How Does The Workforce of a Company Respond to Changes In Demand?

By

Harry Cassady

Wilson High School
Modeling Dynamic Systems

May 17, 2007

System Dynamics Advisor: D. Fisher
Mathematics Department
Wilson High School
Portland, Oregon

Introduction

In order for a business to be profitable, a balance needs to be maintained between the number of employees and the amount of work that must be completed. If a company hires more employees than necessary to complete the work, their costs will go up with no gain in productivity. Conversely, if not enough employees are hired, costs will go down, but profitability should decrease since not all of the work will be done.

I decided I would like to learn more about how different factors can affect the balance of employees, and what measures companies can take to achieve the optimal balance between the amount of work vs. number of employees. In order to better understand this complicated system, I built a computer model that simulates the number of employees a company needs to hire depending on certain factors, such as the amount of work that needs to be completed.

There are some things that one can assume without building a model. For example, if there is a sudden increase in the amount of work that must be completed, more employees will need to be hired. (Figure 15) While this may seem obvious, other factors may not be so apparent. For example, more employees may be needed to compensate for a backlog in work during periods in which too few employees had been hired.

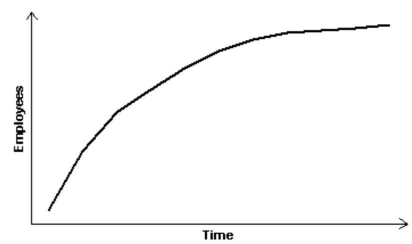

Figure 15: Increase In Work Over Time

Let's say that you are the manager of a factory that produces handcrafted pencil sharpeners. Each week your factory produces 1000 pencil sharpeners, and one worker will build ten pencil sharpeners per week. What happens if a wave of nostalgia for Number 2 pencils sweeps the nation and there is a sudden increase in the demand of pencil sharpeners? How fast will the work pile up? This is the type of question I am trying to answer by building this model.

By building this model, I hope to create a better representation of what really happens when changes happen in the workplace. This can help companies understand how to manage their workforce, and prepare for future changes in the incoming workload.

Building the Model

The first thing I did in building my model was to put down a simple structure containing some of the important stocks and flows I was planning to use.

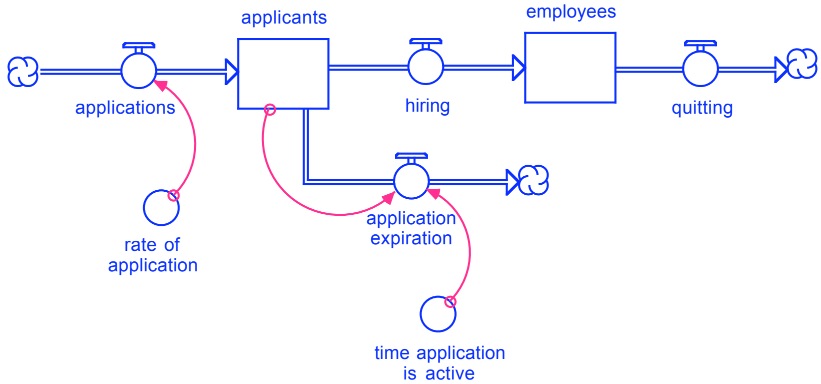

Figure 1: Original Sketch

In my first sketch, I included the current number of employees at the company, along with the current number of active job applicants. At this point, I realized that including the number of job applicants would cause complexity of the model to grow way beyond the scope of the model. Therefore, I decided that for this model I would omit the application sector, along with the number of people quitting. I then created an information delay to represent the fact that the process of hiring new employees, and firing current employees takes time.

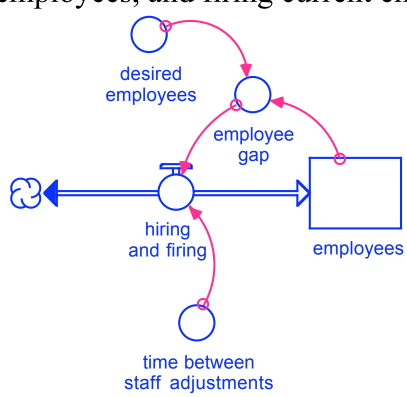

Figure 2: Final Employee Structure

With the employee sector built, I then needed to create the other two main structures in the model. These structures are employee satisfaction, and the amount of work that needs to be done. The work structure contains an inflow of new work, and an outflow of work completion. The new work is a manual test input, while the work completion is affected by how much work one employee does each week multiplied by the total number of employees at the company.

The perceived satisfaction structure is a basic information delay structure. This is because when conditions change for the better (or worse), it takes time before the effect on employees becomes noticeable. Both of these structures can be seen in figure 3 on the following page.

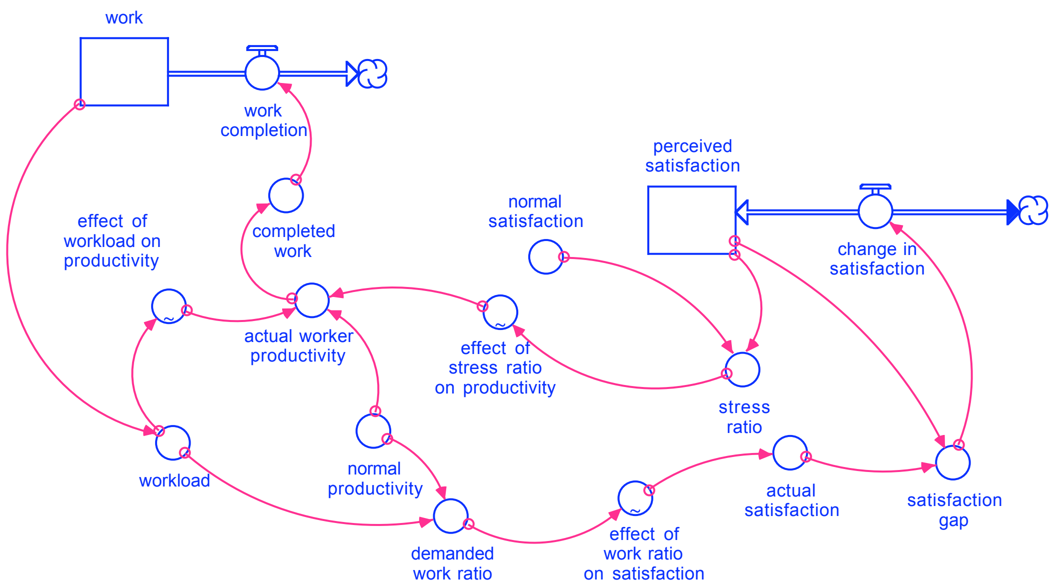

Figure 3: Work and Satisfaction Structure

With the basic model structures set up, I began connecting the different sectors. First, I connected the employee structure to the work structure. [figure 4.] I did this by creating a workload converter. This gives the total amount of work each person would need to do in order for all the work to be completed. The workload then affects the productivity of each employee (see documentation for details.)

I then began to connect the work to the employee sector. [figure 5] The number of desired employees is determined by dividing the amount of work needed to do by the productivity of one worker. This gives us the total number of employees needed to complete all of the current work, and the current rate of productivity.

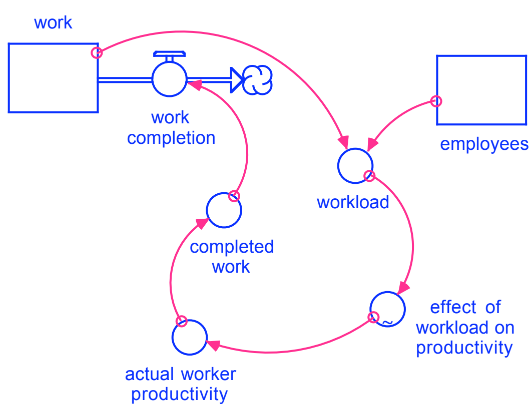

Figure 4: Connecting Employees to Work

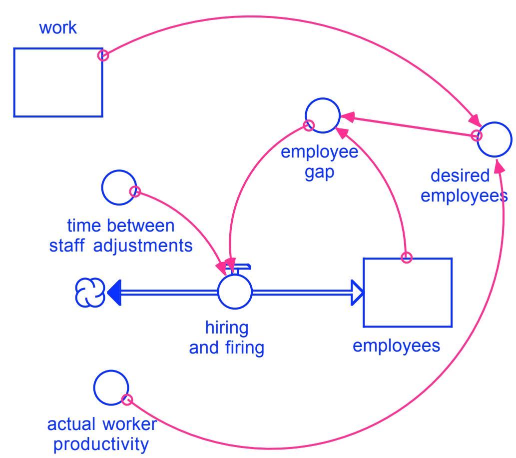

Figure 5: Connecting Work to Employees

The final process in building the model is connecting the satisfaction sector to the work sector (satisfaction does not directly affect the level of employees.) The effect of work on satisfaction is done by setting up a work demanded ratio. This ratio tells us whether more, or less work is being asked of the employees compared to a "normal workload" value. This then alters the satisfaction depending on whether people are being overworked, or under worked.

Satisfaction affects the amount of work done by each employee. I did this by setting up a stress ratio, telling us whether people are more or less stressed than normal. This then affects the average worker productivity, and hence the amount of work completed.

