unwieldy, based on the findings in the report and the author's own extensive experience, he suggests that it might be more expensive to manage this way. However, it is a workable process and the decision making is on firmer ground in terms of getting broader representation.

Another question posed by the author is: Can we learn from 2006? It seems as though we are having a lot of these meetings — there have been at least four Fraser sockeye reviews over the last 15 years, the last one being in 2005 and prior to that a review in 2004 when we lost a 1.3 million fish. The overriding question is: Can we learn from our mistakes?

"The overriding question is: Can we learn from our mistakes?"

Can we do pre-season forecasting effectively? If not, what can we do instead?

Randall M. Peterman, Professor and Canada Research Chair in Fisheries Risk Assessment and Management, School of Resource and Environmental Management, Simon Fraser University

Can we do pre-season forecasting of salmon abundance effectively and if not, what should we do? To answer these questions you need to understand the challenge of coming up with good pre-season forecasts for abundance of salmon. To do this you need to know why good forecasts are needed, how forecasting is done in western North America (not just British Columbia) and how forecasting methods compare. This presentation explored these questions and purposely addressed mostly non-British Columbia examples because Dr. Peterman wanted to focus on the big picture and not get bogged down in details of particular fisheries.

He started by making sure that we were all on the same wavelength with respect to our thinking about what pre-season forecasts are. In reviewing the basic life history of salmon, he noted that spawners produce juveniles that migrate out to the ocean, and depending on the species, they spend anywhere from 16 months to several years in the ocean before heading back to freshwater. The key uncertainty here is the survival rate of the juveniles before they become adults (Figure 1).

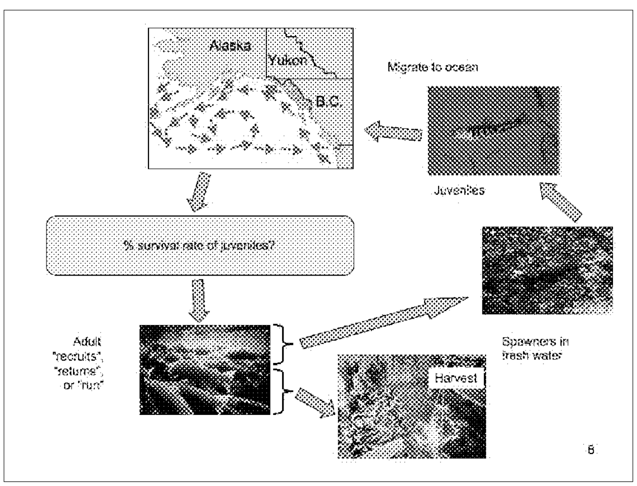

Figure 1: Survival rate of juvenile salmon.

Adult abundances that survive in the ocean may be referred to as adult recruits, returns, or the run size — these three terms are synonymous. These are the fish that are then partitioned either into harvest or those that migrate back toward their spawning areas. Therefore, pre-season forecasting refers to all the adults prior to onset of the fishing. It does not include the fishing — it is all those that are available to be harvested or left to spawn. (Figure 2).

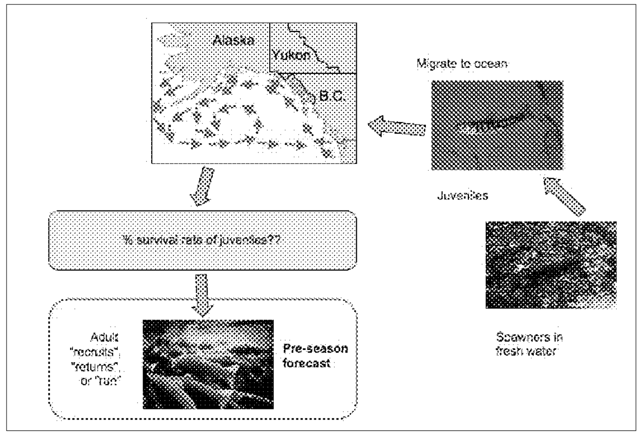

Figure 2: Pre-season forecast.

"The problem is that we have highly uncertain survival rates for the juveniles ..."

The problem is that we have highly uncertain survival rates for the juveniles, which is why there are question marks inserted in the figure. It is quite variable and quite uncertain from one year to the next and, given that uncertainty, we have the difficult job of trying to come up with the pre-season forecast.

