

THE LAST WORD ON PRODUCTIVITY AND TELECOMMUTING

By Gil E. Gordon

Copyright (c) 1997, Gil Gordon Associates. All Rights Reserved.
For more information, visit www.gilgordon.com
This document may be downloaded and/or reproduced freely, but the title, author's name and copyright information must be included as shown here.

1. INTRODUCTION: THE "LAST" WORD?

Since getting started with the telecommuting field in 1982, I have heard or read an endless stream of questions, comments, and allegedly profound pronouncements (including those from my mouth and/or keyboard) about how telecommuting does or doesn't lead to increases in productivity.

Similarly, there has been a parallel stream of discussion about doing cost-benefit analysis for telecommuting. These are two very important topics, and not surprisingly are on the minds of executives and managers trying to figure out what to do about telecommuting.

In the past, I have written on these topics several times in my TELECOMMUTING REVIEW newsletter, and have also addressed them in part in the FAQ section of this Web site. The purpose of this commentary is to collect all these observations, expand on them, and attempt to provide a comprehensive and balanced look at these issues.

Is this really the "last word" - the definitive answer to the productivity dilemma - the ultimate proof about how wonderful telecommuting is? Of course not. I don't think anybody will ever write the last word on these topics; we all continue learning about them as we gain experience. But this is at least a good summary of what I feel are the basic issues involved, and a common starting point for our continued investigation.

2. PRODUCTIVITY DEFINED

Section Summary: The industrial-era models of productivity simply don't apply to today's knowledge workers. As a result, we know less about their productivity - and its measurement - than we'd like to think we do.

The classic definition of productivity is based on a comparison of output produced by a given amount of input - typically, hours of labor. This is

the view of productivity we can all easily understand. For example, a worker stands at a machine and produces an average of ten units of output per hour. If the worker works faster, the tools are improved, raw material defects are eliminated, or anything else changes so that eleven units can be produced per hour, that's a 10% productivity increase. Nothing could be simpler.

If we move off the factory floor and into the office, there are many jobs that lend themselves to this kind of productivity measurement. A person processing invoices or insurance claims forms, for example, can do five, or eleven, or twenty per hour; change the methods, add some training, improve the systems, etc. and those numbers should increase. Once again, we see a very evident productivity increase.

It should be painfully obvious why this approach to productivity, and productivity measurement, is of little value for today's professional-level office workers. It is these workers (such as programmers, accountants, market researchers, lawyers, engineers, designers, and so on) who make up a large and growing segment of almost every developed nation's payroll, and thus they are the focus of much of this commentary:

- Few, if any, of these workers produce any number of "units" of work per hour, or perhaps even per day or per week. Their work output cannot sensibly be measured with the same "unit" mentality as we use for factory workers or even for lower-level office workers.

- There isn't a direct correlation between units of labor and units of output for these professional-level, or "knowledge" workers. The factory worker standing at a drill press can reasonably be expected to produce another ten units if we add another hour to the workday, but the knowledge worker doesn't have this same kind of directly-proportional increase.

- More activity doesn't necessarily mean more productivity for knowledge workers. The professional who works longer hours, attends more meetings, spews out more memos, and otherwise appears to be working harder may in fact be creating less, not more. Does the person in the office next to you who works a sixty-hour week create 50% more value for his/her paycheck than

the person who works a 40-hour week? Probably not.

I'm definitely not a productivity expert, but it's pretty clear to me that we don't really understand what we mean when we use the term "white-collar productivity." That term connects a farm- and factory-era work measurement with information-age occupations, and it just doesn't make sense. It's a habit, it's convenient, it's popular, and it's always trendy to talk about "improving white-collar productivity," but I just can't figure out what it means.

If we can't measure white-collar productivity in the office, we certainly can't measure it away from the office. Herein lies the big problem with measuring the productivity effects of telecommuting among knowledge workers. Keep that in mind; we'll return to this point shortly.

3. WHY "EFFECTIVENESS" MAY BE A BETTER APPROACH

Section Summary: Instead of struggling with the productivity problem, maybe it's time to look for a broader and more sensible way to describe what knowledge workers actually do. A collection of several factors - labeled here as "effectiveness" - might help us better describe and measure what these workers do - both in AND out of the office.

Because I have never been able to resolve this productivity-measurement dilemma for information work, I have chosen to avoid it. I'm sure there are many people who have spent countless hours fine-tuning some very sophisticated approaches for measuring office work like we used to measure factory work; for me, this is not a useful pursuit.

Instead, I argue that we're better off taking a broader view. The fundamental problem with "productivity" is that it always seems to come back to measuring quantity of work produced - but this is grossly inadequate for knowledge work. The sheer amount of work produced is certainly important, but it's not all that matters.

My preference is to talk about "effectiveness" instead of productivity. You can use whatever word you like here, but to me, "effectiveness" is a more inclusive term that conveys the idea that we're trying to measure more

than just the quantity of work. This is a kind of market-basket approach that includes:

- Quantity (how much gets done)
- Quality (how well it gets done)
- Timeliness (when it gets done)
- Multiple priorities (how many things can be done at once)

Whenever I think of an exemplary office worker (let alone an exemplary telecommuter), it is the combination of all four factors that counts - not just any one of them. It's important how much gets done, but it's equally important that the work is done well, done on time, and done in a way that lets multiple projects or responsibilities all get handled at the same time.

