# Forecasting the Impact of Policy Interventions on Societal Behavior

**Brian Armstrong, Charles Day**
**Charles F. Day & Associates**
**West Lafayette, Indiana**
**barmstrong@cfday.net, cfday@cfday.net**

**Alok Chaturvedi**
**Purdue University**
**West Lafayette, Indiana**
**alok@purdue.edu**

## ABSTRACT

To overcome the challenge of heterogeneity in data sources, format, and content, disciplinary theories, and modeling frameworks we describe a methodology that partitions all data into individuals, organizations, institutions, infrastructure, and geographies entities (IOIIGs) and automatically extracts relevant theories for each entity class. We then reconstruct a virtual Reference World (RW) in which IOIIG "agents" interact with each other to produce emergent societal level behavior. Policies are defined in terms of actions by some Institution and Organization agents to shape the behaviors of Individual agents so that desired societal level outcome may be achieved. The RW approach consists of four phases: simultaneous self-ingestion of data with continuous verification and validation, self-extraction of behaviors, self-extraction of policies, and simulation of future states.

RW methodology continuously ingests dynamic data through an onboarding process which relies on semantic definitions of fields. Static traits and dynamic actions of IOIIGs are extracted as sampled data to form behavior maps for the various modeled entities. Data describing intended and observed impacts of policy actions are converted into action meshes. The behavior maps and action meshes are then combined using neural network and machine learning techniques to create plausible future meshes. Agent based simulation is then used to produce future states.

We demonstrate the application of the RW approach to supporting strategic policy design for increasing the post-secondary attainment of the work eligible population nation-wide.

## ABOUT THE AUTHORS

**Brian Armstrong** is a lead developer and researcher in Charles F. Day & Associates. He is an expert in building scalable, integrated frameworks for big data analytics. He is actively involved in designing cloud-based platforms that provide the power of automatic semantic extraction, network inference, and simulation-based forecasting to analysts as dynamic and interactive user interfaces. Dr. Armstrong received a doctoral degree in Computational Science and Engineering from Purdue University in 2010 with a research focus on distributed computing, performance analysis of industrial-grade applications, and automatic high-level code transformations for parallel computing.

**Charles F. Day** is Founder President and CEO of Charles F. Day & Associates. The company consists of two lines of business, Federal Services, and Agent-Based Modeling & Simulation. Dr. Day graduated from Purdue University with Bachelor of Science degree in management. He served on active duty in the Army including as a Research & Development Manager. He completed an MBA at St. Ambrose University and a Doctor of Business Administration from Nova Southeastern University in Fort Lauderdale, Florida. He also has the Owner/President Manager Executive Education certificate from the Harvard Business School. He has been a guest lecturer at the National Defense University in business and organizational topics.

**Alok Chaturvedi** is a professor of Management at the Krannert School of Management at Purdue University. He has been working on Dynamic Data Driven Application System for over a decade and has published over 100 articles. Professor Chaturvedi led the development of Synthetic Environment for Analysis and Simulation (SEAS) platform, which was funded by NSF's successive DDDAS programs and has been field tested for over 15 months by ISAF in Afghanistan as well as used to develop strategies to reduce high school and college dropout rates, improve the well-being of under privileged citizens, and plan for disasters during major spectator events. Professor Chaturvedi was named in Federal 100 by Federal Computer Weekly and was awarded the "Sagamore of the Wabash" by the Governor of Indiana, the highest civilian award for his service to the State.

# Forecasting the Impact of Policy Interventions on Societal Behavior

**Brian Armstrong, Charles Day**
**Charles F. Day & Associates**
**West Lafayette, Indiana**
**barmstrong@cfday.net, cfday@cfday.net**

**Alok Chaturvedi**
**Purdue University**
**West Lafayette, Indiana**
**alok@purdue.edu**

## THE CHALLENGE OF DESIGNING POLICIES THAT MATTER

Policy makers seek to design interventions that will be effective in encouraging positive changes in a population's behavior. Populations are diverse with many different goals among different people, making many policies ineffective because people do not respond even when opportunities are offered, have counterintuitive impact because helping one segment of population oftentimes creates a greater inequality with other segments, and have trouble sustaining an impact over time, which is especially problematic when policies are out of date by the time they are adopted.

Committing to policies such as national policies to increase education attainment and eliminate racial inequality is a large financial investment and is typically ineffective in the long term due to biased intuition and the inability to foresee the counteraction of other trends in society. The desire of this work is to build a data-driven approach to policy design leveraging models that are not inherently biased and to provide this platform as an interactive system which allows users to quantitatively compare future outcomes and sustainability of arbitrary policies before committing to a course of action.

Building such a data-driven platform is challenging because the required data involves multiple sources which are inherently heterogeneous due to the fact that policies involve many perspectives from diverse fields. And, if the platform is to be relevant for the time frame of developing policies, it must remain flexible and maintainable to allow new ideas to be easily incorporated.

We describe a novel platform, Reference World (RW), to address the challenges in bringing policy design into the modeling and simulation arena. RW consists of a unique combination of machine learning techniques, agent based modeling, and neural networks accompanied by a user interface for policy makers to perform "what if" experiments.