Figure 3 describes one example of a case in which there are very good data on variation over time in the survival rate of juveniles in the ocean. Clearly, there is tremendous variation from one year to the next. If you look at the increase in survival rate from 1967 to 1968, in one year it increased from about 5% to over 20%. Subsequently, in the next year, it decreased from over 20% to about 8%. That is a tremendous range — a 4-fold change in abundance of adults from 5% to 20%, given that all else is equal. For a given juvenile abundance this is equal to a 4-fold change, meaning that you would have 4 million sockeye coming back, instead of 1 million.

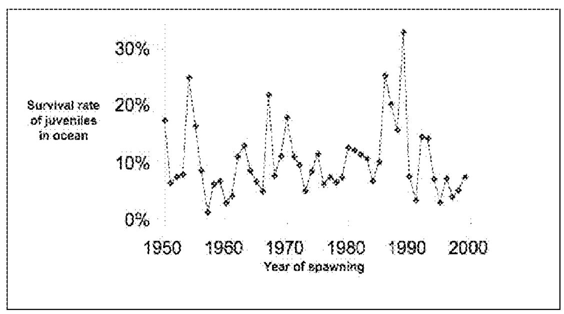

Figure 3: Survival rate in the ocean for sockeye salmon from Chilco Lake, BC, 1950–2000.

Why are good forecasts needed?

Basically there are two reasons to have good forecasts. The first is to help those who catch fish to plan ahead — this refers to everyone who is catching fish, including First Nations, commercial, recreational sectors and all the other people who are interested in fish. This can be illustrated with a specific commercial fishery example, the Bristol Bay sockeye fishery in Alaska, the largest sockeye salmon fishery in the world (Figures 4). On the 'x' axis is the year the fish returned and on the 'y' axis is the number of adult salmon returning, in millions.

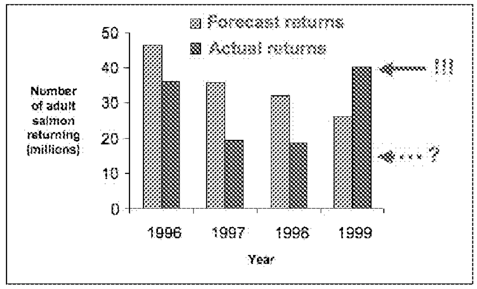

Figure 4: Forecast returns and actual returns for a number of sockeye salmon returning in Bristol Bay, Alaska from 1996–1999.

You can see that from 1996 to 1998 the forecast by the Alaska Department of Fishery and Game turned out to be larger than the actual returns, by a substantial amount. The reason this is so important from a commercial fishery point of view, is that Bristol Bay, Alaska is a very remote fishing area and therefore the industry needs to have a good forecast ahead of time in order to ship millions of cans to Alaska, hire workers for the processing plants, and get enough tenders lined up in order to service the boats. In the three years from 1996 –1998 the commercial industry basically 'ate' a lot of costs related to leftover labour, cans, etc. Then when the Alaska Department of Fishery and Game (ADFG) put out its forecast in 1999, the industry said, "these guys are using a forecasting method that is biased on the high side so we are not going to ship nearly as many cans, nor as many people" and guess what happened? In this case, they had many more fish than even the ADFG had predicted — over 40 million, instead of 26 million. We can easily guess who was caught short. Industry didn't have nearly enough infrastructure and labour in place to process the fish and that was a serious problem.

The second reason why we need good pre-season forecasts is to help managers deal with three issues. One is to determine the appropriate early-season regulations — these can always be changed later but for the first few weeks they rely heavily on the pre-season forecasts. The

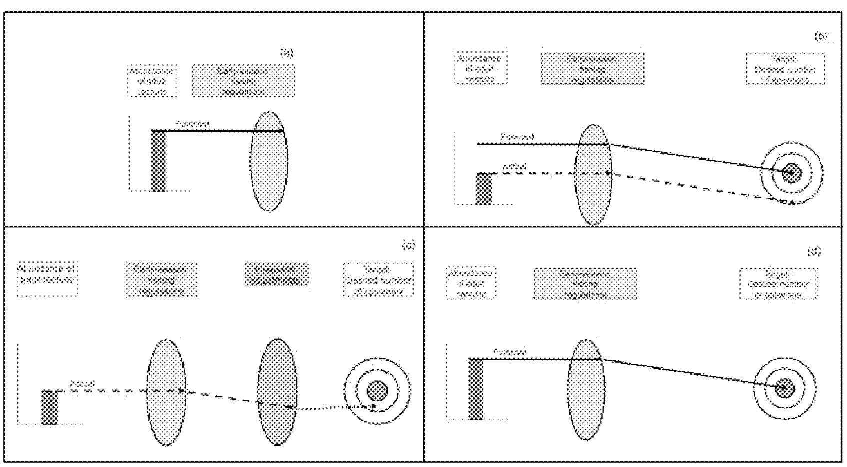

Figure 5: Relationship between the projected abundance and target abundance of spawners — shifting the trajectory to move towards the target.

managers also rely on good forecasts to help meet the target number of spawners or the target harvest rate, depending on which they are aiming for. In addition, of course, they wish to reduce the frequency of conservation problems. Pre-season forecasts that are of high quality can help to meet these needs.