So, an "effective" knowledge worker is one who scores well on these four criteria. You'll note that I have still included the "how much" factor - and everything I said before about the difficulty of measuring white-collar work still applies. But the difference here is that we put that quantity factor in the context of a bigger picture, and thus we aren't so reliant on (or so affected by) a simple increase in the amount of work output.

Another reason why I prefer this collection of measures - and you can certainly change or expand my list of four - is that it forces us to look at the relationship among all the factors. Here's an example: consider an employee who turns out a lot of work with good quality, but is consistently late and has trouble juggling multiple priorities. Would you say this person is a good worker just because of the high-volume, high-quality output? Of course not - you'd appreciate those strengths, but when you look at the four factors together you see the total picture has some serious deficiencies.

This productivity vs. effectiveness discussion is directly relevant to telecommuting because it is the collection of those four factors which seems to account for the gains we see in telecommuters' work. It isn't just the fact that they might be doing more work; it's also the fact that they're doing better work, in less time, and with a better ability to handle multiple assignments. If we're really interested in seeing whether telecommuting has any payoffs, we need to look at the full scope of the work performed - just as we should be doing for everyone in the office.

And that observation leads directly to ...

4. THE PRODUCTIVITY (OR "EFFECTIVENESS") MEASUREMENT MYTHS AND TRAPS

Section Summary: No matter how you define it, there are some inherent myths and traps whenever we try to measure what knowledge workers do. These problems affect our understanding of this work in an office - and make it even harder to determine exactly what telecommuters are doing away from the office.

The first myth or trap to recognize is that we can't measure productivity, or effectiveness, or whatever you choose to call it, away from the office any better than we can in the office. I have witnessed the search for this Holy Grail of telecommuting measurement for years - a search for the ultimate metric, tool, or scorecard to measure how well telecommuters are performing. If we don't know what they're doing when they come into the building five days a week, how can we possibly know with any more precision what they're doing when they telecommute a few days a week?

Interestingly, this realization is one of the "accidental" benefits of telecommuting. The search for that perfect measurement method for telecommuters sometimes causes smart managers to realize that they need to become more disciplined and serious about work measurement in the office - something that is long overdue in most organizations.

Ready for some more myths and traps? Here they are:

Confusing Means vs. Ends - Even if you can figure out what to measure, there's a risk that you are only measuring the work processes that are assumed to lead to the desired results, not the results themselves. A classic example of this means-end inversion is the counting of lines of code produced daily by a programmer. Since the goal is to produce software, which is made up of lines of code, it seems logical to assume that the programmer who generates more lines of code per day is more productive and more valuable.

Not really - for two reasons. First, more code isn't necessarily better code; sometimes a more efficient processing solution relies on fewer lines, not more. Second, if you emphasize lines of code as the measure of an employee's value, pretty soon you will have trained your programmers to

push out more and more lines of code because they think that's what they will be rewarded for - even though all those lines together add up to lousy software.

The High-Stakes Poker Game - Measuring productivity or effectiveness in a telecommuting program is an "in for a dime, in for a dollar" kind of game. I don't think you can do this unless you're willing and able to do it the right way - which means measuring both the telecommuters and a non-telecommuting control group, and measuring both of them before and after you introduce telecommuting. No matter how sophisticated a series of measures you develop, you can't measure only the telecommuters and attribute their performance gains solely to telecommuting unless you also see what happened with everyone else's performance.

Do It Right or Don't Bother - Because it isn't easy to do good measurement, I believe you should either commit to doing it the right way or simply forget about it - assuming you're trying to come up with some defensible measures. If you use faulty measures or faulty methods, you leave yourself open to some very valid criticisms. Also, many organizations choose not to do rigorous evaluation even if they know how, because they feel it will result in one group of employees looking like "failures" through no fault of their own. Think about that: if the telecommuters are shown to be undeniably more effective (however you define it), does that mean that your otherwise hard-working office workers now look like they're slacking off?

5. COST-BENEFIT ANALYSIS MYTHS AND TRAPS

Section Summary: It's important to compare the costs and benefits of telecommuting - but it's even more important to define both of them correctly and comprehensively, and to steer clear of some common problems.

Let's continue with this theme of separating fact from fallacy and flawed thinking. You can't talk about measuring productivity or effectiveness, which are presumed to be among the benefits of telecommuting, without also talking about the costs involved in creating those benefits. I have found

that cost-benefit analysis of telecommuting can be as tangled as the productivity issue is.

What Benefits Are You Measuring? - I'm often asked, "what benchmarks or evaluation measures should we use to determine if our telecommuting program is a success?" This should be an easy question to answer - if the program was implemented with a sound business rationale.