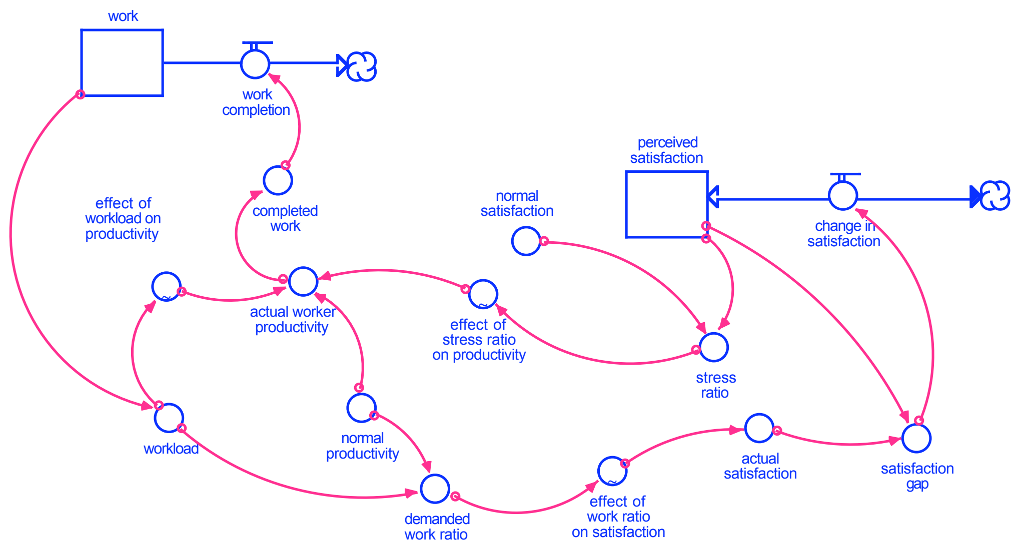

Figure 6: Connecting Work and Satisfaction

The final thing that I did was connecting the wage structure (a structure I decided to use only as a test input) to the satisfaction.

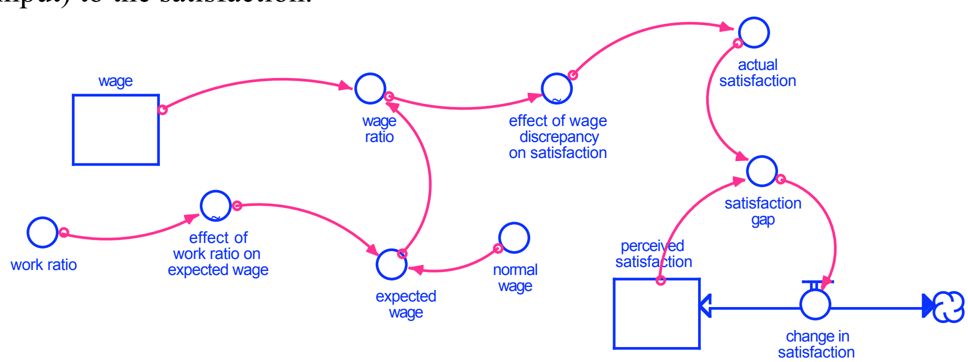

Figure 7: Connecting Wage to Satisfaction

I did this by creating a converter that represents the expected wage, defined by the amount of work that each employee is doing. I then created a wage ratio that tells us whether employees are getting paid more or less than they think they should be getting paid. In turn, this then affects the actual satisfaction of each employee.

Now that all of the sectors are built and connected, the model is complete. The full model appears on the next page; the equations with documentation are listed in the appendix.

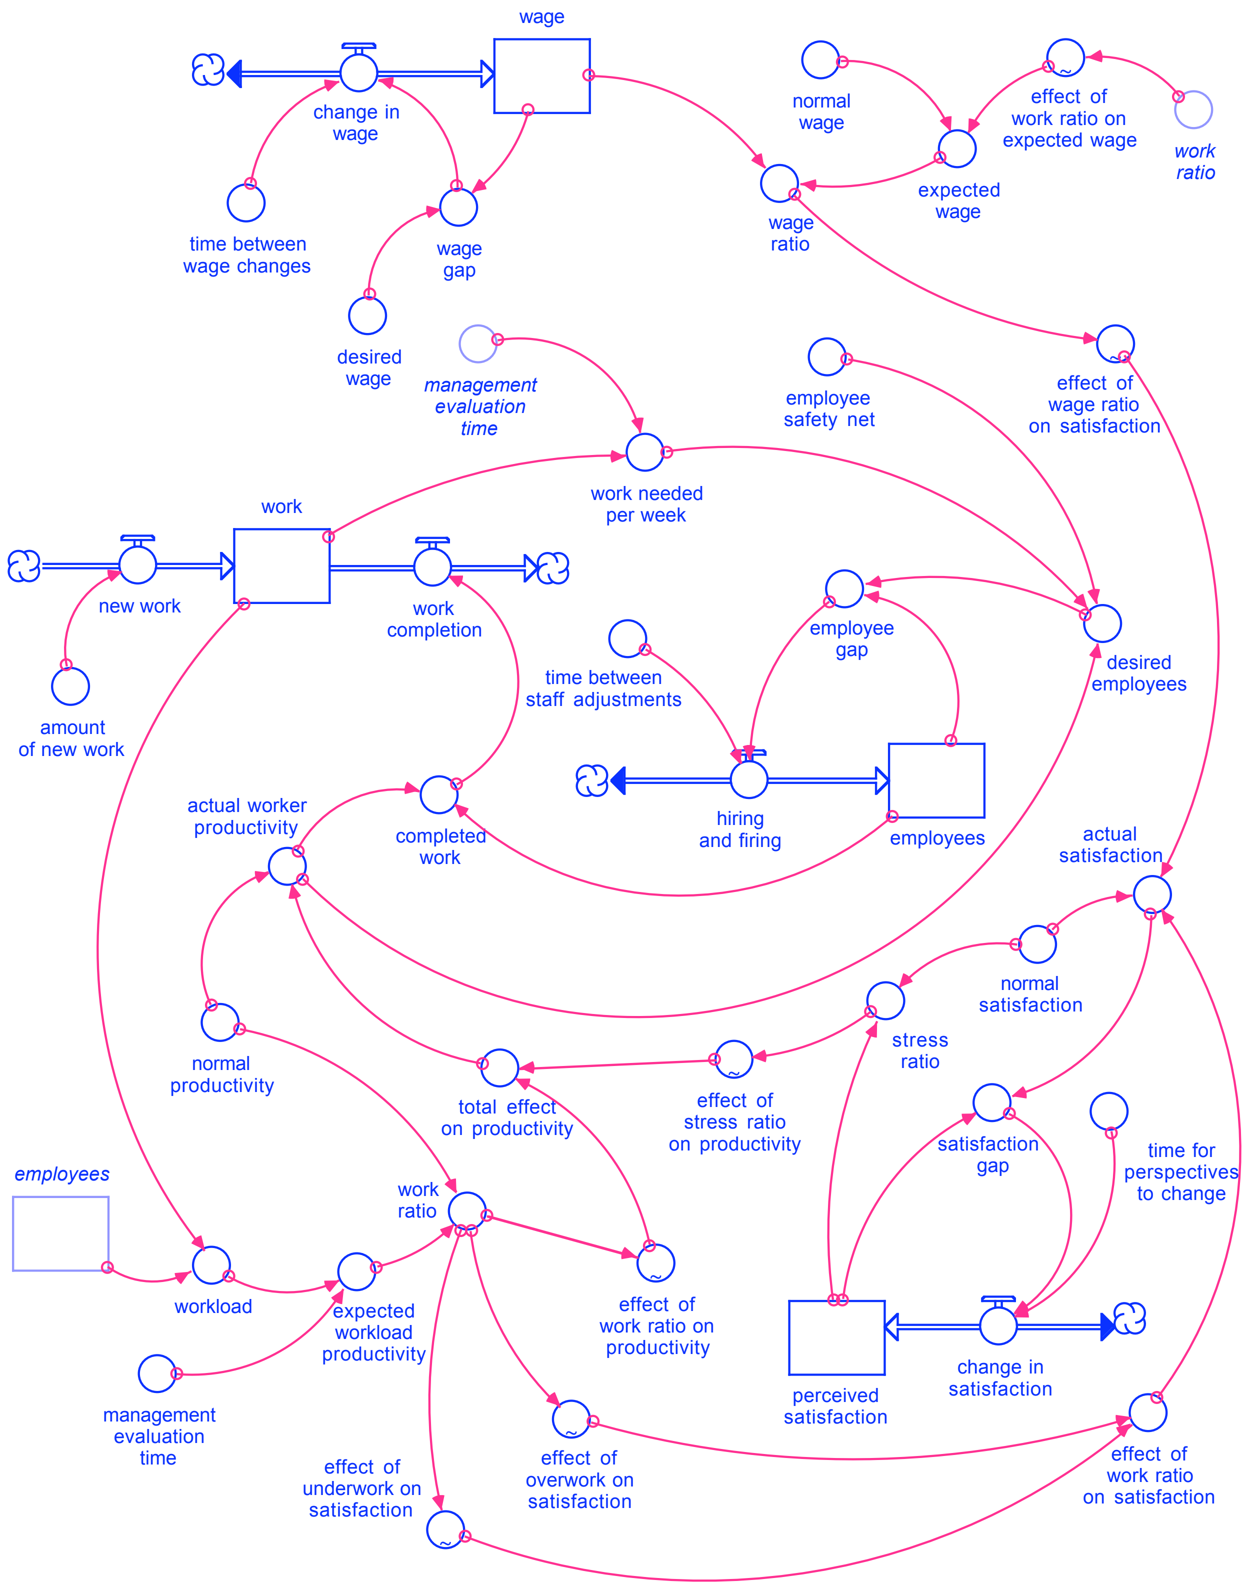

Figure 8: The Completed Model

Simulation Time and DT

For all of my tests, I chose to run the model with one unit of time equal to one week, and a simulation time of 52 units of time, or one year. The DT is set to 0.25. The reason I chose a time unit of a week is because with this model, we want to observe the effects on the system over between one to two years. Since you need multiple calculations per year, you need a time unit less than one year, but if you were to make the time unit one day, you would have to many calculations per year.