You can think about how to manage problems these using a conceptual diagram such as the series in Figure 5. An abundance of adult recruits is forecast at some level and early season fishing regulations are set (a). One can think of this as a lens that changes the projection or trajectory of the abundance so that it meets the target abundance of spawners (b).

If the forecast is wrong and the actual numbers come in much less than forecast, and then if those same early season fishing regulations are applied, the target won't be hit and it will come in too low. Thus, managers use extensive in-season adjustments to correct for errors in the pre-season forecast (c). Again, this is shown as a lens that shifts the trajectory to move towards the target (d).

"... the history of data for any population of salmon on the Pacific west coast shows that we rarely hit the target and there is a frequency distribution of actual levels of spawners around the target."

However, the history of data for any population of salmon on the Pacific west coast shows that we rarely hit the target and there is a frequency distribution of actual levels of spawners around the target.

The data in Figure 6 represent a randomly chosen example from the Nass River sockeye for a period of 28 years. On the 'x' axis is the ratio of the actual number of spawners at the end of the season to the target number of spawners. When that ratio is 1, the target has been met. Notably, there is tremendous variation from one year to the next in how close we are to meeting the target; there is one year in which there were three times the number of spawners as desired and two years when it was down to about 20% of the target. This is the typical type of frequency distribution for most fish populations and this is a reality along the entire Pacific west coast.

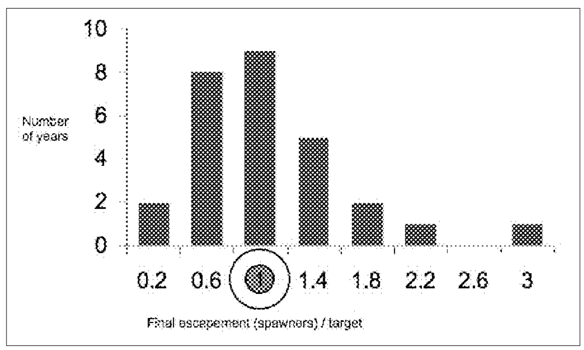

Figure 6: *Final escapement to target ratio for Nass River sockeye over 28 years.*

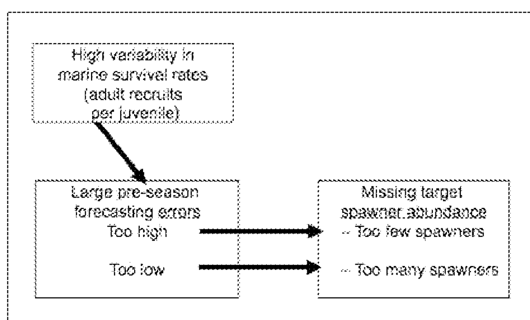

Figure 7: *Summary diagram showing that high variability leads to pre-season forecasting errors, which lead to missing the target.*

"... high variability in marine survival rates and the contributing variability from the juvenile abundances lead to large pre-season forecasting errors that in turn lead to usually missing the target spawner abundance."

In summary, high variability in marine survival rates and the contributing variability from the juvenile abundances lead to large pre-season forecasting errors that in turn lead to usually missing the target spawner abundance (Figure 7).

If the forecasts are too high, then that typically means that there are too few spawners, because by the time it is realized that there are many fewer fish than expected, it is too late to allow enough fish past the fishery to meet the target. Conversely, when the forecasts are too low, by the time there is an in-season update that says there are a lot more fish than expected, the harvesting intensity will not take the spawner abundance down to meet the target and so we end up with more spawners than the target.

Table 3: Pre-season forecasting models

1.	Simple (e.g., same returns as last year, ...)
2.	Project forward the recent time trend of returns
3.	Abundance of parental spawners
4.	Abundance of juveniles (very few cases)
5.	Sibling age-class (sockeye and chum)
6.	Between-species correlation (pink-chum)
7.	Multi-stock models based on populations with correlated survival rates
8–11.	Some of above with environmental indices added (e.g., ocean temperature)

How is forecasting done in western North America?

