30
PROVOCATION

Distilling Career Advice from the Happiness Literature

Robert H. Frank

For the past several decades, social scientists have been examining statistical measures of human well-being in an attempt to identify the conditions that promote satisfying lives. As a long-time spectator of the happiness literature and an occasional contributor to it, my aim in this essay will be to summarize how what we've learned might help people better plan their lives and careers. When I use "you" in what follows, I'll be addressing a young person who's currently trying to decide what to do in life.

In plotting your career, one of the first questions you'll confront is, "How much money will I need to be happy?" Existing studies don't serve up as clear an answer as you might hope, but there are nonetheless useful insights to be gleaned from a careful reading of those studies.

One pattern observed early on was that average income levels in a country tend to rise much more rapidly over time than the corresponding levels of average happiness. Many interpreted that finding to mean that having extra income didn't make people appreciably happier. But that interpretation was quickly challenged by the economist Richard Easterlin, who noted that if we examine the income-happiness relationship in a second way, income seems much more important.[1] Easterlin noted that in studies comparing *groups* of people at different points along the income scale, average happiness levels increased sharply as income increased. But the individual data points were very noisy, meaning that studies of individual happiness and income data didn't show nearly as clear a link between the two. There were plenty of poor people describing themselves as extremely happy, for example, and likewise many rich people describing their lives as miserable. But once income and happiness values were averaged for large groups of people, it became clear that the wealthier groups were much happier, on balance, than those with lower incomes.

Easterlin's explanation was simple and intuitively appealing. He said that the discrepancy between the two ways of looking at the income-happiness relationship stemmed from the fact that relative income differences account for more of the observed variance in happiness than absolute income differences. What matters most, he argued, is a person's percentile position in the income distribution at any moment in time. When average income in a country grows over time, people's respective positions in the income distribution tend to remain stable, which means we shouldn't expect average happiness levels to change much with income growth. But when we look at people at different points along the income scale at a single moment in time, we're by definition comparing the happiness levels of people with different rankings in the income distribution. And when we look at the data that way, people with higher incomes are clearly happier on average than their counterparts with lower incomes.

These findings have become a little less clear in the light of more recent work. The economists Betsey Stevenson and Justin Wolfers, for example, suggest that average happiness levels in a country do tend to rise significantly over time as average income grows.[2] They argue that in earlier studies, authors failed to detect this tendency because there weren't enough data to facilitate accurate comparisons. Scholars continue to debate the validity of these more recent findings.[3]

But even if measured happiness levels did not increase with income growth over time, it would be a mistake to conclude that income growth isn't beneficial. After all, societies with higher incomes tend also to enjoy cleaner air and water, better schools, less noisy environments, safer working conditions, longer life expectancy, and many other obvious benefits.

One way to summarize what we've learned about the income-happiness relationship is to ask how it might inform your choice between two parallel worlds that were alike in all respects except for any differences associated with their different income levels. A clear implication of existing studies is that if you'd occupy the same position in the income distribution in both worlds (say, the 50th percentile), you'd have compelling reasons for choosing the richer world. You'd probably feel at least a little happier in that world, but even if not, you'd be likely to live significantly longer.

What if you faced with a choice between being in the 25th percentile in a high-income society or being in the 75th percentile in a society in which your income would be lower in absolute terms? Here, the answer is less clear. If the income difference were very small, the 75th percentile position would probably be more satisfying. But for sufficiently large income differences, that conclusion could easily flip.

The choice just described closely resembles the actual choices you'll confront when you face several competing job offers. Figure 1 depicts three hypothetical offers, one from a firm with highly talented workers (Firm 1), another from a firm with workers of intermediate talent levels (Firm 2), and a third from a firm whose workers have relatively low talent levels (Firm 3). The dark lines in the diagram show the pay schedules for the three firms, which tell us how much workers with differing productivity levels will be paid. Like the schedules we observe in practice, they are upward-sloping, meaning that the more productive a worker is, the more he or she gets paid. But note that in each case, pay rises less than dollar-for-dollar with increases in productivity. This too is a feature almost universally observed in actual pay schedules.

Figure 1.

THE TRADEOFF BETWEEN LOCAL RANK AND PAY.

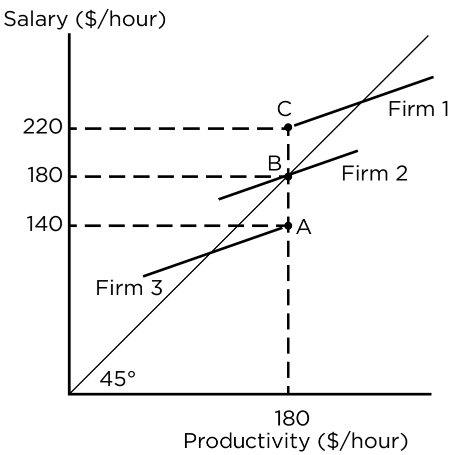

In this example, I've assumed that your productivity is $180/hr, meaning that if you worked for free, any firm that hired you would enjoy $180/hr in additional net revenue. If you chose to work in Firm 1, you'd occupy the job labeled C, making you the least productive worker in that firm. In that job, your pay would be $220/hr. If you instead chose to work in Firm 2, you'd occupy the position labeled B, placing you in the middle of that firm's productivity distribution, where your pay would be $180/hr. Finally, you could choose to work in Firm 3, where you'd occupy the job labeled A, be paid $140/hr, and be the most productive worker in your firm.

