$$

where  $y_t^d$  is the disposable income whose value depends on the status of the agent:

$$y_t^d = \begin{cases} (1 - \tau)y & \text{if employed } (w = e) \\ (1 - \tau)\theta y & \text{if eligible to UI } (w = i) \\ (1 - \tau)\psi y & \text{if unemployed and not eligible } (w = u) \end{cases}$$

where  $\tau$  is a tax rate used to raise the necessary revenue to finance the unemployment insurance program and  $w$  is an indicator of UI labor market status. Eligibility for unemployment insurance benefits may be dictated by various indicators here  $\theta$  and  $\psi$ , summarized by  $\alpha$  that will be specified for each model below. For the moment let us simply say that eligibility depends on a vector of variables  $s_t$  that evolves according to some, potentially endogenous, law of motion:

$$s_{t+1} = \chi(s_t)$$

Finally, households obtain every period a draw from a job opportunity lottery, following a binomial Markov process. The complete household problem can be

represented in recursive form, thus the Bellman equation of a worker with an employment offer is:

$$\begin{aligned}
 V(m, s|e; \alpha) &= \max \left\{ \begin{array}{l} \max_{m'} u(c, 1 - \hat{h}) + \beta \int_{s'|e} V(m', s'; \alpha) d(s'|e), \\ \max_{m'} \int_w u(c, 1) dw + \beta \int_{s'|u} V(m', s'; \alpha) d(s'|u) \end{array} \right\} \\
 S.T. \quad & m' = m + y^d(w, s; \alpha) - c \\
 & m' \geq 0 \\
 & s' = \chi(s)
 \end{aligned}$$

Note that when a worker turns down an offer, he may still get unemployment insurance benefits, depending on monitoring and his luck at defeating it. The probability to “beat” the system, i.e. collect benefits after turning down an offer, is common knowledge and measured by  $\pi$ .

For a worker without an employment offer, the Bellman equation can be written as follows:

$$\begin{aligned}
 V(m, s|u; \alpha) &= \max_{m'} u(c, 1) + \beta \int_{s'|u} V(m', s'; \alpha) d(s'|u) \\
 S.T. \quad & m' = m + y^d(i, s; \alpha) - c \\
 & m' \geq 0 \\
 & s' = \chi(s)
 \end{aligned}$$

### Equilibrium

For each period in the sample, a steady-state equilibrium is an allocation of work, asset and consumption for all households, a value function  $v(\cdot)$ , a distribution of assets  $f(\cdot)$ , and a tax rate  $\tau$  such that:

1. households solve their individual intertemporal problems, given  $(\alpha, \tau, \pi)$  and labor market characteristics;
2. the unemployment insurance agency balances its budget;
3. there is an invariant distribution of agents.

It is important to note that under this definition, all households assume that the current unemployment insurance and labor market characteristics in any given period will remain unchanged forever. One may argue whether this myopic view



of the world is a valid approximation or not. One could think of taking into account how agents may be forward looking in these dimensions as well, that is, how they may anticipate changes in the parameters of the UI system as well as in unemployment rate and duration. Modelling these expectations, however, is very difficult, in particular expectations about changes in the UI system. Doing so would make the computation of the equilibrium several times more difficult, as one would not be able to rely on invariant distributions anymore.

## 2.2 The simplified UI program

We need to make specific what makes an unemployed worker eligible for unemployment insurance, that is we need to specify what lies under what we have so far referred to as  $\alpha$ . For the simplified UI program, we assume that unemployment benefits can be obtained immediately and that unemployed workers stay eligible forever and obtain every period the same proportion  $\theta$  of their income. The only times when a worker does not receive benefits is when she has become ineligible, either by getting caught cheating the system or simply by not asking for benefits. In such case, a household gets a share  $\psi$  of past income.

Finally, monitoring is characterized by a probability of success in shirking of  $\pi$  that is positive only when the worker has been previously unemployed. In other words, a quitter cannot shirk successfully, but a searcher can with probability  $\pi$ . The simplified UI program thus has the following vector of parameters:

$$\alpha = (\theta, \psi, \pi).$$

This is the set of parameters to which we want to map the detailed UI program.