I chose to run my tests for one year, because this seemed to be a long enough time for everything to fall back into equilibrium, without there being a large period of equilibrium, distorting the curves. For different situations, one might choose to raise or lower the run time. For example, if you were studying the long-term effects of a very slow decreasing productivity, you would probably raise the runtime to five, or even ten years.

In my tests, I used a DT of 0.25. The reason I chose this, is because it gives an accurate representation of the curves being generated, without overloading the model with too many calculations. It would also make sense to use a DT of either 1/5 or 1/7, so that each calculation was reported on either a day of the year, or business day. I chose not to do this due to the fact that DTs not equal to the inverse of a number squared tend to create nasty rounding errors, due to the binary nature of the modeling program.

How I Know it Works

A model is invalid if one does not have confidence that it works properly. There is a number of different things that give me confidence that my model is working as it should.

In building the model, I made sure to set all of the stocks to allow non-negative integers (a cheat some system dynamists use which allow some more lazy structures). This means that the modeling program will not prevent stocks from becoming non-negative, and the actual structure of the model has to do this instead. Since at no point during my tests did any of the stocks drop below zero, I have confidence that the structure of the model is working as it should to prevent these negative values.

During my tests, the model never exhibited impossible behavior. Impossible behavior would be things such as 6 million employees needed to do 100 units of work (or exponential, unrestricted growth of employees), or a negative number of employees. The fact that the structure of the model keeps all values within acceptable values gives me confidence in the structure of the model.

Another important thing (that many people overlook) that gives me confidence in the model is that all units match up. You could go through the entire model, and through each calculation, every unit makes sense. This gives me confidence in the model, because it means that everything included in the model logically makes sense. Many times during the process of building the model, I had to change structures in order to make the units make sense, and these changes far improved the model once I thought about it.

The final reason I have confidence in the model is because all of the logic makes sense. I understand, and can explain why each part of the model makes sense, and needs to be there. Without logic to back you up, you can be sure that no one will be willing to take your model seriously. Therefore, the fact that all the logic makes sense, gives me confidence that the model is running as it should

Model Feedback and Loops

Examining feedback loops can greatly help you understand the structure and behavior of a model. Therefore, I will explain some of the most important feedback loops in my model.

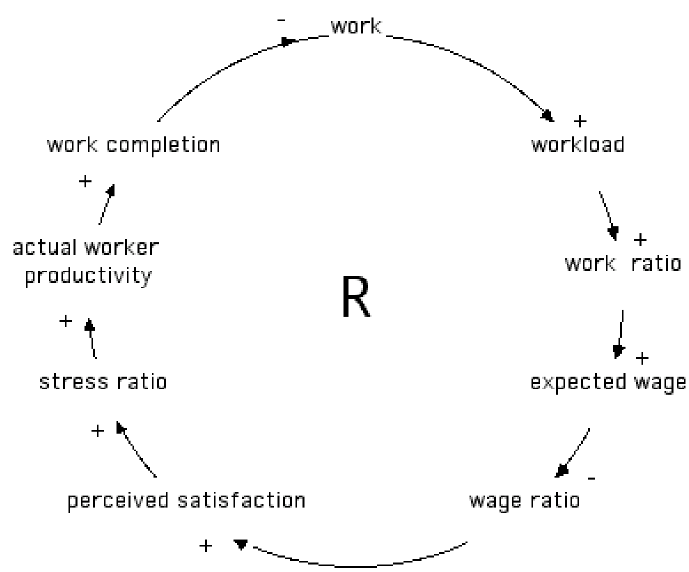

Figure 9: Work-Satisfaction Loop

The work-satisfaction loop is a loop where the amount of work to be done affects the satisfaction, which affects the worker productivity, affecting the amount of work to be done. So what happens when the amount of work increases? When the amount of work increases, employees need to do more of it to get the job done, therefore increasing the workload. This causes the work ratio to rise, which means that people are being asked to do more then they normally do. This causes the expected wage to rise, since employees expect to be paid more money to do more work. But because the actual wage has not changed, the wage ratio then decreases, meaning people expect more money than they are getting paid. This causes the perceived satisfaction to decrease. This then causes the stress ratio to decrease, meaning that people are less happy than normal. When people are unhappy, they get frustrated and tend to spend more time grumbling and use more sick leave, and performing with less motivation. This causes the actual worker productivity to decrease, which in turn decreases the total amount work that is completed each time period. When less work is completed, the total amount of work to do increases. Since increasing the total amount of work to do eventually comes back around increases it even more, this is a reinforcing feedback loop.

Another loop is the satisfaction-employee loop. It takes the exact same route as the work-satisfaction loop [figure 9] up until actual worker productivity, at which point it splits off. Therefore, I am omitting in-between work and actual worker productivity.

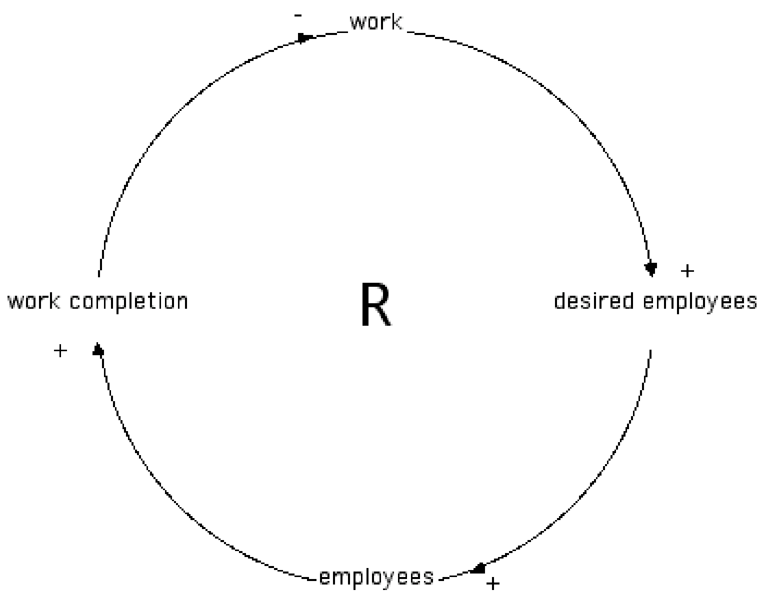

Figure 10: Work- Employee Loop

As we know from the work-satisfaction loop, when the amount of work goes up, worker productivity goes down. Since it now takes more employees then it did before to complete the same task, you now need more people to do the job. Therefore, desired employees goes up with number of employees not far behind. Since you now have more employees at your company, more work can be completed. Therefore both completed work, and work completion go up. Since more work is completed, the total amount of work goes down. Since an increase in the amount of work eventually causes the amount of work to decrease, this is a counteracting feedback loop.

The work-productivity loop is part of the relationship between the work, and number of employees.

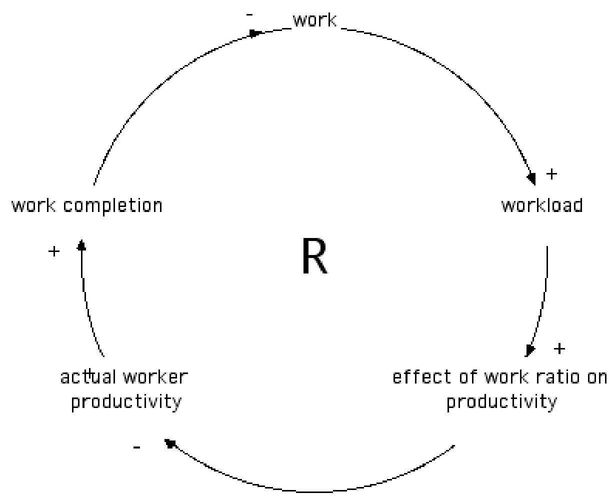

Figure 11: Work-Productivity Loop

When the amount of work goes up, the total amount of work that needs to be done each week also goes up. This increases the number of desired employees, which then increases the actual number of employees at the company, assuming the personnel department is doing its job. With more employees, more work is completed and work completion goes up. This causes the total amount of work to decrease. That means that this is a counteracting feedback loop.

The final feedback loop that I am going to explain in detail is an odd one. This feedback loop is the workload-productivity loop. The reason that this feedback loop is different is because it is actually a combination of both a counteracting, and reinforcing feedback loop. There two reasons why I decided to combine these two loops. The first reason is because they are both modeling the exact same thing: How overworking workers affects their productivity. The second reason I combined these two loops is because the other is essentially “turned off” when the other is acting upon the model. In other words, nether of these loops can effect the model at the same time.

