

Where Should We Live?

Andrew E. Clark*

Paris School of Economics and IZA

April 2009

Abstract

One of the central topics in recent empirical work on subjective well-being is that of comparisons to a reference group, over a variety of domains of economic and social life. One such reference group is neighbours. Any resulting spatial spillovers that are identified have potential implications for the welfare-maximising degree of spatial segregation. In this paper, we summarise some recent findings with respect to geographical comparisons of income, unemployment, health and religion, and present some new results regarding spatial spillovers in marriage. The resulting predictions regarding spatial segregation differ sharply according to the domain under consideration. While work in this area remains very preliminary, subjective well-being data may well help to both identify spillovers from neighbours, and inform about individuals' location decisions.

Keywords: Comparisons, Spatial Spillovers, Income, Unemployment, Marriage, Health, Religion.

JEL Codes: D01, D31, D62, I10, I31, J12, J64, N3, R23.

* PSE, 48 Boulevard Jourdan, 75014 Paris, France. Tel.: +33-1-43-13-63-29. E-mail: Andrew.Clark@ens.fr. I am grateful to audiences at numerous seminars for developing these ideas, and to an anonymous referee for helpful suggestions.

Where Should We Live?

1. Introduction

Where should people live? The typical answer is wherever suits them, which is perhaps not very interesting economically. The situation is different when there are spillovers or externalities between individuals. Then my location decision may impinge on your utility, and therefore on your choice of where to live.

These spillovers might be purely size-related, reflecting increasing returns to scale (with respect to public transport, for example), overcrowding, or some kind of thick-market externality; alternatively, they could be dependent on the individual characteristics of those who live in a certain area, such as age, income, education, marital status or religion. It is this the latter kind of externality that is considered here.

The Economics of Happiness has addressed many issues, but has not paid overmuch attention to the question of location. One aspect of the latter is the value of amenities, which has been analysed by van Praag and Baarsma (2001) in the context of aircraft noise, Clinch *et al.* (2007) for green spaces and other amenities, and Dolan and Metcalfe (2007) for urban regeneration. A recent Inter-American Development Bank (IADB) project has used subjective well-being data to evaluate the impact of various housing and neighbourhood characteristics across cities in Latin America (see Powell and Sanguinetti, 2008). This work reveals the price at which individuals are prepared to trade different amenities off against each other. If all individuals are identical, then, absent any moving or other costs, all locations should procure the same utility. If individuals are heterogeneous, then those who value green spaces will locate in greener areas, and so on.

Here we consider interactions between individuals, and their implications for location choice. Again, we will use subjective well-being data to essentially value locations, but now in terms of the kinds of people who live there, rather than the innate nature of the housing or the neighbourhood. The varying effects of other occupants' characteristics on individual well-being has direct implications for location choice.

The next section of the paper presents some evidence of spillovers between neighbours, and their implications on location choice, with regard to five key individual characteristics: income, unemployment, health, religion and marriage.

2. Spillovers from Neighbours

Individual well-being in a particular location depends not only on the area's characteristics in terms of natural resources or infrastructure, but also on the type of people who live there. While most people would not require empirical analysis to be convinced that individuals don't like living close to a landfill or under a flight path, the situation with respect to neighbours requires perhaps more statistical analysis. This section summarises some recent work on well-being and neighbours' characteristics, and draws out the conclusions with respect to location choices.

Income

A body of research across the social sciences has related various measures of well-being to not only individual income, but also the income of the individual's reference group. A common finding in this literature (surveyed in Clark *et al.*, 2008a) is that individual well-being is negatively correlated with reference group income. This is interpreted as support for a utility function which depends on relative income, so that part of well-being is determined

by the level of my own income relative to that of others; or, in other words, my utility falls as the income of my reference group rises.

One of the definitions of reference group is geographic: Do you want your neighbours to be rich or poor? Analyses using reference group income at a fairly aggregated level, such as Luttmer (2005) and Ferrer-i-Carbonell (2005), have concluded that life satisfaction is positively correlated with own income, but negatively correlated with neighbours' incomes. The obvious implication for locational choice is that, for a given level of own income and neighbourhood amenities, individuals will be better off in areas with poorer neighbours, as they will then feel relatively rich.