Which job should you choose? If all you care about is money, the job at C in Firm 1 would clearly be your best option. But in choosing that job, you'd also be the least productive worker in the firm, and if that prospect would make you sufficiently uncomfortable, you might want to consider one of the other options. By choosing job B in Firm 2, for example, you'd sacrifice $40/hr in pay but would move up to the middle of the co-worker productivity distribution. Or for an additional $40/hr sacrifice in pay, you could choose Job A in Firm 3 and be the most productive worker in that firm. In short, you face a difficult tradeoff between two things you value: your hourly pay, on the one hand, and your rank among your co-workers on the other.

The upshot is that your best choice among these three options depends on how strongly you feel about these two job dimensions. If you don't care much about your local rank among your co-workers, job C is probably your best option, since the extra pay would be enough to compensate for whatever negative feelings you might have about being the least productive worker. But if you're highly sensitive about interpersonal comparisons, the other extreme option might be best. In that case, the satisfaction you'd derive from being the most talented worker in Firm 3 might be more than enough to compensate for any misgivings you might have about your lower rate of pay. If you fall between these extremes, perhaps job B at Firm 2 would be your best bet.

One of the most important sources of compensating wage differentials involves how workers feel about their employer's mission

My point in presenting this example is not to identify one of the options as being better than the others, but rather to emphasize that your choice among jobs merits careful consideration of the many important ways that jobs differ along dimensions other than pay. The economist's theory of compensating wage differentials holds that the more pleasant a job is generally regarded to be, the less employers will have to offer in order to fill it.[4] (Or equivalently, the more unpleasant the job is, the more they'll have to offer.) The specific non-salary dimension highlighted in that example was your local rank among your co-workers. On the plausible assumption that most workers like high rank better than low rank, an employer that offers you the option of high rank can get by with paying you less.[5] Analogous compensating wage differentials have been documented for autonomy, opportunities for learning, workplace safety, and a host of other desirable job characteristics.

One of the most important sources of compensating wage differentials involves how workers feel about their employer's mission. Consider this thought experiment: You have two job offers, one to work for the American Cancer Society writing advertising copy for a campaign to discourage teenagers from smoking, the other for a tobacco company writing ad copy for a campaign to encourage teen smoking. If both positions paid the same and offered otherwise identical working conditions, which would you choose?

Long ago, when I posed this question to a group of Cornell University seniors about to enter the job market, almost 90 percent said they'd pick the American Cancer Society position. And when I asked them to report how much more the tobacco company would have to pay before they'd change their minds, they demanded a salary premium of more than 80 percent on average.[6]

These magnitudes shouldn't surprise you. If you're like most people, when you leave work in the evening, you'd feel better if your day's efforts had made the world better in some way or at least hadn't made it worse. As it turns out, compensating wage differentials for morally satisfying jobs are among the largest of any we observe in practice.

Many job seekers mistakenly focus primarily on the salaries of competing offers. That's a natural error, since salary is not only an important feature of any job, it's also by far the easiest one to observe and compare among jobs. But a central lesson of the happiness literature is that money explains only a small proportion of the observed individual differences in happiness levels. Your efforts to focus on the numerous other job characteristics that people find important will yield big dividends. If an employer is offering what seems like an unexpectedly high salary for someone with your experience and skills, consider the possibility that that she or he may want you to do things that most people would find noxious. In short, when shopping for the right job, be sure to pay careful attention not just to the salary you're offered, but also to the long list of other working conditions you care about.

A central lesson of the happiness literature is that money explains only a small proportion of the observed individual differences in happiness levels

In the same vein, I'll note that *Business Week, The U.S. News & World Report*, and other college ratings services might better serve their readers if they too paid closer attention to job characteristics other than salaries. One component of the ranking formulas employed by those services is average graduate starting salaries, which gives schools an incentive to encourage students to pursue high salaries to the exclusion of other goals. We should instead be trying to encourage students to focus on a much broader suite of relevant job characteristics.

Findings from the happiness literature can also be combined with observations about recent labor market trends to gain insight about how better to position yourself to land a job you'll find satisfying. As Charlie Munger, the vice-chairman of Warren Buffett's Berkshire Hathaway, has written, "The safest way to try to get what you want is to try to deserve what you want."[7] But what must you do to deserve the job you want, one whose combination of high salary and attractive working conditions you're most likely to find personally satisfying?