## 2.3 The complex UI program

Now we want to describe a real world UI program as completely as computationally feasible. It has the following components:

1. A waiting period  $a$ , i.e., unemployed workers have to wait some time before becoming eligible for full benefits. Partial benefits may be given, though.
2. An eligibility period  $z$ , i.e., how many periods an unemployed worker can obtain benefits if she remains jobless.

3. The proportion of income that unemployed workers obtain as benefits,  $\theta(j)$ , which may vary through the unemployment spell, including the waiting period ( $j = 1, \dots, z$ ).
4. The proportion of income unemployed workers receive after losing eligibility,  $\psi$ .
5. The probability of shirking success for searchers,  $\pi$ .

Thus, the set of policy parameters we want to calibrate from the data is:

$$\alpha = (a, z, \{\theta(j)\}_{j=1, \dots, z}, \psi, \pi).$$

We can now turn to finding those policy parameters for the economies of interest.

### 3 Calibration

The calculations for income replacement rates in the United States are based on the unemployment insurance benefit system in force in 2005 for the state of Ohio. Unlike most European countries, each state in the United States has its own specific rules for unemployment insurance benefits. States differ considerably with respect to the replacement rate and the maximum benefit amount. Ohio was chosen as the reference state, because it is similar to France in terms of average wages and its manufacturing share. Also, Ohio is often considered the most “American” state in its demographic structure, including in its racial composition. New York and California, in the opposite, are too rich relative to the mean and not as manufacturing-based.

The unemployment insurance benefit in Ohio (US\$ 17,212) is the annualized benefit that would be paid to someone earning the average insured wage (US\$ 36,500) (United States Department of Labor, 2005). This is the maximum weekly benefit payable. (In other words, those earning more than the average wage would receive an identical benefit amount.) The benefit is payable weekly for up to six months. The net benefit (US\$ 16,111) and the net wage (US\$ 29,034) were computed by subtracting relevant federal and state taxes, assuming standard deductions and allowances for a single person. Federal and state social assistance programs (such as Temporary Assistance for Needy Families) provide no cash or housing benefits for able-bodied singles in Ohio. However, the Federal Government’s Food Stamp program, provides a cash equivalent benefit for food. This

benefit was counted as social assistance benefit in our calculations, using its maximum benefit of US\$ 1,944 per year.

In the case of France, the unemployment benefit (€16,529) is based on the annualized benefit in the national unemployment insurance system. This amount is 54.7% of the insured wage. The insured wage (€30,219) corresponds to the average production worker wage in the OECD's 2005 Benefits and Wages Statistics. The relevant taxes and social charges, using standard deductions and allowances, were subtracted from the gross benefit and gross wage to produce the net benefit (€13,645) and net wage (€21,470). The net social assistance benefit amount (€8,300) combines the RMI and (zone II) housing benefits.

For labor market indicators, we use the 2005 average unemployment rate of 5.9% in Ohio, as well as an unemployment duration of 21 weeks. In 2005, the average duration of UI benefit spells was 15.5 weeks, with 29% exhausting their benefits after 26 weeks. Assuming a Poisson process, we find an average duration of 21 weeks. In France, the unemployment rate was 9.8% in 2005. For unemployment duration, we average the numbers found for 1994–2002 by Brunet, Clark and Lesueur (2006), which implies 36 weeks.

Finally, the remaining parameters are set identically in both countries. First, we need to take a stand on the utility function. Following the literature that has made use of this type of models, we select the following CES utility function:

$$u(c, l) = \frac{(c^{1-\sigma} l^\sigma)^{1-\rho}}{1-\rho}.$$

Again, taking conventional values from the literature, we set  $\sigma = 0.33$ ,  $\rho = 2.5$  and  $l = 1$  for an unemployed agent and  $l = 0.55$  for a worker, which is consistent with the fact that workers on average tend to spend 45% of their available time at work. Finally, we set  $\beta = 0.999165$  which, with a model frequency of a week, corresponds to a yearly discount rate of just above 4%. We carry robustness exercises to check how these choices influence our results.