There is a variety of methods (models) that are used for pre-season forecasting. Table 3 describes 11 types of models.

The first is a very simple model in which the forecast is equal to last year's returns (a basic assumption), or what you had two years ago, or four years ago, or it is the average over the last four years. Proceding through the list of models, it is clear that the forecasting methods become more complicated. The second approach is to simply project forward the recent time trends of returns. The third is to use the abundance of spawners that gave rise to the adults that are returning this year. In a few cases where there are available data on juvenile abundance, they can be used to forecast the adult abundance. The next type of model is a sibling age-class model — it is commonly used for sockeye and chum salmon. This is where you use the forecast of age 4-year-olds this year based on the abundance of age 3-year-olds that came back last year. They are siblings and they shared every part of their life up to the last year. Then there are between-species correlations where you can use, in some cases, the correlation between pink survival rate and chum survival rate for pinks and chum cohorts that went to sea in the same year. In addition, there are some more multi-stock complex models that are based on populations with correlated survival rates. Models 1–11 are based on some of my group's work on trying to improve forecasting. The next four models (8–11) are similar to the above models but include environmental variations such as ocean temperature, to reflect ocean conditions, added as a predictor.

How do these forecasting methods compare?

With funding from an NSERC research grant, the Peterman research group is investigating the question: How do these forecasting methods compare? They used a database that they have compiled over approximately a decade of 120 Pacific salmon stocks along the west coast. The average duration of data sets was about 31 years and they included three species (see Figure 8).

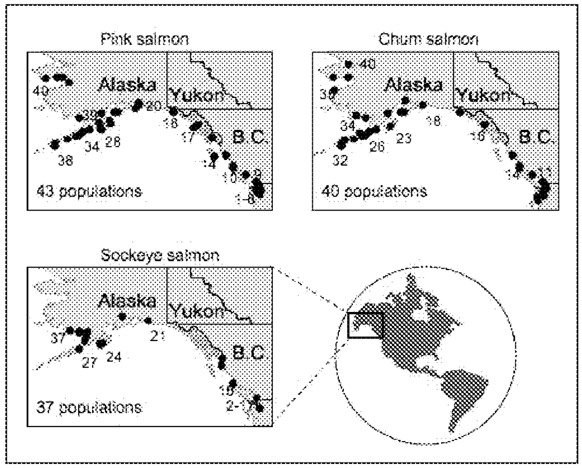

Figure 8: *Database of 120 Pacific salmon stocks over an average of 31 years.*

What if a given model had been used in the past?

- Used data through year *t (e.g., 1960)* to estimate parameters of forecasting model
- Forecast abundance of adult recruits in *t*+1 *(1961)*
- Compared forecast to actual recruits in *t*+1 *(1961)*
- Added actual recruits from *t*+1 *(1961)* to data set

Repeated over all years

Figure 9: *Retrospective evaluation of each pre-season forecasting model.*

For pink salmon and all the other species, each black dot represents the point of ocean entry for the seaward migrating juveniles. Included in the database are 43 populations of pink, 40 populations of chum and 37 populations of sockeye salmon. We conducted a retrospective evaluation of each forecasting model (see Figure 9). In essence they asked: What if a given model had been used in the past? How would it have performed?

To answer these questions, we started with a given block of data, for example, the first 10 years of data up through 1960, and we estimated parameters of a particular forecasting model. We then forecast the abundance of adult recruits in the next year — in this example, 1961 — and then compared that forecast with the actual returns in that year (1961), calculated the error in the forecast, and then took that data observed point and added it to the data set. They repeat this process over all years. This iterative process adds one new data point each year and re-estimates the forecast for the next year to move forward.

Figure 10 (a-e) illustrates how this works graphically. These data were used to estimate the parameters through to 1986 in order to forecast the abundance for 1987 (a). When using a sibling model, you can see that the red data point is much higher than the actual returns, resulting in a very large positive value for forecasting error in 1987 (b). A 1987 data point was then added to the data set, and we re-estimated the parameters for