This does yield the seemingly counterfactual conclusion that individuals should flock to low-income neighbourhoods: houses there should actually sell for a premium. This is analogous to the analysis of wages within a firm in Frank (1985), where status-conscious individuals will accept lower wages to be at the top of the pecking order within a firm: people are ready to pay for status. Here individuals would be ready to pay for status via presumably higher house prices in low-income neighbourhoods.

We probably do not observe this phenomenon: the kind of flocking that we do seem to observe is of relatively wealthy people together. Why does location choice not follow the implications of well-being analysis? One way of reconciling implication and observation is to note that it is remarkably difficult to make the "*ceteris paribus*" assumption hold. Neighbours' income may well be correlated with other aspects of behaviour that individuals appreciate (like being able to pay for local public goods, or contributing to local social capital). Recent work by Clark *et al.* (2009) using Danish administrative data does indeed find, using reference groups based on the very local level (often down to blocks of a few hundred metres), that individual satisfaction is positively associated with very close neighbours' incomes.

Why do Clark *et al.* (2009) find results diametrically opposed to those in Luttmer (2005) and Ferrer-i-Carbonell (2005)? One immediate difference between the analyses is the size of the “local area”. In Clark *et al.*, there are a few hundred households in the reference group; in Luttmer there are on average 150 000 individuals per reference group, and in Ferrer-i-Carbonell, the reference group is defined by German Länder. Having such large geographical areas may well abstract from the local public goods and good neighbourliness aspects. Having a poor region 10 or 20 kilometres away may then make me feel relatively rich, but without imposing any negative externalities via local public goods or something else. As such, estimating well-being equations using wide geographical reference groups allows us to identify the pure income externality between individuals. But for location choice, all externalities are taken into account, including the positive ones from close neighbours. It seems as though the net effect is positive. People like having rich close neighbours, and are ready to pay a premium to live in rich areas. This is indeed what we observe.

Unemployment

Although income comparisons have attracted a great deal of attention, it is likely that comparisons take place over many different aspects of economic and social life (see Clark, 2008): all of these will potentially impact on location choices. We now turn our attention from income to labour market status, and consider the implications of unemployment, both own and others.

Probably one of the least controversial correlations to come out of the well-being literature is the strong negative one between own unemployment and own well-being. But what of the effect of others’ unemployment on my own well-being? Others’ bad experiences on the labour market might be thought to inform us of our own future prospects. As such, anyone who is active on the labour market, whether employed or unemployed, is likely to

receive news of greater unemployment with some despondency. However, in line with the idea of income comparisons sketched above, there may well be a status or comparison component as well. As such, having more unemployed people around you may make you feel better off, relatively speaking.

The sum of the signal effect (about future prospects) and status effects (comparing to others) is *a priori* ambiguous. Research on the first seven waves of data from the British Household Panel Survey (BHPS: see <http://www.iser.essex.ac.uk/survey/bhps>) in Clark (2003) suggested that the sign of the effect of others' unemployment on the individual's own well-being depended on the individual's own labour-force status. Greater regional unemployment was shown to be negatively correlated with the well-being of the employed, but positively correlated with the well-being of the unemployed.

The graph in Figure 1 illustrates the findings. The measure of well-being here is the 12-item version of the General Health Questionnaire (the GHQ-12). This consists of twelve individual questions, covering, for example, trouble sleeping, belief in self-worth, enjoying day-to-day activities, and playing a useful role. All of these questions are answered on a one-to-four scale. Two of the answers are negative and two are positive. With respect to trouble sleeping, the positive answers are "*Not at all*" and "*No more than usual*", and the negative answers are "*Rather more than usual*" and "*Much more than usual*". The "Caseness" measure of GHQ adds up the number of questions that are answered using one of the negative responses. We here invert this score, producing an index on a scale of 0 to 12, where higher numbers (i.e. more questions answered positively) correspond to better psychological functioning.

Figure 1 shows the bivariate relationship between this Caseness score and regional unemployment. The Y-axis shows the average difference in GHQ scores between the employed and the unemployed, by region (11) and by year (7): this is a measure of the

psychological cost of unemployment. The Labour Force Survey unemployment rate, again by region and by year, is represented on the X-axis. The resulting scatter of points is consistent with the psychological effect of unemployment being lower in regions where the unemployment rate is higher. Clark *et al.* (2008b) find somewhat similar results in German Socio-Economic Panel (GSOEP) data, although they emphasise the key role of labour-market attachment in mediating the effect of others' unemployment on individual well-being.