To merit high pay, you have to be able to produce goods or services that buyers value highly. That can happen even if only a few buyers value what you do, if those buyers value it highly enough. But you're more likely to become highly valued if what you do is valued by many buyers. Changes in technology have been making it easier than ever for individuals to serve larger numbers of buyers. But those same changes have also been concentrating demand on a dwindling number of the best performers in every arena.[8] For example, the best federal income tax accountant in a city could once serve only a small number of that city's most important clients, but the author of the best income tax software program can now serve an unlimited number of clients nationwide. That's why the author of the best tax software program earns so much more than the best local accountant used to. One side effect, however, has been to reduce the demand for local accountants. So as technology has extended the reach of the best performers in almost every arena, their income growth has far outpaced that of others.

One implication is that if you want to earn a lot of money, your first priority should be to find something that you can develop deep expertise at doing. In almost every domain, people who get really good at what they do capture a much larger share of total incomes and leave correspondingly smaller shares available for others. Become an expert at something!

That's obviously easier said than done. Those who have studied expert performance estimate that it takes many thousands of hours of difficult practice to develop true expertise at any task.[9] That's why my first response when students ask me for advice about how to succeed is to ask whether there's any activity they've ever engaged in that they found completely absorbing. On reflection, most answer affirmatively. I then suggest that they try to prepare themselves for a career that entails tasks as similar as possible to the activity they named. Even if it's one that doesn't normally lead to high financial rewards, I tell them not to worry. My point is that becoming an expert is so challenging that you're unlikely to expend the effort necessary to do it unless the task is one that you love for its own sake. But if it's a task you love, the process of becoming an expert at it will be highly rewarding quite apart from whether it leads to high pay.

In the happiness literature, "flow" has been identified as one of the most deeply satisfying human psychological states.[10] Flow occurs when you're immersed in an activity to such an extent that you lose track of the passage of time almost completely. If you can land a job that enables you to experience substantial periods of flow, you'll be among the most fortunate people on the planet. What's more, as the years pass, you'll almost surely develop deep expertise at whatever it is you've been doing.

Flow occurs when you're immersed in an activity to such an extent that you lose track of the passage of time almost completely. If you can land a job that enables you to experience substantial periods of flow, you'll be among the most fortunate people on the planet

At that point, you may find that your services have become extremely valuable economically, even if not many people in any given location place high value on what you do. Again, that's because technology has been steadily extending the geographic reach of those who are best at what they do. If even a tiny fraction of a sufficiently large set of buyers cares about your service, you may be worth a fortune.

There's of course no guarantee that you'll become best at what you choose to do, or that even if you do you'll find practical ways of extending your reach enough to earn a big paycheck. But if you'll think about your choices carefully, you'll see that the downside risk of my recommended course of action is limited. By choosing to concentrate on a task you love, you'll be able to report truthfully that you enjoy the substantial proportion of your life that you spend at work, which is much more than billions of others can say. And the happiness literature should also reassure you that it's possible to live a very satisfying life indeed even if you don't earn a lot of money.

Bottom line: One of the best ways to spend the extra money you could have earned by taking an unpleasant job is to savor the experience of the more satisfying one you chose instead.

NOTES

1. See Richard Easterlin, “Does Economic Growth Improve the Human Lot?” in *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz,* ed. Paul David and Melvin Reder (New York: Academic Press, 1974) and Richard Easterlin, “Will Raising the Incomes of All Increase the Happiness of All?” *Journal of Economic Behavior and Organization* 27 (1995): 35–47.
2. Betsey Stevenson and Justin Wolfers, “Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox,” NBER Working Paper No. 14282, August 2008.
3. Eric Falkenstein, “Stevenson and Wolfers’ Flawed Happiness Research,” Falkenblog (blog), July 7, 2013, http://falkenblog.blogspot.com/2013/07/stevenson-and-wolfers-flawed-happiness.html.
4. This venerable theory traces all the way back to Adam Smith, *The Wealth of Nations* (New York: Modern Library, 1937), 100.
5. You may be wondering how Firm 1 can pay you $220/hr if your productivity is only $180/hr. The answer is that high rank, a job characteristic that workers value, could not exist if others were unwilling to accept positions of low rank. The $40/hr premium you get at C compensates you for the negative consequences of your holding a position of low rank, and Firm 1 recovers that premium by paying its top-ranked worker $40/hr less than the value of his productivity. He is willing to accept that pay cut because of the value he places on high rank.
6. Robert H. Frank, *What Price the Moral High Ground* (Princeton: Princeton University Press, 2004), Chapter 5.
7. Charles Munger, “USC Law Commencement Speech,” University of Southern California School of Law, Los Angeles, CA, May 13, 2007.
8. Robert H. Frank and Philip J. Cook, *The Winner-Take-All Society* (New York: The Free Press, 1995).
9. K. Anders Ericsson and Jaqui Smith, eds., *Toward a General Theory of Expertise: Prospects and Limits* (New York: Cambridge University Press, 1991).
10. Mihaly Csikszentmihalyi, *Beyond Boredom and Anxiety: Experiencing Flow in Work and Play* (San Francisco: Jossey-Bass, 1975).