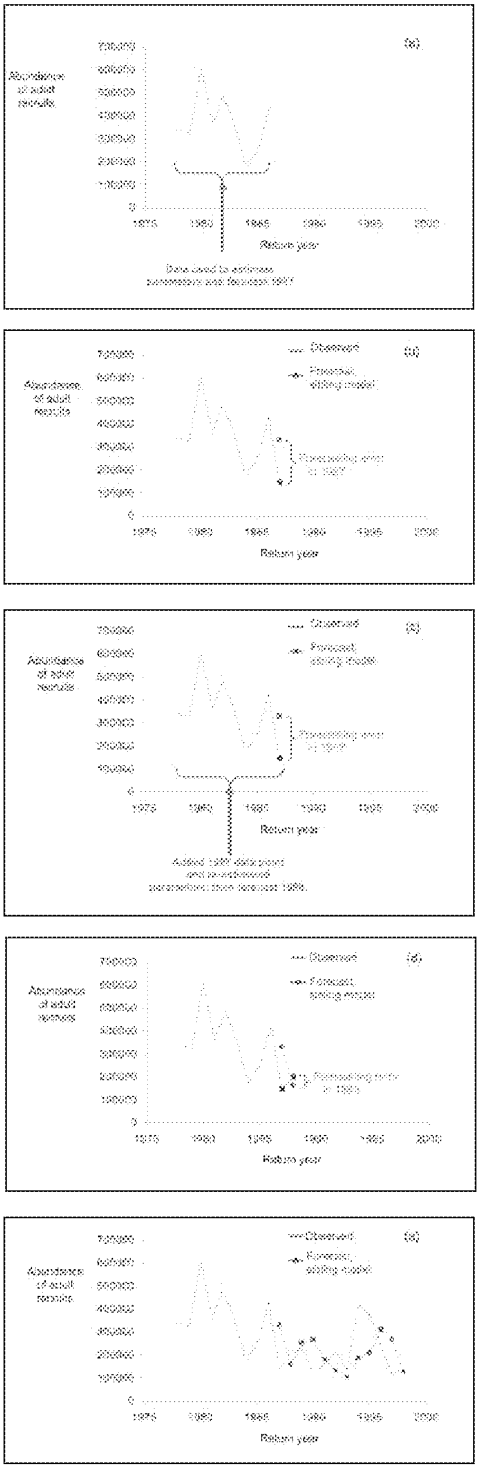

***Figure 10 (i–v):** Results of application of an observed abundance of adult recruits for chum salmon of Andreafsky River, Alaska from 1975–2000.*

that model and forecast 1988 (c). In this case, the forecast was too low — a small enough forecasting error but different in sign than the previous year (d). This process was repeated through the series and this resulted in the differences between the forecast values of abundance in red and the observed values in blue (e). Therefore, by accumulating all those differences in forecasts and actual returns over time, we obtained a frequency distribution from which we could calculate a series of statistics summarizing the performance of the fore-casting methods, calculating these forecasting errors for all forecasting methods, one at a time for each salmon population of the 120, and for each historical year.

Table 4 summarizes the results of this work. There were 37 populations for which we did these calculations. We used 11 pre-season forecasting models, and had over 1,000 cases of annual forecasts and comparisons with actual recruitment. For chum salmon, also had those same 11 forecasting models and about 600 cases for calculating forecasting errors. For pink salmon, only eight models were used — three of the models did not apply to them. We then ranked the forecasting models based on several ranking statistics, one of which is mean error, that is, the bias in the forecast, as well as whether or not

Table 4: Restrospective analyses conducted

Species	Populations (stocks)	Pre-season forecasting models	Cases (stock-years) calculated forecasting errors for each model
Sockeye	37	11	1,081
Pink	43	8	783
Chum	40	11	665

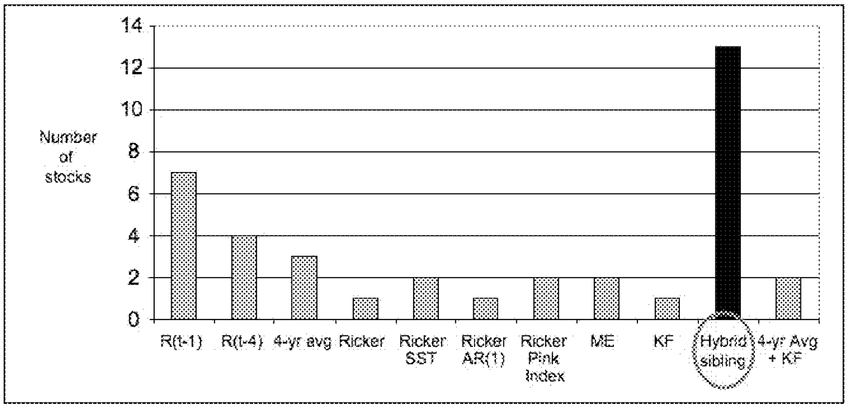

***Figure 11:** Ranking of models based on bias and magnitudes of errors for 37 sockeye populations located between Washington and Alaska.*

they tended to be too low or too high. Among other measures, we also calculated the mean magnitude and direction of error.