Taking these results at face value, what spatial distribution of employed and unemployed would maximise welfare? In the previous section on income, high earners made everyone miserable, because they drove relative income down. Equally, low earners in the reference group were welcome, as they made everyone feel better off. With the BHPS unemployment results described above, we lose this homogeneity in the spillover effect: others' unemployment makes the employed feel worse off, but makes the unemployed feel better off. The implication is that geographical sorting by labour-force status will maximise welfare: the unemployed should live with the unemployed, and the employed with the employed. The unemployed will almost always report lower levels of well-being than do the employed, but this kind of labour-force segregation will reduce the gap between the two groups, as Figure 1 suggests.

Health

We next consider one of the most essential elements of individual well-being: health. Some recent work has highlighted possible geographical spillovers in this realm too, although much remains to be done in order to be sure of the causal mechanisms. Following on from the discussion of unemployment above, we ask whether the well-being impact of poor health is attenuated by poorer health in the reference group. Several recent papers suggest that this might be the case.

The first of these considers Body Mass Index (BMI) and a geographical reference group. Blanchflower *et al.* (2009) use Eurobarometer data to show that self-reported feelings of being overweight increase with own BMI, but fall with the average BMI of the peer group (measured at the country*age group*sex level). Clark and Etilé (2008) is along somewhat the same lines, but uses a much-reduced “geographical” definition of the reference group: the individual's own household, and specifically their partner. The relationship between individual well-being on the one hand, and both own and partner's BMI on the other hand is analysed in GSOEP and BHPS data. While own BMI and well-being are negatively related, the strength of this relationship is mitigated by partner's BMI. In the BHPS data, an individual in a couple where both are obese reports a similar mental well-being level (as measured by the Caseness GHQ index described above) as an individual in a couple where neither is obese. The “worst” situation is being obese and having a non-obese partner.

Powdthavee (2009) also considers the household dimension of health spillovers, and shows that the strength of the relationship between own health problems and own subjective health is lower when the individual lives with other people who also have health problems. This is consistent with the individual comparing her health problems to those of other household members (although it should be noted that this is not the only interpretation, and that these spillover effects are fairly small in size).

Again, taking the results at face value, welfare would be maximised by geographical segregation in health: while those in better health don't want to live with the unhealthy, those who are in worse health do. It should be emphasised that research in this area is still very preliminary, however, and much more work is needed before we can definitively conclude as to the presence of geographical comparison effects in health, and thus any kind of health segregation.

Religion

In Clark and Lelkes (2009), we appeal to data from the first three waves of the European Social Survey (ESS) to analyse the existence of religious spillovers. While there is now a very substantial literature showing that the religious report higher levels of subjective well-being, only little work has considered whether my neighbour's religion affects my own well-being, and, critically, whether the size (and sign) of any such effect depends on how religious I am myself.

In the ESS data, own religious behaviour is positively correlated with individual life satisfaction, controlling for "standard" demographic characteristics and country and year fixed effects. Religious behaviour here is alternately defined by prayer or churchgoing. We introduce religious spillovers by calculating average religious behaviour at the regional level (which latter is defined at the NUTS2 level). When added as explanatory variables to a life satisfaction equation, these also attract positive coefficients: individuals are more satisfied on average, *ceteris paribus*, in more religious regions.

This spillover does not work in the same way as those for health and unemployment described above, and is more analogous to the generic spillover of others' income: the positive effect of others' religion is found both for those who are religious themselves and for those who are not religious. To illustrate the results, Catholics report lower levels of life satisfaction as the percentage of "atheists" (those who do not identify themselves as belonging to any religion) in the region rises. The same is true for Protestants and for those of "Other" religions. However atheists themselves are also less satisfied as the percentage of atheists in the region rises.

The spatial implication of this generic positive spillover from others' religion is that, all other things equal, people will prefer to live in religious rather than non-religious regions. The fact that some people most definitely do live in relatively less-religious regions must

therefore reflect some kind of compensating differential for their neighbours' godlessness (like lower house prices) or heterogeneity in preferences (of which more below).