## 4 Results

We solve for the generosity of the UI program in the following way. First, we write the Bellman equation for the detailed economy, where we discretize the asset space in a grid. Using typical dynamic programming techniques, we iterate on the Bellman equation and the distribution of agents in the state space. Once convergence is achieved, we verify whether the UI agency's budget is balanced and

adjust the tax rate if needed. Once convergence is also obtained on the balanced budget requirement, we use the invariant distribution of households and the value function to compute an average welfare value.

Next we turn to the simplified economy, proceeding in the same way for an initial guess of the generosity parameter  $\theta$ . The average value from this model is then compared to the previous one, and  $\theta$  is adjusted. We repeat the exercise until the average welfare values are equalized across economies.

Alternatively, one could also proceed in a much simpler manner, by computing what we call a naive measure of generosity. Under such a scheme, we determine the perpetual benefit  $\theta$  that would provide the same life-time expected benefit from a single unemployment spell:

$$\sum_{t=1}^a \beta^{t-1} \theta_t + \sum_{t=a+1}^{a+z} \beta^{t-1} \theta_t + \sum_{t=a+z+1}^{\infty} \beta^{t-1} \psi = \frac{1}{1-\beta} \theta.$$

With such a measure, of course, we would not be taking into account labor market conditions, the endogenous self-insurance decisions of households through asset accumulation, or the impact of the presence of the UI system on those that are not unemployed. We compute this measure below, to show the importance of labor market conditions and endogenous choices in the assessment of the relative generosity of a UI program.

Table 1 provides our results, including robustness exercises. We see first that while program benefits in Ohio and France look rather close (56% and 64%), the naive measure finds a huge difference in generosity (8% and 41%). This, and the fact that replacement ratios are much lower than the program benefits, stem from the waiting period and the benefits after UI, which are much lower in Ohio. In fact, one could conclude from this that the UI system may not even be worth its cost.

The benchmark numbers (third line of Table 1) provide the results from the full model runs without moral hazard ( $\pi = 0$ ). We find that the replacement ratios  $\theta$  are higher than the naive measure, but still lower than program benefits. 0.15 in Ohio is very low, while 0.50 is actually in the ballpark of what some studies have used for France or Europe. These are the numbers we recommend to use when calibrating a model with a simple unemployment insurance system. They properly account for eligibility rules and labor market conditions.

Obviously, we made some choices when setting up or calibrating the model. The following experiments try to shed some light on how these choices may have influenced the results. In the first one, we remove leisure from the utility function,

Table 1: Unemployment insurance generosity results

	Ohio	France
Actual program	0.56	0.64
Naive measure	0.08	0.41
Benchmark	0.15	0.50
No leisure	0.15	0.47
$\rho = 10$	0.17	0.57
$\rho = 1.1$	0.33	0.39
$\beta = 0.99833$	0.15	0.50
$r = 0.0004$	0.31	0.31
$\sigma = 0.67$	0.15	0.47
Other's job market	0.27	0.58
No income security	0.08	0.15

Note: The table provides the replacement ratios that characterize the unemployment insurance programs of France and the US state of Ohio under various experiments.

as in fact many other models do. It turns out that the generosity measures are not affected much by this change. The reason for this is that leisure is removed simultaneously in both models, the complex and the simplified one. This is also the reason why other calibration choices lead to small changes in the numbers, like increasing risk aversion ( $\rho = 10$ ), impatience ( $\beta = 0.99833$ ), halving the share of consumption in the utility function ( $\sigma = 0.67$ ).

One change is significant though. Reducing the risk aversion parameter to 1.1 makes both regions similar in generosity. The basic intuition is that when agents do not care much about fluctuations in consumption and leisure, fluctuations in income matter little as well and whether the labor market conditions are different or the system is designed in various ways has little impact.

Other robustness exercises do not show significant departures from the benchmark numbers. These include introducing moral hazard in the model where workers can turn down job offers. Again, as generosity is by definition the same in both the complex and simplified models, household tend to shirk in similar numbers in both cases and model outcomes in terms of  $\theta$  are not significantly altered.