The results for these retrospective analyses were three-fold. First, no single model of a species performed well across all stocks and all ranking statistics. Figure 11 provides an example for 37 sockeye populations along the west coast, from Washington through to Alaska. The hybrid sibling model was the best model for about half the populations. But note that every model listed on the 'x' axis also had at least one stock for which it

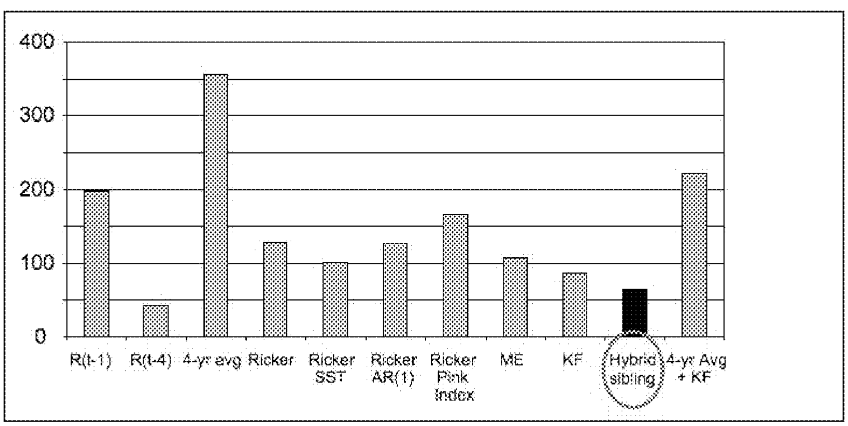

***Figure 12:** Average magnitude of error in forecasts across 37 sockeye stocks.*

was the best model. Thus, you cannot justify applying a single forecasting model to all stocks. For the average magnitude of error, Figure 12 shows the mean percentage error of forecasts across all 37 stocks. The sibling model, is the second best one in terms of this ranking statistic. However, there is quite a bit of variation. Notably, even though the hybrid sibling model is the second best, its average forecasting error is over 60%.

The second result of this work is that typically the best model for each stock only explained a small portion of variation in annual recruitment overtime. Large uncertainty in abundance remaining unexplained. It is also important to note that the quality of forecasts in British Columbia is just about the same as elsewhere.

Figure 13 shows results for using the best forecasting model for each stock of pink and sockeye salmon rather than applying one forecasting model across all populations of a species. For a given species there are results for each stock, from 1 to 43, where 1 is in Washington and 43 is in western Alaska. The dots show the proportion of variation in adult recruitment explained by the best forecasting model. Note that the Alaskan pink salmon stocks did have slightly higher values than BC or Washington