The measurement should be based on the reasons for starting the program in the first place. For example:

- If you're trying to reduce employee turnover, measure whether or not turnover has dropped;
- If you're trying to cope with a shortage of office space, measure whether or not you can accommodate your staff adequately in your existing space at existing expense levels;
- If you're trying to boost sales to your catalog-order customers who call late at night (and who previously had to wait a long time on hold for an agent to answer), measure how those sales have changed.

In all cases, keep in mind the points made earlier about measurement myths; your turnover might have decreased after you implemented telecommuting, but it may have decreased for ALL employees, not just telecommuters. If that's the case, there is probably some other explanation besides telecommuting.

Sharpen Your Pencils - When comparing costs and benefits, it's almost always easier to measure the former than the latter. That is, you can tell to the fourth decimal point how much it costs to provide a PC and software for a telecommuter; you probably can't tell with the same precision how much it costs you to recruit and retain that telecommuter. In that example, it's likely that the benefit of avoiding employee turnover via telecommuting might save anywhere from 50% to 150% of the employee's annual salary - which is a lot more than the PC and software cost.

The basic problem is that most organizations are better at tracking costs than benefits; as a result, the costs are more obvious and appear to be overwhelming at times. The key to doing cost-benefit analysis is to determine what tangible and intangible benefits result from telecommuting,

and making some conservative assumptions about the dollar impact of those benefits. Then and only then can you do a reasonable comparison of what it costs vs. what you gain with telecommuting.

Initial Vs. Ongoing Costs - The last, and most pervasive, cost-benefit measurement problem is based on the difference between the visible start-up costs and the less visible recurring costs. If you don't understand this difference, your cost-benefit analysis will be flawed.

Let's assume a telecommuter needs a computer and related supplies and equipment that cost US\$3500. If you start with ten telecommuters, that's a start-up cost of US\$35,000 for equipment alone - which sounds like quite a bit of money. But you must remember that these equipment costs are almost always treated by the accountants as a capital investment, and with the value of depreciation don't actually cost your organization US\$35,000 off the bottom line. Remember too that these are almost always one-time costs - you pay the bill to set up the telecommuters and that's it.

By comparison, consider the monthly telephone bills you'll undoubtedly face with telecommuting. If you spend US\$200 per month per telecommuter, that adds up to US\$2400 per year, or US\$24,000 for those same ten telecommuters.

The big difference is that the phone bills are treated as expense items, and have no depreciation value. So, a dollar of phone bill does cost you a dollar off the bottom line. And these are recurring costs - month after month you face these phone bills, for as long as the telecommuters are working.

The implications should be clear. First, you must analyze not only the start-up costs but also the recurring costs - and the accounting implications of each. Second, your efforts to control the costs of telecommuting should be directed at both start-up and recurring costs. If you don't, you may find that your cumulative phone bills after 12-18 months of telecommuting can easily exceed the equipment costs. Thus, it's important to look for ways to keep those ongoing expenses under control, even if it might mean a higher initial equipment investment.

There's much more to consider when doing your cost-benefit analysis, but these three traps are the biggest ones to avoid. Doing a cost-benefit

analysis isn't hard; doing a good one is. Take the time to do your analysis correctly and your assessment of the payoffs will be better respected and understood.

6. THE POLITICALLY-CORRECT NEED TO MEASURE

Section Summary: There may sometimes be a big gap between what managers SAY they're looking for and what they REALLY want. If you don't separate the two, you might waste a lot of time and effort - and your telecommuting program won't be successful.

My last comment on this entire topic of productivity and cost-benefit analysis has nothing to do with money, and everything to do with attitudes.

It has been my experience that some managers and executives disguise their concerns about telecommuting. Even though they might understand the potential business benefits, these managers have an underlying resistance to the concept. They don't believe that people will work unless they are watched, and otherwise are not as willing to trust people to telecommute successfully.

However, here in the late 1990's it is definitely not politically correct to simply say, "I just don't trust my people to work at home." So, the manager who feels this way might express that concern in terms of more palatable questions like "what's the return on our equipment investment for telecommuting?" or "what kind of productivity gains can we expect?"

While these are often legitimate questions, they may not be the right questions. It's easy to spend hours or weeks of time chasing after the answers to these questions, only to find that the manager then asks, "Okay, but what about the long-term effects of telecommuting on customer service?" or some other question that sends you off on yet another chase for more data.

There's no easy answer to this dilemma. One approach I've used is to pose this question to the potentially mistrustful manager: "If I can show you that you will at least break even, and perhaps gain something, with telecommuting, will you be ready to implement it immediately?" If the answer is a resounding "Yes!", then you have the obligation to do your numbers and make the business case. But many times, the response will be

something like, "Well, I'm not quite sure I still have some questions ... we might not be quite ready for it yet." These are exactly the responses you want, because you can now deal with the hesitation and concerns behind them instead of spending lots of time going after data that won't affect the decision.

7. IN CLOSING ...

One of my favorite teachers in graduate school used to describe some of our writing as "squid-like" - meaning that we, like a squid, were moving backward while discharging vast quantities of ink. I've discharged vast quantities of electrons, not ink - and I hope you feel it has been in the right direction. Please don't let my "last word" be the end of this discussion - let me hear from you with your own comments.