Marriage

The employed are more satisfied living with other employed rather than other unemployed, and those in ill health are more satisfied living with others in ill health rather than others in good health: we can imagine carrying out analogous analyses with respect to any number of potential different behaviours or characteristics. This last sub-section will consider a characteristic that has not yet to my knowledge been addressed in the Economic well-being literature: spillovers between marital statuses.

To do so I use the first sixteen waves of the BHPS and run a series of well-being regressions. The results of these are summarised in Table 1. This table is divided up into two panels with three regressions each. The first panel refers to the results using the 0-12 GHQ score (as described in the unemployment section above) as the dependent variable; the right-hand uses the overall life satisfaction variable, which is measured on a one to seven scale, with higher numbers indicating greater satisfaction. The number of observations is lower in the right-hand panel, reflecting the fact that the life satisfaction question was only asked in waves 6-10 and 12-16.

The first regression in each panel shows the results of "standard" well-being regressions. While there are a variety of control variables (as listed at the foot of the table), only the estimated coefficients relating to marital status are produced here. In the standard regression, the married report higher levels of well-being than do singles (the omitted category), although only significantly so for life satisfaction. The separated, divorced and widowed all report lower well-being scores by both measures.

The second column in each panel adds information on the regional marital composition of the respondent's region. These are calculated by using weighted data directly from the BHPS itself. This is not totally satisfactory, and future work will match in marital status information from an external data source, such as the Census or the Labour Force Survey. I calculate the regional percentage of married, separated, divorced and widowed. All four of these percentage variables are added to column 1's specification. All percentage variables have been divided by 100 in this table, and standard errors are clustered at the region*year level. None of the estimated coefficients on these regional variables are significant, although that on the regional percentage married is positive in both panels.

The third column interacts own marital status with the regional percentage married. Remarkably, this produces a negative and significant estimate in both panels, with the main effects of both own marriage and the regional marriage rate now being positive and significant. The results in Table 1 thus suggest that the well-being of the non-married is highest in high-marriage regions; on the contrary the well-being of the married is highest in low-marriage regions (although it should be pointed out that the hypothesis that the sum of the coefficients on Percentage(Married) and Married*Percentage(Married) is zero is not rejected by the data).

Taken literally then, the married impose (weakly) negative externalities on each other, but positive externalities on the non-married. In terms of welfare, an obviously dominated choice would seem to be to have all of the married located separately from all of the non-married, as they would then only produce negative externalities for each other, but no positive externalities for the non-married. In the more realistic eventuality of a certain distribution of the non-married across regions, it turns out to be remarkably difficult to predict the "best" distribution of the married.

Some simulations carried out for the two-region case produced divergent predictions, dependent on where the non-married were located, and on the percentage of the population that was married. As a thought experiment, imagine this two-region world, and an initial situation where the non-married are equally spread out between the two regions, but all of the married live in Region 1. What happens to well-being if we move one married person from Region 1 to Region 2?

- I. The unmarried in Region 2 are better-off, as the marriage rate in Region 2 rises
- II. The married in Region 1 are better-off, as the marriage rate in Region 1 falls.
- III. The non-married in Region 1 are worse-off, as the marriage rate in Region 1 falls
- IV. The (married) individual who moves is better-off, as the marriage rate in Region 2 is lower than that in Region 1, and the effect of the marriage rate on the married is negative.

The sum of these four effects can be positive or negative, depending on the number of people concerned and the point from which the “redistribution” of the married starts. The optimal distribution of certain groups of individuals who impose negative externalities on each other, but positive externalities on the outside group deserves further work. In particular, it would be useful to see whether these externalities produce a coordination problem, such that the free choice of location by individuals may not produce the social optimum.

Another useful line of research would be to think a little more about the mechanism behind this marriage spillover. Why should more married people bring benefits to the non-married, but harm (or at least leave indifferent) the married? The discussion in this paper has been mostly couched in terms of social comparisons or status. However, more prosaic transmission channels are also likely. In particular, having more married people in the neighbourhood may result in there being more children. This will impose negative

externalities on those who dislike children; and will specifically be harmful for individuals with children if it increases competition for places at “good” schools, which are fixed in the short-run.

To test this idea, I re-ran the regressions in Table 1 separately for those with and without school-age children. While there were some differences in the estimated coefficients across these regressions, there was no strong evidence that the negative externality from the percentage married applied more strongly to those with school-age children. The identification of the channels through which these externalities work is another area where more research would be helpful.