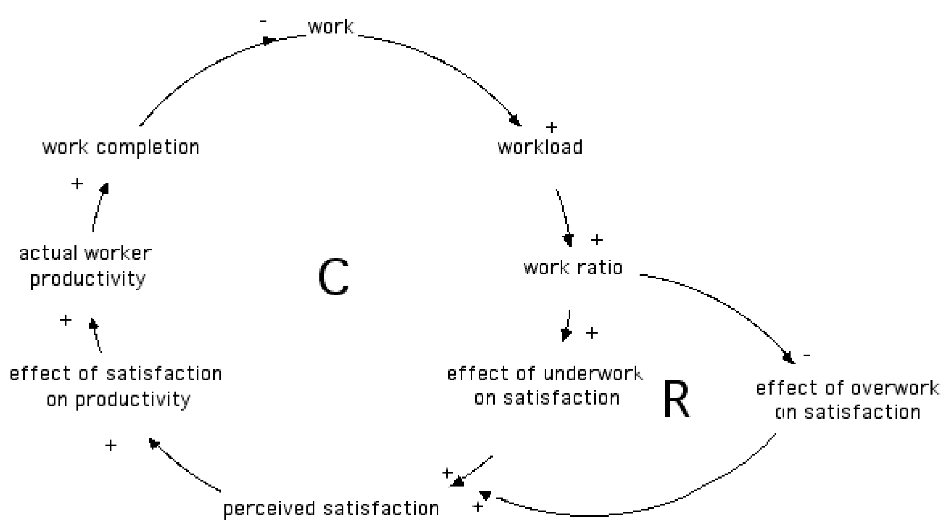

Figure 12: Workload-Productivity Loop

Since this is a combination of two different feedback loops, I am going to start with two different scenarios. The amount of work there is to be done will either increase, or decrease. If the amount of work increases, the workload will increase, and the work ratio will increase above one, meaning people are overworked. If the amount of work decreases, the workload will decrease, and the work ratio will decrease to below one, meaning people are under worked. Here is where things get a little tricky. If the work ratio is greater than one, then the effect of overwork on satisfaction will decrease, causing perceived satisfaction to decrease. If the work ratio is less than one, then the effect of under work on satisfaction decrease, causing the perceived satisfaction to decrease. This means that whenever the work ratio gets farther away from one, the satisfaction will decrease. When satisfaction decreases, the effect of satisfaction on productivity decreases, causing the actual worker productivity to decrease. When the worker productivity decreases, work completion also decreases, causing the total amount of work to increase.

It would not be entirely accurate to say that anytime the amount of work changes, work will go up. The correct way to phrase it would be: anytime the amount of work changes in a way that moves the workload farther away from the normal productivity, the amount of work will go up; while any time the amount of work changes in a way that moves the workload closer to the normal productivity, the amount of work will go down.

There are a couple of other feedback loops in this model that I did not describe in detail, although this is not to say they are unimportant. These would include the workload-satisfaction loop, the satisfaction-employee loop, the direct workload-productivity loop, the direct workload-employee loop, the wage loop and the satisfaction loop. These loops are outlined in the appendix.

Looking specifically at one feedback loop is a great way to gain a better understanding of a model, but in order to get the whole picture; you also have to understand how these different feedback loops affect the model as a whole. The model is nothing more than a number of counteracting and reinforcing feedback loops working to keep the entire system in equilibrium.

A perfect example of this is the workload-productivity loop that I described earlier. If you recall, this one loop is actually a combination of both a reinforcing, and counteracting feedback loop. When the work changes from the equilibrium amount in either direction, one of these feedback loops will kick in to try and bring the amount of work back to the equilibrium value. This is essentially what is happening throughout the entire model; except there are eleven different feedback loops all working at the same time to get the system back into equilibrium.

Model Boundaries

There are a couple of factors I was originally planning on including in my model. After some thought about these factors however, I realized that they would have caused the model to grow way beyond the scope of what I set out to do.

When I made the first sketch of my model, I had originally included a material delay in the employee sector. I had the stock, “applicants”. This had the inflow “new applications”, and two outflows: a hiring flow that fed into the employees and an application expiration time, which allowed applications to expire after a certain amount of time. I also had an outflow from employees that represented people quitting.

In order to include these factors, it would have been necessary to include a sector called “public perception”. This would be affected by the current wage, worker satisfaction, workload as well as many other factors (in fact, it could be a completely different model.) The public perception would then affect the rate of new applicants, along with rate of quitting. I envisioned a sort of cat and mouse battle to try and keep the level of employees at a stable level, when you have to worry about people leaving, and the number of available experienced pencil sharpener craftsmen running dry due to the company’s reputation as a lousy place to work.

This is also why I initially included a developed wage structure, despite the fact that it is not affected by any other parts of the model. I left it in so that if someone decides to expand upon the model, they can use that structure that had already been built.

Ultimately, I decided that including these factors would cause the model too grow to large for what I was trying to do. Therefore, I removed these factors, and am holding them constant for my model.

Model Testing

Now that the model is built, and we have examined some of the important structures within the model, we can begin to run some tests using the model. There are a variety of different situations that can be tested by changing different values within the model; however, I am going to test how an increase in demand affects the system.

Think back to that pencil sharpener company that you own. Let’s say that for the last three years there has been an extremely stable demand of 1000 pencil sharpeners per week. In week four however, a large pen company goes out of business and pen prices skyrocket. As people begin using pencils in place of pens, the demand for pencil sharpeners increases to 1500 pencil sharpeners per week. How does your company react to this sudden demand for pencil sharpeners?

Figure 13: Increase in Demand

At the time of the demand increase, there are not enough employees to complete all the work, so as one would expect, the amount of work increases sharply. Right after the work increase, the number of employees begins to increase (although not as quickly as the work), and as the number of employees increases, the speed in the increase of work slows down, eventually peaks, and then begins to fall almost as quickly as it rose. This is because as the number of employees increases, more of the work is getting done, and so the amount of work rises slower, and eventually when there are enough employees to do all of the work, it begins to decrease. This decrease in work is followed by a slow decrease in employees, until they both level out at equilibrium.

The actual worker productivity line made an interesting shape. It fell slightly, before rising a small amount, then falling just quite a bit more, before slowly growing back up to normal. The reason the productivity makes this shape can be explained by thinking of the two factors that affect it, the stress and work ratios. The stress ratio tells us whether people are more or less satisfied than normal. When the work sharply increases, workers become overworked and they become more stressed. This causes the productivity to decrease (the first dip), although begins to rise as satisfaction increases due to the increasing number of employees. The second dip is caused by the fact that once enough employees are hired, the work quickly falls, and suddenly there are a lot more employees than necessary. Once again, this causes productivity to fall (the second dip), and then slowly grow back to normal as the number of employees reaches the optimum level.

Always wanting to save money, you realize that when the demand increases you have to hire a large number of employees to do the backlog of work that had accumulated. You then speculate that perhaps hiring a small number of "extra" employees would reduce the number of employees that you would need to hire after the increase in demand.

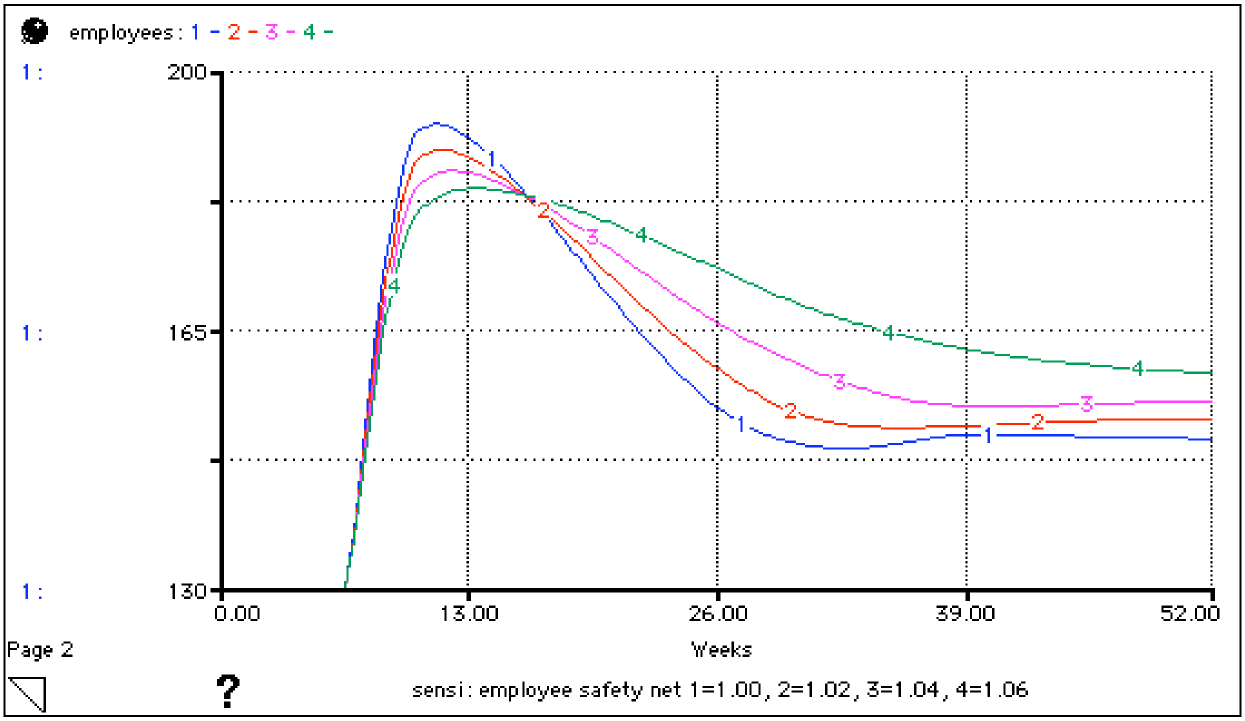

Figure 14: Employees with Increase in Demand and Employee Safety Net

By running this sensi- analysis, you can notice a couple of different things about having a safety net. One thing is that the higher your safety net, the lower your peak number of employees is going to be. Also, the higher safety net causes less oscillation, but also remains at a higher value for a longer period of time. The safety net reduces provides a cushion that absorbs small changes, causing the model to become more stable. Since you are hiring a larger percentage of employees, it also equalizes at a hire amount, and is the cause for the slower decay time. I would recommend that a company used a small safety net of around 2%. The reason I recommend this is because it seems to absorb a decent amount of oscillation without hiring too many more employees than necessary.