Of particular interest is that our procedure also allows us to conduct counterfactual experiments. For example, if we transpose Ohio's labor market conditions in France but keep France's UI system, we find that the measure of generosity reaches 0.58, which is higher than anything else we have found so far. What we

de facto is combining a UI system that is already close to optimal — thus inducing little precautionary asset accumulation — with a labor market characterized by lower risks (and again less incentives for savings). This leads to a situation where people save very little, which brings it even closer to the optimal  $\theta$ , which is typically in the 0.5-0.7 range. More intuitively, adopting the French UI system in Ohio raises its generosity, but not that much as unemployment duration is shorter.

In another experiment, we remove all income support after the UI eligibility period. We notice that the generosity in France drops severely. The system was generous because of long unemployment durations combined with generous income support. This highlights the fact that it is not necessarily the unemployment insurance that is generous in France, but rather what comes after it. Interestingly, generosity is now higher in Ohio under this experiment. This can be explained by the fact that the benefits are only a little lower in Ohio when eligible, but much fewer people run out of eligibility in Ohio as unemployment duration is shorter.

## 5 Conclusion

We have used a microfounded model in which workers have the opportunity to partially self-insure against unemployment risk to measure and compare the generosity of a European and an American unemployment insurance systems. Applying our methodology to France and Ohio, we find the first to be indeed much more generous. We quantify this unemployment insurance generosity by computing the replacement ratio in perpetuity that makes agents indifferent with actual programs in all their complexities. Using this metric, we can say that France is about three times more generous than the state of Ohio (equivalent replacement ratios of 50% in France versus 15% in Ohio).

Our measures are robust to most of the assumptions we have taken when building our model. Additionally, our methodology allows us to understand why the generosity differs across the two regions. We find that the large discrepancy is due mostly to a combination of higher benefits after unemployment insurance eligibility and a longer unemployment duration in France.

## 6 References

- Brunet, Carole, Andrew Clark and Jean-Yves Lesueur, 2006, Statut résidentiel et durée de chômage: une comparaison microéconométrique entre la Grande-Bretagne et la France, document de travail GATE 06-13, Université Lumière Lyon 2.
- Den Haan, Wouter, Haefke, Christian and Ramey, Garey, 2001. "Shocks and Institutions in a Job Matching Model," NBER Working Papers 8463, National Bureau of Economic Research, Inc.
- Department of Labor. 2005 Comparison of State Unemployment Insurance Laws: 2005. <http://www.workforcesecurity.doleta.gov/unemploy/uilawcompar/2005/comparison2005.asp> (accessed November 15, 2007)
- Hansen, Gary D. and Imrohoroğlu, Ayse, 1992. "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard," *Journal of Political Economy*, University of Chicago Press, vol. 100(1), pages 118–42, February.
- Ljungqvist, Lars and Sargent, Thomas J., 1998. "The European Unemployment Dilemma," *Journal of Political Economy*, University of Chicago Press, vol. 106(3), pages 514–550, June.
- Martin, John P., 1996. "Measures of Replacement Rates for the Purpose of International Comparisons: A Note," *OECD Economic Studies*, OECD, number 26, pages 99–115.
- OECD, 2007. Benefits and Wages Statistics: Country Specific Files. [http://www.oecd.org/document/29/0,3343,en\\_2649\\_34637\\_39618653\\_1\\_1\\_1\\_1,00.html](http://www.oecd.org/document/29/0,3343,en_2649_34637_39618653_1_1_1_1,00.html) and <http://www.oecd.org/dataoecd/37/36/39676814.pdf>. (accessed November 15, 2007)
- Pallage, Stéphane, Scruggs, Lyle and Zimmermann, Christian, 2008. "Measuring Unemployment Insurance Generosity", working paper 2008-42, University of Connecticut Department of Economics.
- Pallage, Stéphane and Zimmermann, Christian, 2001. "Voting on Unemployment Insurance," *International Economic Review*, vol. 42(4), pages 903–23, November.
- Scruggs, Lyle, 2006. "The Generosity of Social Insurance, 1971–2002," *Oxford Review of Economic Policy*, Oxford University Press, vol. 22(3), pages 349–364, Autumn.