3. Conclusion

This paper has appealed to recent empirical work on spillovers between individuals to produce a number of predictions about welfare-maximising spatial distributions. These spatial predictions differ remarkably according to the characteristic under consideration. Some of the spillovers that we have identified seem to be generic. Having religious neighbours makes everyone more satisfied, regardless of their own religiosity. More religious neighbourhoods should then be sought-after. The situation with respect to income is a little more complicated, although analysis using large geographical reference groups suggest that individuals like having high income relative to their neighbours, some recent work has suggested that it is in practice difficult to dissociate income from the local tax base or local public goods or social capital. As such, people may in general prefer having rich close neighbours (although they would like non-local members of their reference group to have lower income). The key distinction here is in terms of who counts for income comparisons and who counts for local public goods and social capital. These two effects pull in different directions, and the net effect may well be that rich neighbourhoods are preferred.

Other spillovers depend on the individual's own characteristics. In British data, others' unemployment was shown to reduce the well-being of the employed but increase the well-being of the unemployed. As such, any spatial segregation of the employed from the unemployed will increase the welfare of both groups. A similar conclusion pertains to health.

Last, different results are found for marriage. Here higher regional marriage rates are associated with greater levels of well-being for the non-married, but lower levels of well-being for those who are married themselves. The welfare-maximising spatial distribution of marriage does not necessarily call for complete segregation, as for unemployment and health above, but turns out to be difficult to characterise other than that, and requires further work.


It should be emphasised that some of these results are still preliminary, and require further testing on different datasets. We have also taken no account of preference heterogeneity. It is entirely possible that some groups of individuals be very sensitive to income and income comparisons, for example, while others may be fairly indifferent to them. Most work does not take such heterogeneity into account, and assumes the same preferences over the whole sample. This is unlikely to be true (some recent examples include Clark *et al.*, 2005, Mancini *et al.*, 2008, and Pinquart and Schindler, 2007). In this case, we may see spatial sorting even when the spillover is generically positive or negative. To take the example of income, those who are especially income-comparison sensitive will be especially attracted to neighbourhoods with lower income reference groups. Equally not all of the unemployed may compare their labour-force situation to that of other local unemployed: these individuals will presumably be less tempted to locate together with other unemployed, and may prefer to live in low-unemployment areas, where job prospects are better.

The analysis described above related to spatial spillovers with respect to five different characteristics. There is however no reason to stop at these five, and the research could be extended to many other domains of individual life. It would also be of great interest to relate

the spatial predictions from well-being analysis to both individual-level location decisions (do these kinds of regional-level variables help us to explain why people move?), and observed distribution patterns, as measured by segregation indices.

Figure 1. The Well-Being Gap between those in Work and the Unemployed ($GHO_E - GHO_U$) and Regional Unemployment Rates.

BHPS Waves One to Seven. (Eleven Regions)



Key: GL = Greater London, RS = Rest of the South East, SW = South West, EA = East Anglia, EM = East Midlands, WM = West Midlands, NW = North West, YH = Yorkshire and Humberside, NT = North, WA = Wales, SC = Scotland.

Source: Clark (2003).

Table 1. Well-Being and Marital Spillovers in the BHPS.

	<i>GHQ-12</i>			<i>Life Satisfaction</i>		
<i>Married</i>	0.016 (0.009)	0.015 (0.009)	0.201** (0.050)	0.238** (0.011)	0.246** (0.010)	0.638** (0.061)
<i>Separated</i>	-0.371** (0.020)	-0.377** (0.022)	-0.377** (0.022)	-0.241** (0.021)	-0.230** (0.024)	-0.229** (0.024)
<i>Divorced</i>	-0.143** (0.011)	-0.138** (0.011)	-0.138** (0.011)	-0.095** (0.014)	-0.074** (0.014)	-0.072** (0.014)
<i>Widowed</i>	-0.147** (0.014)	-0.142** (0.015)	-0.140** (0.015)	-0.054** (0.018)	-0.045* (0.018)	-0.041* (0.018)
<i>Percentage(Married)</i>		0.063 (0.113)	0.236* (0.119)		0.203 (0.129)	0.571** (0.146)
<i>Percentage(Separated)</i>		-0.429 (0.358)	-0.429 (0.358)		0.096 (0.405)	0.083 (0.407)
<i>Percentage(Divorced)</i>		0.015 (0.225)	0.003 (0.225)		0.094 (0.208)	0.075 (0.209)
<i>Percentage(Widowed)</i>		0.125 (0.171)	0.108 (0.172)		0.180 (0.222)	0.130 (0.222)
<i>Married*Percentage(Married)</i>			-0.340** (0.090)			-0.725** (0.111)
Observations	189145	174217	174217	127917	115571	115571