Key Learning

Now that the model is completed, I can list a couple of things that I learned throughout the process.

- The number of employees does not increase in a convergent way as one might expect with an increase in work; rather it increases and decreases in a complex way through a variety of different feedback loops.
- It does pay off to have a very small employee safety cushion; however, if it is too big it can be very costly.
- We also have a better understanding about how the level of employees in a company is affected by different factors in the workplace.

The model is fairly open ended as well. For example, one could plug this into a supply and demand model where the employee satisfaction, and employee per customer ratio affected the demand. This could introduce some pretty interesting feedback. I would also love to see somebody add company perception, and applicants that I talked about in the model boundaries section. Then you could see how changes in company perception affected the satisfaction of the normal employee.

I also learned a good amount about the process of building models in making this model. Some of the things I learned are:

- Do not try to focus on everything at once. It took me about three weeks before I actually had any stocks or converters in my model. I simply did not know where to start. I could imagine all of the feedback that I wanted to include in my head, but I didn't know where to start the model. Eventually, I decided that I needed to simply start. I threw down some stocks, and began to connect small things together. You can't get to concerned with the model as a whole, because that is simply too much information to handle at once.
- It is also easy to become overly concerned with small details in the model, without looking at the model as a whole. It is easy to become very obsessed with the exact shape of a graphical converter, or be overly concerned with adding fairly insignificant details. Sometimes these small details are unimportant, and you need to step back; making sure you are still heading in the direction of answering the original question posed.
- I also learned how a model is really a large number of delays and feedback all working together to bring things back into equilibrium. By studying these delays and the feedback, you can really begin to understand not only how things work within your model, but also how almost everything in the world interrelates through countless numbers of feedback loops.

So how does a company's workforce respond to changes in demand? Well, it depends on the circumstances. I tested a couple of different circumstances previously in the model, but the possibilities are limitless. In order to answer that question, you must plug your circumstances into the model, and it will do the rest.

Appendix

Equations:
Omitted

The Final Model

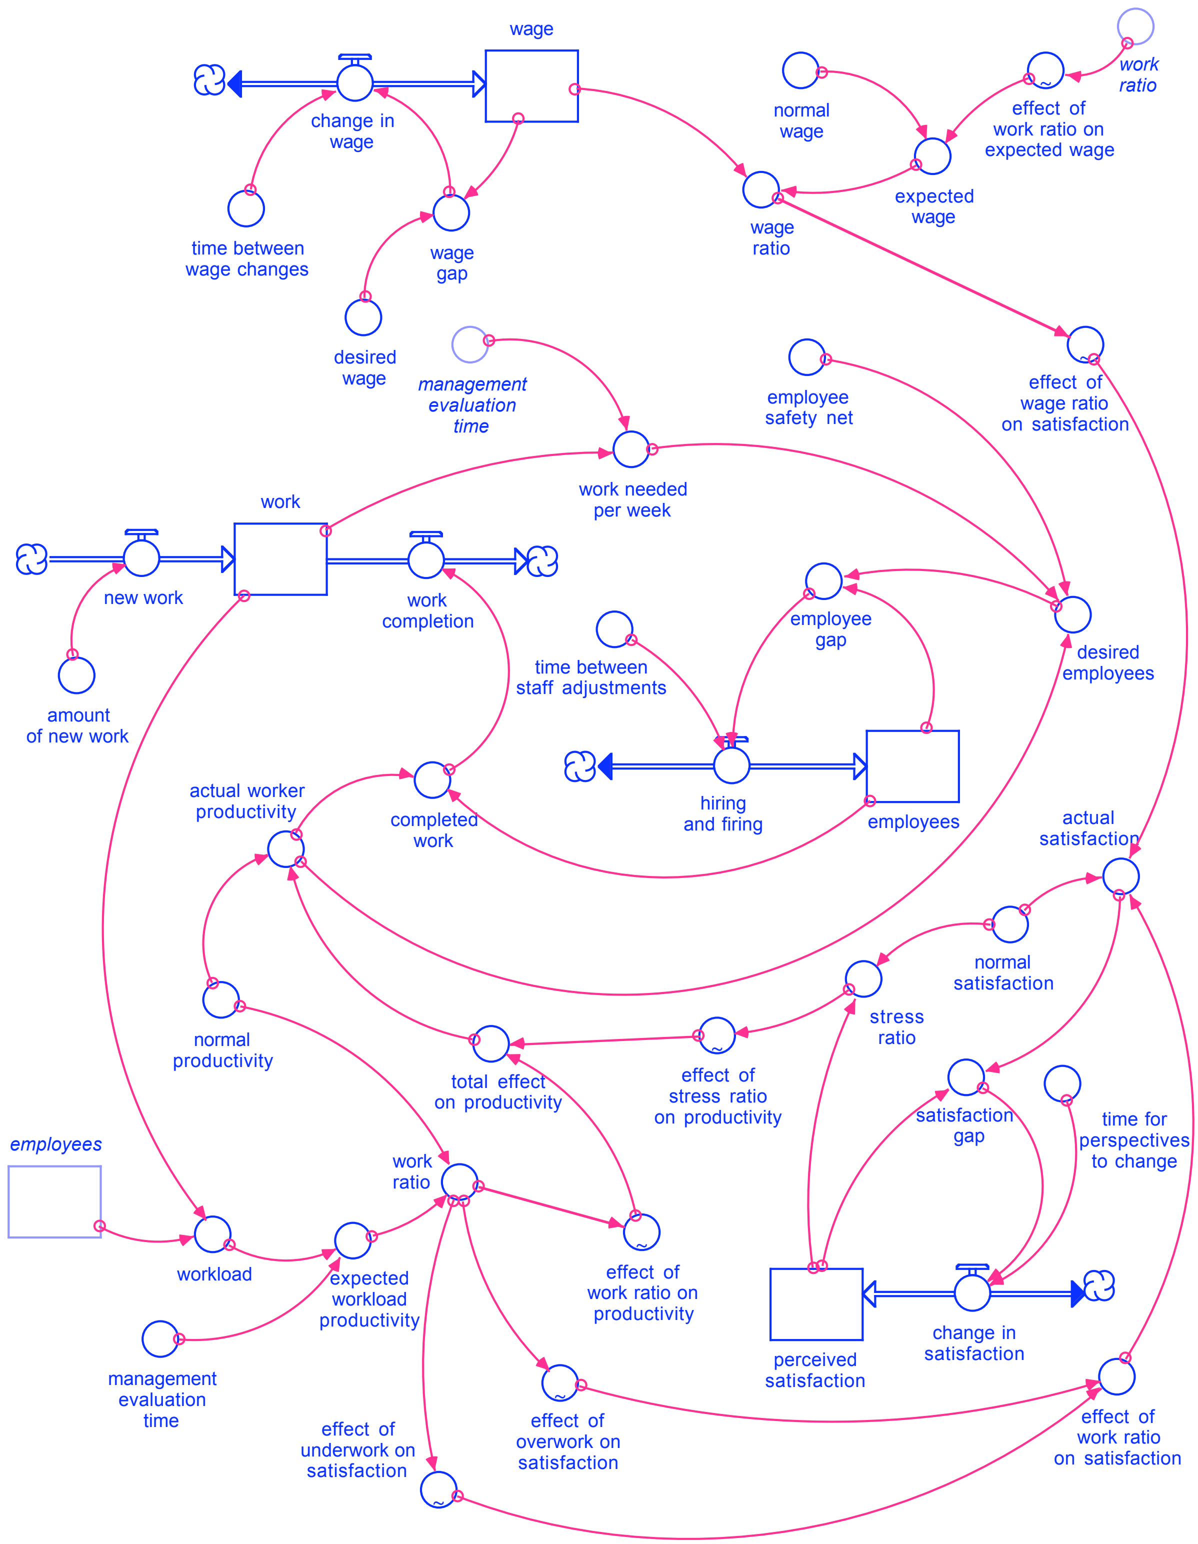
wage
change in wage
time between wage changes
wage gap
desired wage
normal wage
work ratio
effect of work ratio on expected wage
expected wage
wage ratio
management evaluation time
employee safety net
effect of wage ratio on satisfaction
work
work needed per week
new work
amount of new work
work completion
time between staff adjustments
employee gap
desired employees
actual worker productivity
completed work
hiring and firing
employees
actual satisfaction
normal satisfaction
stress ratio
normal productivity
total effect on productivity
effect of stress ratio on productivity
satisfaction gap
time for perspectives to change
employees
work ratio
workload
expected workload productivity
effect of work ratio on productivity
perceived satisfaction
change in satisfaction
management evaluation time
effect of underwork on satisfaction
effect of overwork on satisfaction
effect of work ratio on satisfaction