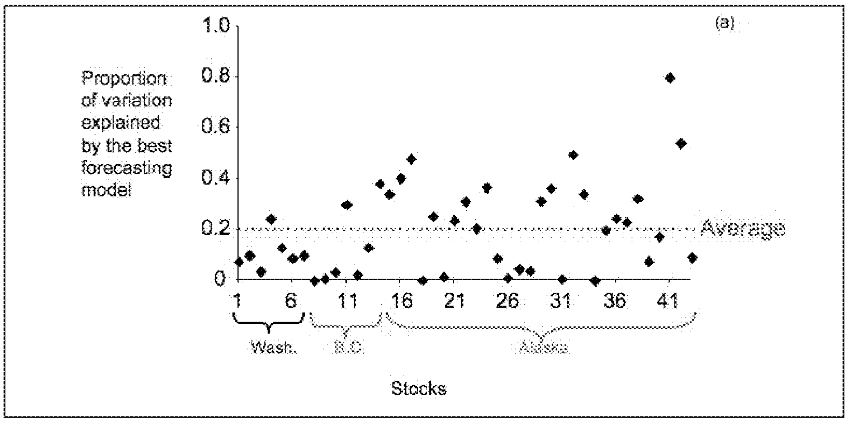

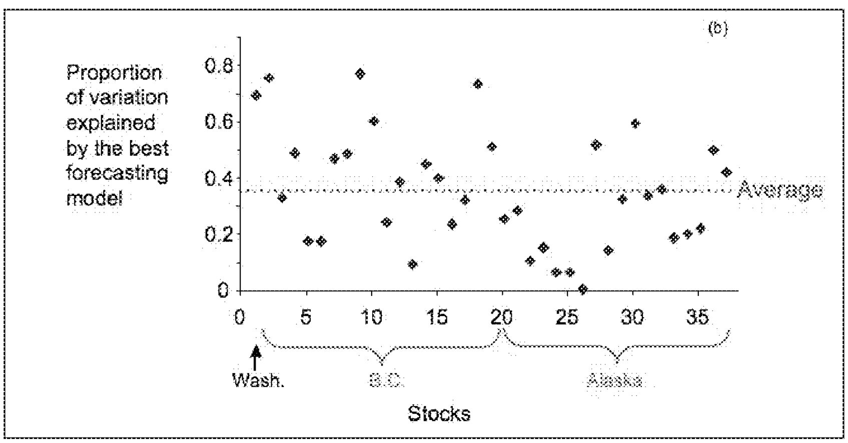

Figures 13 a&b: *Proportion of variation in adult recruitment explained by the best forecasting model for pink (a) and sockeye (b) salmon stocks, from Washington to Alaska.*

stocks but the average proportion of variation accounted for is only 20% (Figure 13a). For sockeye salmon it is slightly better. The average is 36% and in this case the BC forecasts tend to be better than those in Alaska (Figure 13b).

Taking the average values represented by the dashed lines and showing them in a different way, the data in Figure 14 show the average proportion of variation explained by the best forecasting model for each stock. The result for pink salmon is 20%, chum salmon, 21% and for sockeye salmon, 36%. The key point here is that there is a lot of unexplainable variation. Pre-season forecasts generally are not good anywhere on the west coast.

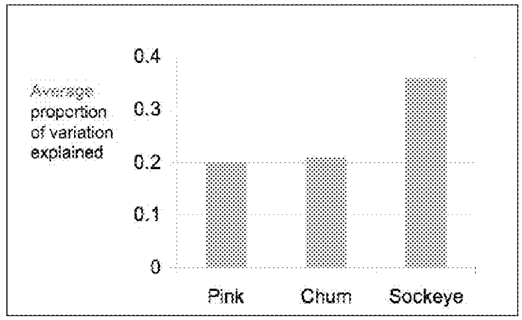

Figure 14: *Average proportion of variation explained using the best pre-season forecasting model for pink, chum and salmon stocks.*

"Pre-season forecasts generally are not good anywhere on the west coast."

This brings us to some important implications for everybody interested in salmon, including scientists, managers, and harvesters. The question is: What should we do with the types of results described above? Table 5 describes six suggested answers.

The first suggestion is to improve in-season monitoring, updating forecasts in-season, and better link in-season decisions with those updated forecasts, taking uncertainties into account. These are well-known needs. Earlier, Craig Orr referred to the in-season updates that need to be done. There are advances made at least yearly, and there are new techniques based on genetic information to allow us to separate fish in real time, and help us estimate abundance of different stocks. This is

"What should we do to address the large errors in pre-season forecasts?"

Table 5: Implications of large errors in pre-season forecasts—what should we do?

1. Improve:
   - In-season monitoring
   - Updating of forecasts in-season
   - Linking in-season decisions to those updated forecasts (taking uncertainties into account)

   Work is already ongoing
2. Increase monitoring of ocean environment (satellites, at-sea sampling, tagging, ...)
3. Conduct more research on links between ocean and salmon survival rates.
4. Reduce loss of fish (harvesting, en-route mortality, ...).
5. Consider comparisons between weather forecasting and pre-season forecasting of salmon abundance.
6. Reduce expectations about accuracy of pre-season forecasts of salmon abundance!

   Uncertainties are large! This essage is for everyone:
   - Users
   - Managers
   - The public
   - The media

an important aspect for dealing with large uncertainty from pre-season forecasts. The way to think about this is to go back to Figure 5 — where in-season adjustments can be applied to help meet the target. We need to make adjustments, to change the shape of the lens, so that we get the projection for the trajectory closer to the target. This is no small task.

The second and third suggestions in Table 3 are related. We need to increase monitoring of the ocean environment. Since 1997 there has been a very important satellite in operation, the SeaWifs satellite. It looks at sea-surface chlorophyll levels, a measure of primary productivity. Information about other aspects of the ocean is becoming available through radar sensing. This means that we can now get a much better picture of the dynamic nature of the ocean over a broad scale and numerous locations. However, we need more 'at sea' sampling — yet this is an extremely expensive process. The major projects that have led to our current understanding of migration of salmon in the ocean were carried out in the 1960s. Tagging studies also are needed. Given the new information, we now need to conduct more research on the links between ocean conditions and salmon survival rates to try to reduce the uncertainties that are very large.