Note: Ordered Probit Estimations. Standard errors (clustered at the region*year level) in parentheses. * significant at 5%; ** significant at 1%. All regressions also control for labour force status, individual yearly income, age and age-squared, sex, education, health, number of children, and region and year dummies. All percentage variables have been divided by 100.

References

- Blanchflower, D.G., Oswald, A.J., and Van Landeghem, B. (2009). "Imitative Obesity and Relative Utility". *Journal of the European Economic Association*, forthcoming.
- Clark, A.E. (2003). "Unemployment as a Social Norm: Psychological Evidence from Panel Data". *Journal of Labor Economics*, **21**, 323-351.
- Clark, A.E. (2008). "Happiness, Habits and High Rank: Comparisons in Economic and Social Life". PSE, Discussion Paper No. 2008-61.
- Clark, A.E. and Etilé, F. (2008). "Happy House: Spousal Weight and Individual Well-Being". PSE, mimeo.
- Clark, A.E., Etilé, F., Postel-Vinay, F., Senik, C. and van der Straeten, K. (2005). "Heterogeneity in Reported Well-Being: Evidence from Twelve European Countries". *Economic Journal*, **115**, C118-C132.
- Clark, A.E., Frijters, P. and Shields, M. (2008a). "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles". *Journal of Economic Literature*, **46**, 95-144.
- Clark, A.E., Knabe, A. and Rätzl, S. (2008b). "Boon or Bane? Others' Unemployment, Well-being and Job Insecurity". PSE, Discussion Paper No. 2008-67.
- Clark, A.E., Kristensen, N. and Westergård-Nielsen, N. (2009). "Economic Satisfaction and Income Rank in Small Neighbourhoods". *Journal of the European Economic Association*, forthcoming.
- Clark, A.E. and Lelkes, O. (2009). "Let us Pray: Religious Interactions in Life Satisfaction". PSE, Discussion Paper 2009-01.
- Clinch, J., Ferreira, S., Brereton, F., Moro, M. and Bullock, C. (2007). "Understanding and Measuring Quality of Life in Ireland: Sustainability, Happiness and Well-Being". University College Dublin, mimeo.
- Dolan, P. and Metcalfe, R. (2007). "Valuing Non-Market Goods: A Comparison of Preference-Based and Experience-Based Approaches". Tanaka Business School, Imperial College London, mimeo.
- Ferrer-i-Carbonell, A. (2005). "Income and well-being: an empirical analysis of the comparison income effect". *Journal of Public Economics*, **89**, 997-1019.

- Frank, R.H. (1985). *Choosing the Right Pond: Human Behaviour and the Quest for Status*. London and New York: Oxford University Press.
- Luttmer, E. (2005). "Neighbors as Negatives: Relative Earnings and Well-Being". *Quarterly Journal of Economics*, **120**, 963-1002.
- Mancini, A., Bonanno, G. and Clark, A.E. (2008). "Stepping Off the Hedonic Treadmill: Individual Differences in Response to Marriage, Divorce, and Spousal Bereavement". Columbia University, mimeo.
- Pinquart, M. and Schindler, I. (2007). "Changes of Life Satisfaction in the Transition to Retirement: A Latent-Class Approach". *Psychology and Aging*, **22**, 442-455.
- Powdthavee, N. (2009). "Ill-Health as a Social Norm: Evidence from Other Peoples' Health Problems". *Social Science and Medicine*, **68**, 251-259.
- Powell, A. and Sanguinetti, P. (2008). "What Determines Neighborhood Quality of Life: A Summary of the Results from Case Studies". IADB, mimeo.
- Van Praag, B.M.S. and Baarsma, B.E. (2005). "Using Happiness Surveys to Value Intangibles: the Case of Airport Noise". *Economic Journal*, **115**, 224-246.