Feedback Loops

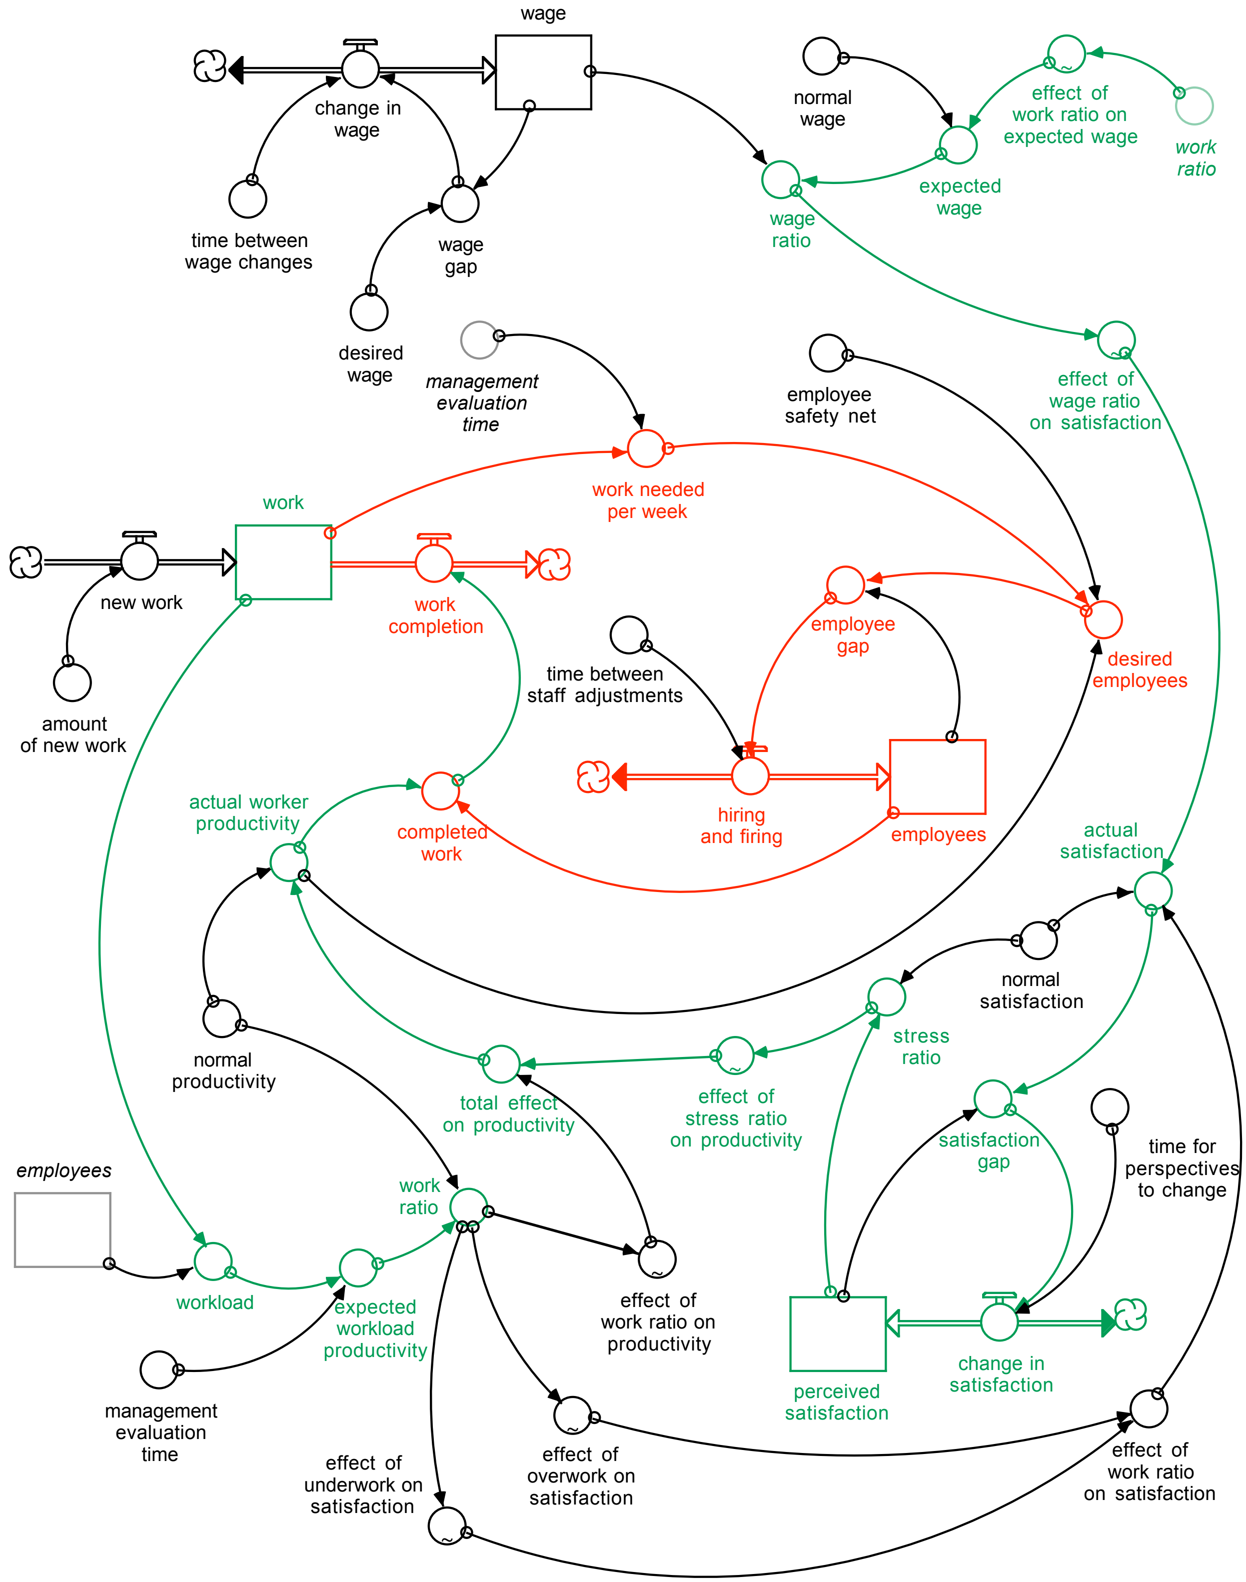

Green = Work-Satisfaction Loop

Red = Work-Employee Loop

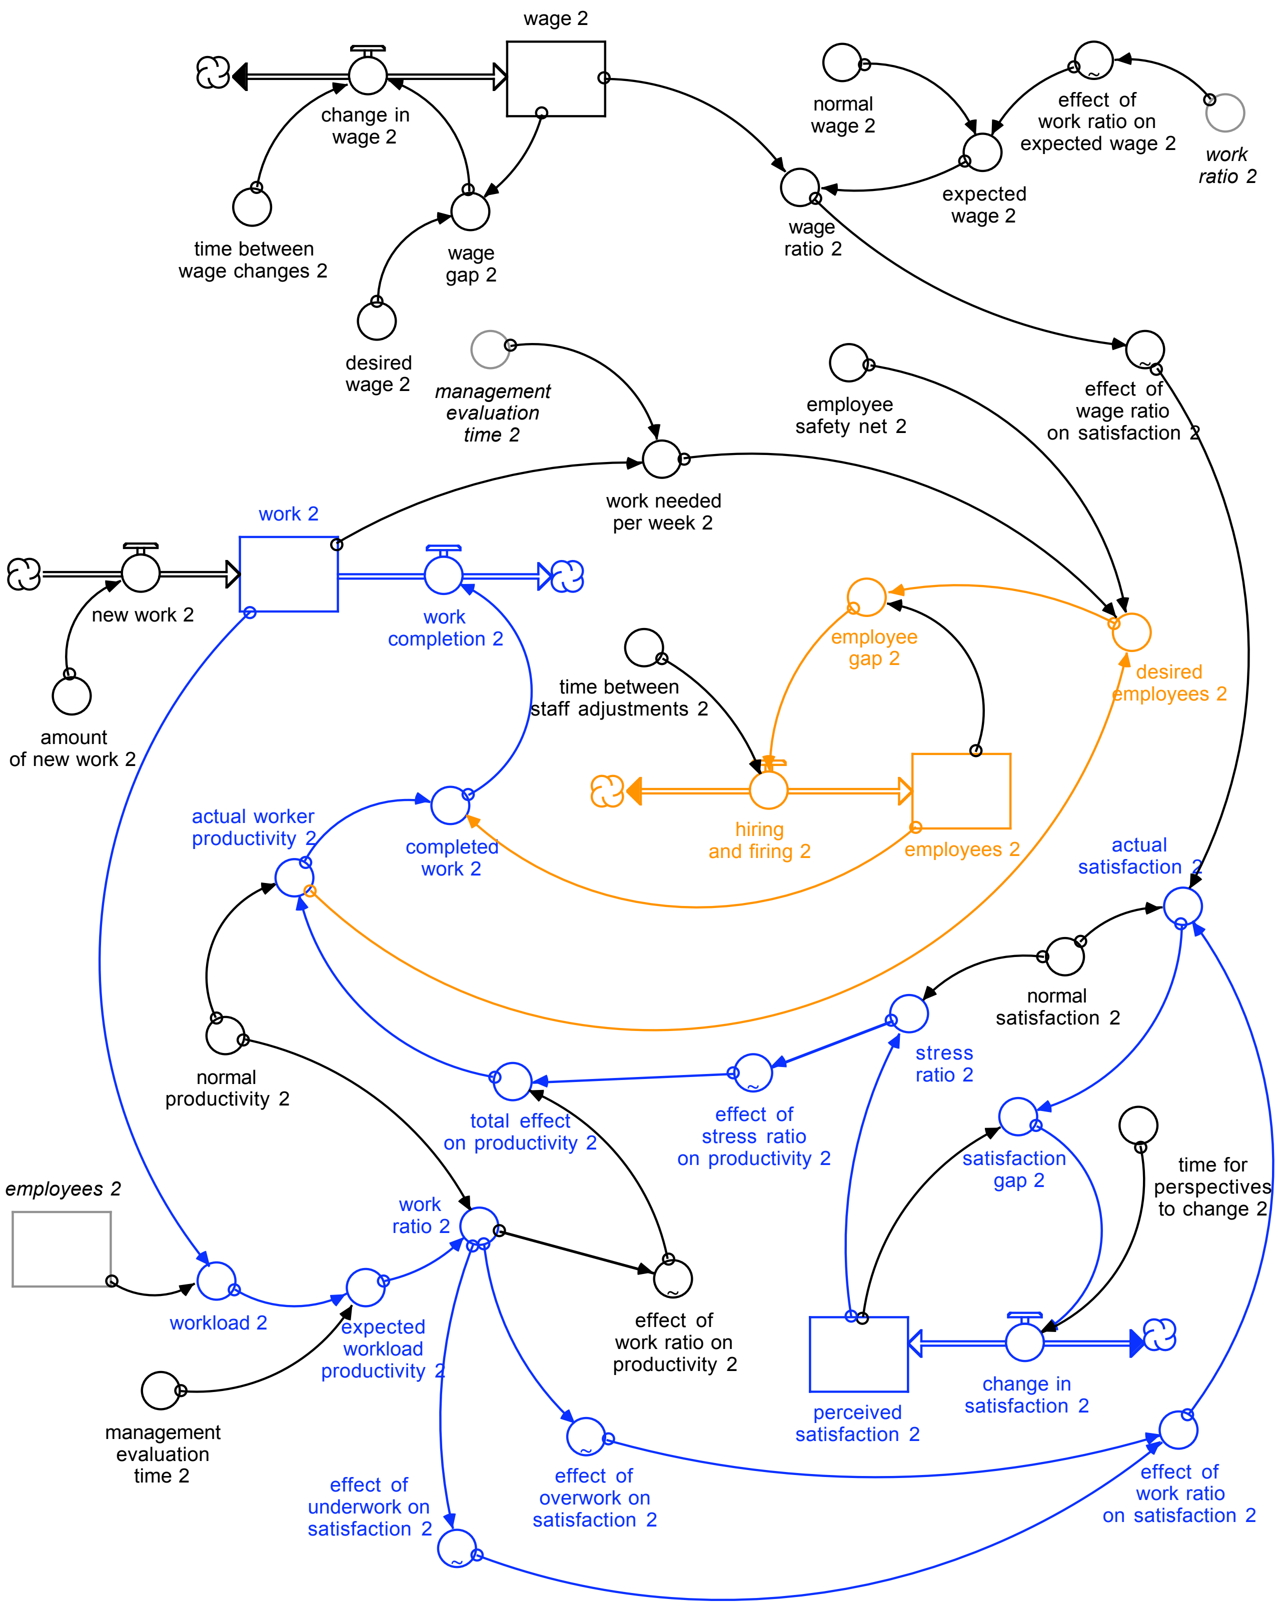

Blue = Workload-Satisfaction Loop

Blue to Orange = Satisfaction-Employee Loop

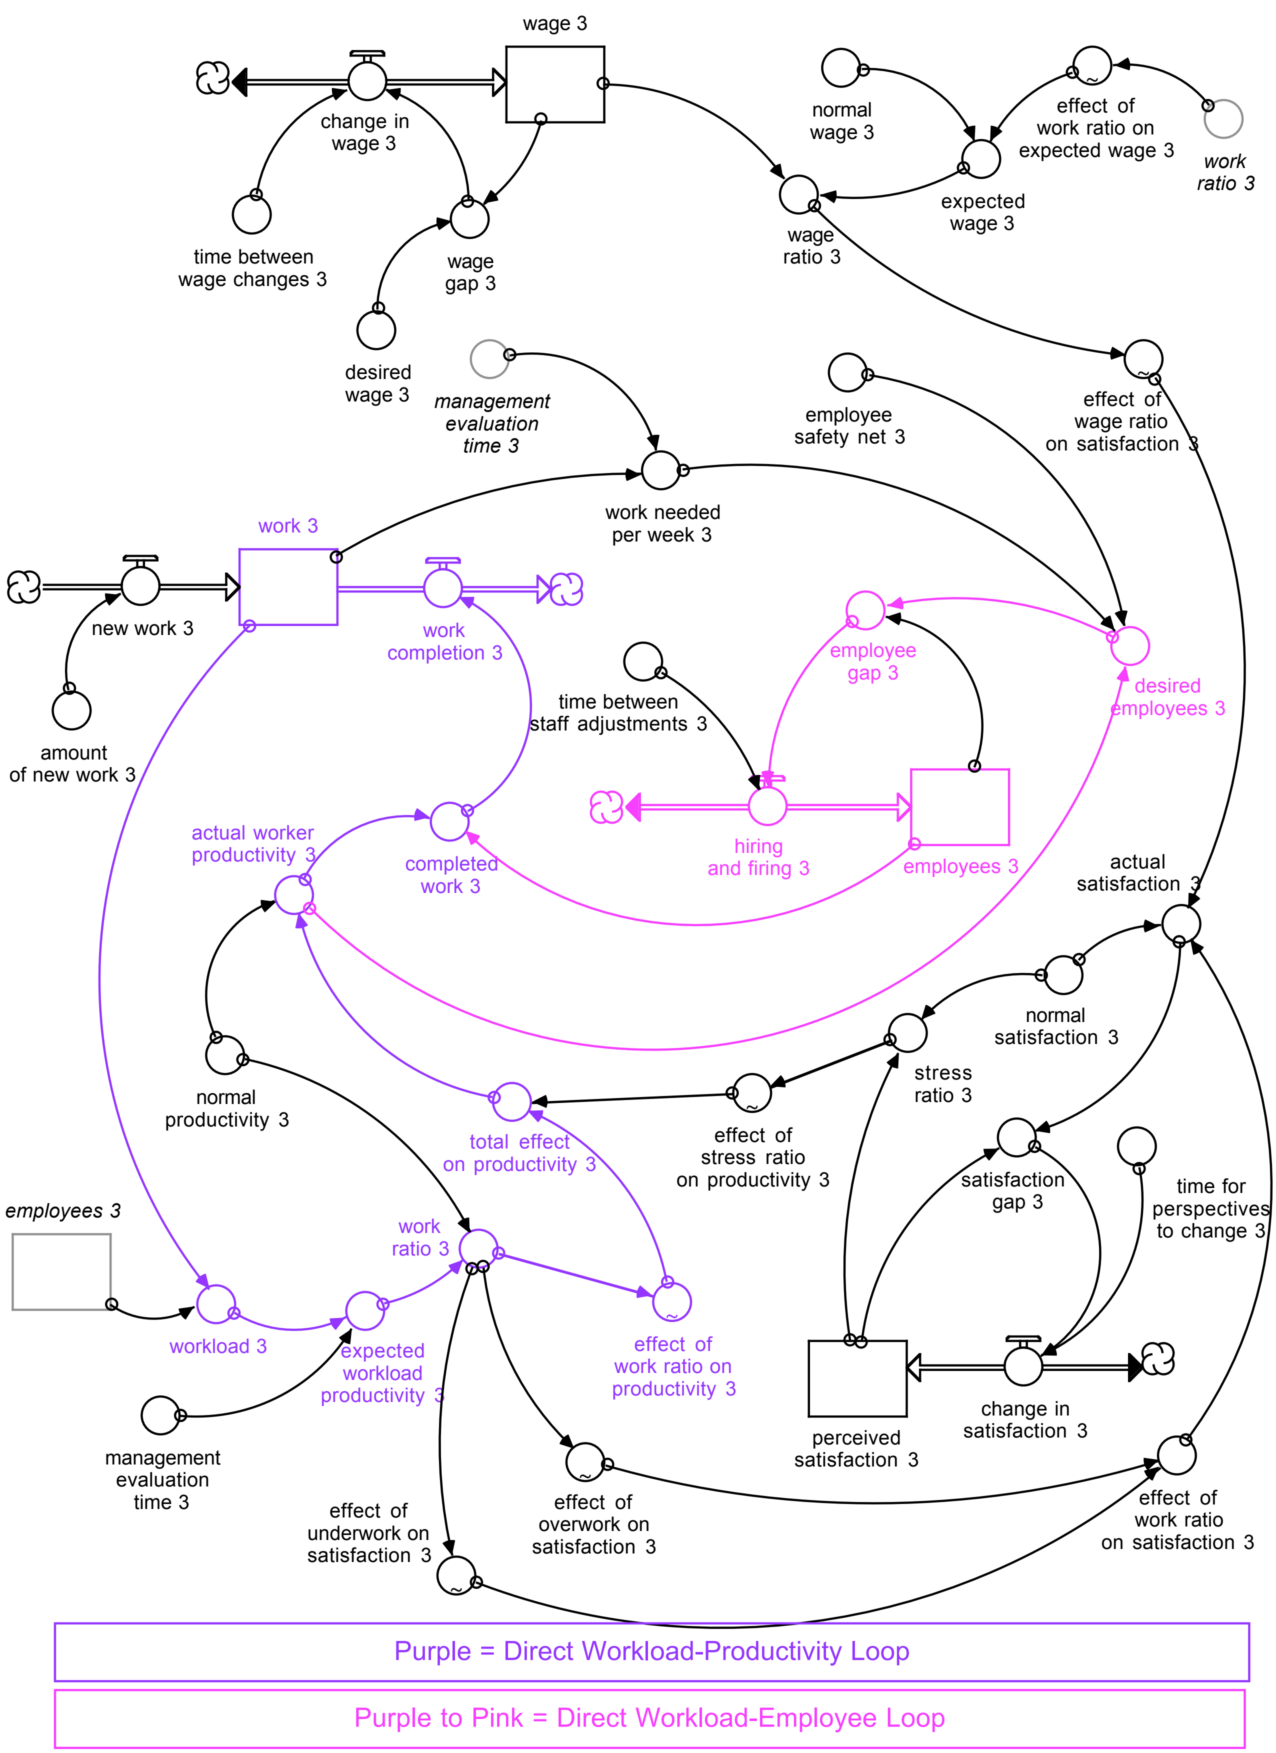
wage 3
change in wage 3
time between wage changes 3
wage gap 3
desired wage 3
normal wage 3
effect of work ratio on expected wage 3
work ratio 3
expected wage 3
wage ratio 3
effect of wage ratio on satisfaction 3
management evaluation time 3
work needed per week 3
employee safety net 3
work 3
new work 3
amount of new work 3
work completion 3
employee gap 3
desired employees 3
time between staff adjustments 3
hiring and firing 3
employees 3
actual worker productivity 3
completed work 3
actual satisfaction 3
normal satisfaction 3
stress ratio 3
normal productivity 3
total effect on productivity 3
effect of stress ratio on productivity 3
satisfaction gap 3
time for perspectives to change 3
employees 3
work ratio 3
workload 3
expected workload productivity 3
effect of work ratio on productivity 3
perceived satisfaction 3
change in satisfaction 3
management evaluation time 3
effect of underwork on satisfaction 3
effect of overwork on satisfaction 3
effect of work ratio on satisfaction 3
Purple = Direct Workload-Productivity Loop
Purple to Pink = Direct Workload-Employee Loop