"We need to increase monitoring of the ocean environment."

The fourth suggestion is to try to reduce the loss of fish through harvesting and en-route mortality. There are large uncertainties in pre-season forecasts, so that means we should manage in a way that reflects that we could have either overestimated or underestimate abundance. Being wrong on the high side usually means that we end up with too few spawners, which raises conservation concerns and reduces the chance of having long-term productive salmon stocks, which is not in the interest of anyone.

"There are large uncertainties in pre-season forecasts, so that means we should manage in a way that reflects that we could have either overestimated or underestimate abundance."

Let us consider the comparisons between weather forecasting and pre-season forecasting of salmon abundance (Figure 15). If you look at the statistics, weather forecasting is actually very good on a short-term basis. One-day forecasts have extremely high quality, and five-day forecasts are very good and much better than they used to be. However, if we extend the time over which we are making weather forecasts, then reliability goes down. Note that weather forecasts are made on the basis of thousands or tens of thousands of observations. There are satellite data, numerous ground stations, and huge computer models evaluating these data and making predictions one or five days in advance. When you look at the data sources for making salmon, forecasts if it was drawn accurately you wouldn't have been able to see the right hand bar in Figure 15, it would be so small in comparison to the thousands of observations for

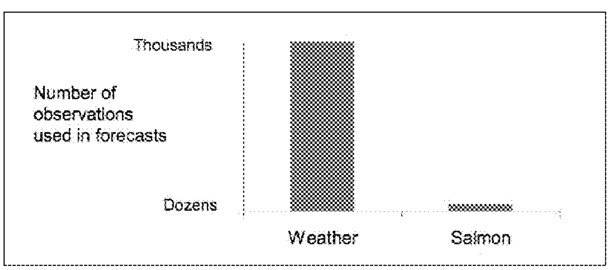

Figure 15: Comparison of number of observations used in forecasts for weather and pre-season forecasts for salmon.

weather forecasts (literally in the order of dozens of data points). Maybe we should think of it this way — we are seeing juveniles disappear into the black box of the ocean and they are facing some unknown survival rate to come up with the adult abundance or recruits for which we are expected to provide forecasts. This is a daunting process. If we think about it in the context of weather forecasting, then we should be amazed that we do as well as we do with pre-season forecasts. At the very least, we should expect that forecasts are going to be highly uncertain.

The sixth and last suggestion in Table 3 is that we need to reduce expectations about the accuracy of pre-season forecasts. Salmon abundance is a highly variable quantity, uncertainties are huge, and everybody needs to recognize this. Users of the resource need to take that into account in their planning. Managers need to take that into account — they already know it well. The public and especially the media have to expect that these forecasts are not going to be perfect.

"... we need to reduce expectations about the accuracy of pre-season forecasts. Salmon abundance is a highly variable quantity, uncertainties are huge, and everybody needs to recognize this."

"The public and especially the media have to expect that these forecasts are not going to be perfect."

WHY CAN'T WE GET IT RIGHT?

Brian Riddell, Division Head, Salmon and Freshwater Ecosystems Science Branch, Pacific Biological, Fisheries and Oceans Canada

Our problems today are certainly not for a lack of awareness of the people involved in the assessments at the Pacific Salmon Commission, not for a lack of data crunching and data volume and certainly not for a lack of effort. That is what drew me to ask this question: With so much effort, why can't we get this right?

Randall Peterman has provided an excellent introduction to one of the primary problems that we face; that is, the level of uncertainty of our information, after all our efforts, is far greater now than it has been in the past. Part of this reflects the change in fishing. In the past, fishing provided important income for people on the coast and also large signals of information, which was subsequently used in management. With reduced fishing effort we have lost a large volume of information that could be used in getting information out the next day, or the next week. We now have to accept that uncertainty is a fact of life. The problem is first that the uncertainty is unpredictable, and second, it is growing.

Fraser stock assessment has been going on since the beginning of the Pacific Salmon Commission, formerly the Pacific Salmon Fisheries Commission, and data have been collected since 1948. It is anything but a simple analysis of trend, however; it is extremely complicated and in fact has been quoted as one of the most complex resource management issues in the world.

"The problem is first that the uncertainty is unpredictable, and second, it is growing."

Getting it right

What do we mean by 'getting it right'? Being 'right' really reflects further on the expectations. It is probably associated with the question: Do we have adequate planning? I believe that our recent efforts in planning are actually quite good. Once we do get the fish back to the coast then we have pre-season plans with respect to the allocation of production — how