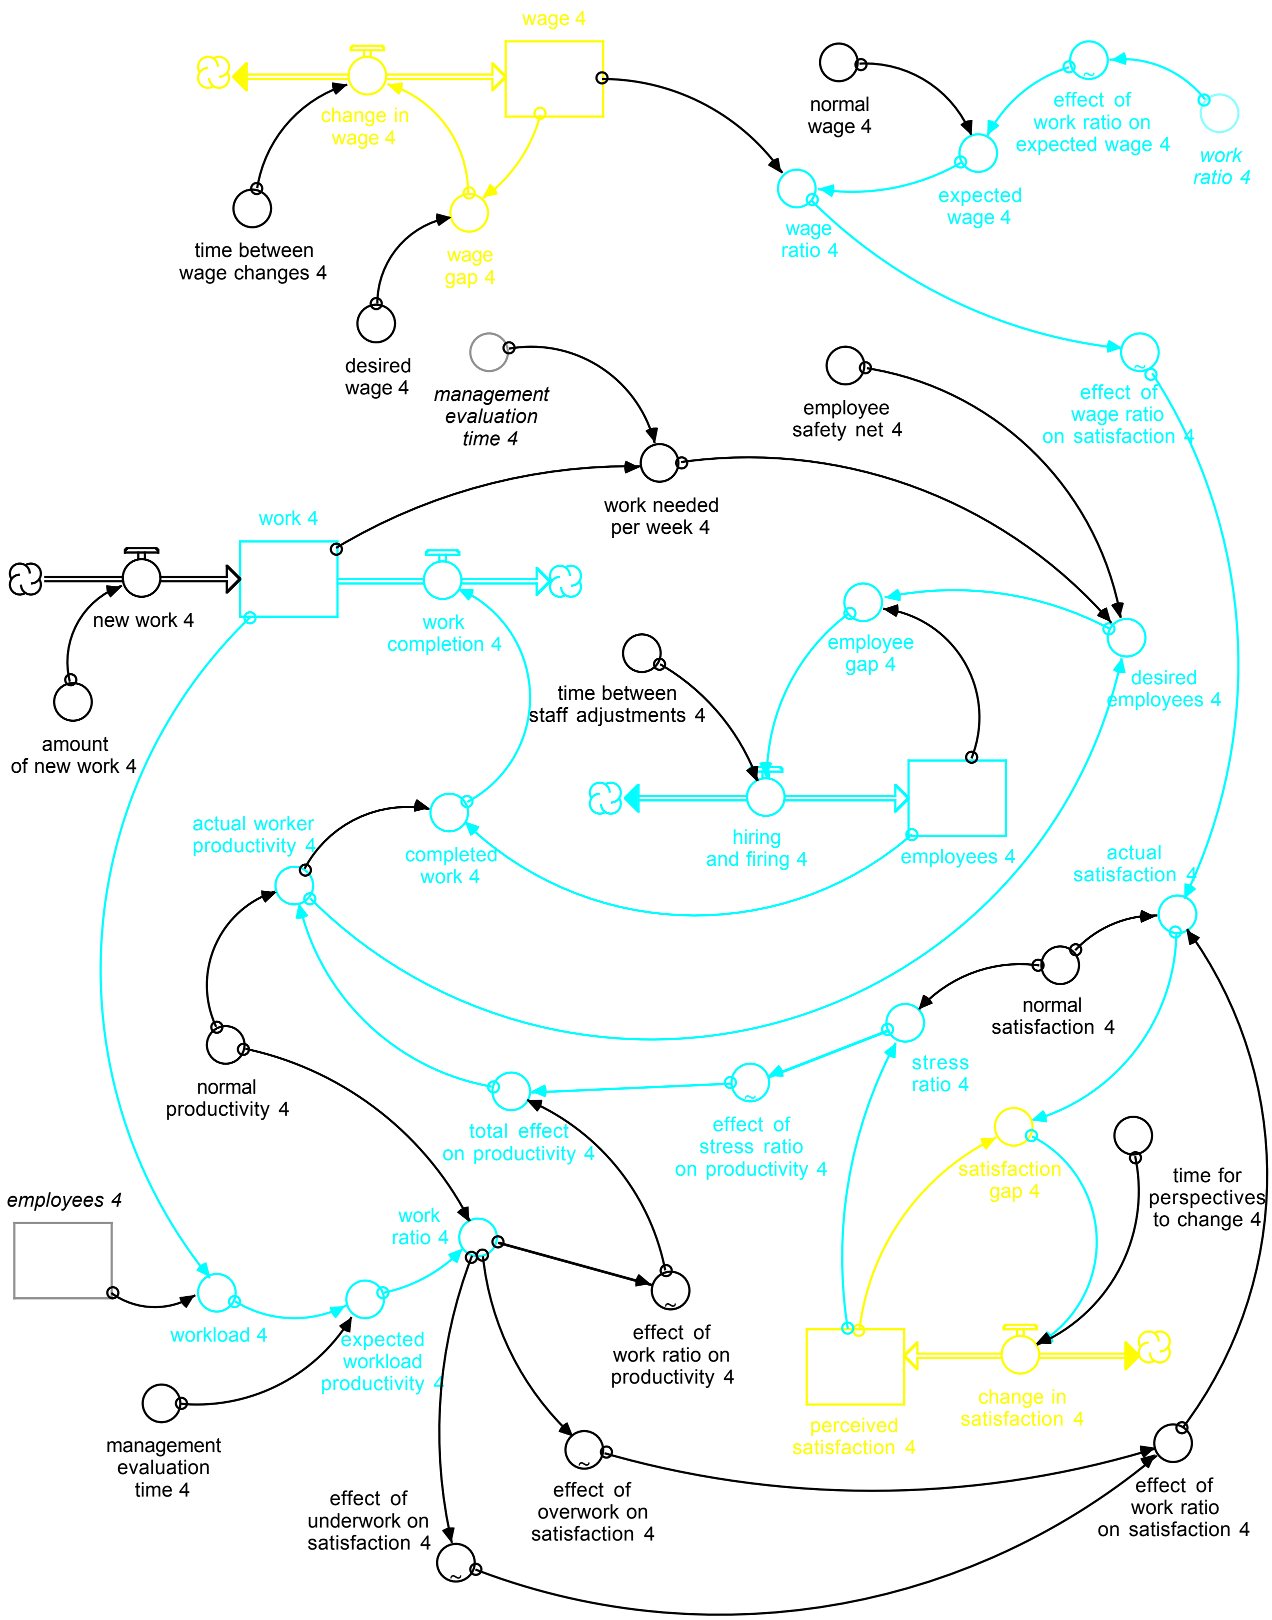

Light Blue = Wage-Employee Loop

Yellow = Other Loops