RW first integrates data to a generic and fundamental model through semantic relationships. Then, policy actions are represented in RW using a data-driven method that captures desired impacts of arbitrary interventions. Finally, RW provides the user with a simulation that allows forecasting future outcomes of status quo or user-designed policies.

### Summary of Contributions

- We will describe the RW platform applied to enable policy design, analysis, and forecasting.
    - The flexibility required to model and forecast the impact of policies is enabled using agent based modeling (ABM).
    - RW uses deep learning and convolutional neural network techniques to extract "features" and trends among the behavior of a heterogenous population.
- We will also share a use case demonstrating how RW is used to develop a complex policy addressing post-secondary attainment, well-being, and racial inequality.

## RELATED WORK

Analyzing and predicting the impact of policies on society-level phenomena, such a post-secondary education attainment among working class Americans, is of vital interest to policy developers because of the potential these impacts have in enabling the United States (U.S.) economy to be successful and the cost in terms of time as well as monetary resources involved in implementing such policies. Yet, understanding the impact of policies on society is challenging due to diversity of demographics, as is noted by McLendon and Perna (2014). Building upon the prior work, Chaturvedi et. al (2011, 2013, 2014) RW, is designed with such challenges in mind—to capture the dynamic behavior among diverse demographics and forecast the sustainability of policies.

Garcia and Bayer (2005) take a closer look at Latino groups with respect to post-secondary attainment in the U.S. They argue that viewing behavior at aggregate levels, divided by racial and ethnic traits alone, fails to capture the reasons for variations in behavior of subgroups within the Latino population. Their work focuses on key factors that distinguish behavior within the Latino population: programs within a school, school location (urban/rural), school type, parental education, nativity, family structure, gender, race, and academic performance. RW captures the factors relating to characteristics of people as static and dynamic traits of individual agents. The characteristics of schools are captured in the way the actions are applied in the simulation (described in the Reference World section below) since institutions are one of the actors in the RW system.

The way that RW represents the diversity of a population in an efficient manner is using agent based modeling (ABM). ABM has been applied to modeling emergent human behavior, as Wu and Sekiguchi (2020) describe. Their work is focused on representing relationships among groups. RW captures the indirect influence among individuals through the use of neural networks (described in the Reference World section below). As probabilities of certain paths within the neural network increase, the probabilities of other paths are decreased. Additionally, the changes to probabilities in subsequent years are impacted by previous years with a dampening function, allowing both an "inertia" of change to occur, similar to the "conflict inertia" that Wu and Sekiguchi represent, as well as emphasizing more recent changes over past changes.

Gu and Blackmore (2015) review the application of ABM to modeling student behavior in higher education. They make use of ABM because it is best suited to modeling the "unpredictable and dynamically changing" phenomena of student enrollment. RW relies on ABM because the choice of individuals to obtain post-secondary attainment is complex, involving more than just access to education and availability of tuition funding, such as whether employed individuals who support a family have the time to take advantage of the opportunities or transportation to attend classes.

Neural networks have been used successfully in many application areas, such as pattern recognition in images. RW relies on neural networks to extract features from Behavior Maps, which can be represented as a sort of image. Duan and Li (2018) detail how a convolutional neural network can be used in combination with classification techniques to recognize age and gender of people in images. Wen and Shi (2018) describe how they use deep learning neural networks to build a model for brain responses to images. Sullivan and Winsnes (2018) define how they train neural networks to be accurate and fully automated in classifying images using a human classification effort. RW techniques from convolutional neural networks and deep learning to identify features and higher level features.

Tavanaei and Ghodrati (2019) focus on more precisely representing actual brain function using spiking neural networks. The neural networks in RW (i.e., the Behavior Maps) can be loosely thought of as a brain that is trained with repeated stimuli, where each stimuli is extracted from an individual's census data record. A Behavior Map is then used to "recall" static traits given dynamic outcomes or vice versa (predict outcomes given partial traits).

## THE ATTAINMENT EXAMPLE

Throughout this paper, we will refer to the "Attainment Example", which is a real-world implementation of the RW platform for a client involved in policy development in the education and workforce development domain. We use this platform to design a strategy with the goals of increasing the percentage of people who have post-secondary attainment (referred to as "attainment" in this paper), improve people's overall well-being, and reduce the racial discrepancies in attainment and well-being. Attainment is defined as any post-secondary education degrees, certificates, or licenses offered by accredited education institutions.

## DEFINITIONS

- **entity:** a representation of an individual person, organization, institution, infrastructure, or geographical region;
- **agent:** an entity that senses its environment and responds due to autonomously making decisions;
- **trait:** a characteristic of an entity, such as the ethnicity, sex, or age of a person; traits can be static or dynamic. Changes in static traits are not of interest to the specific application, such as change in marital status for the attainment example. Traits which may be impacted by policy actions are grouped into dynamic traits, such as education achievement or salary.
- **sampled data:** data represented as snapshots in time as opposed to longitudinal data which tracks observables of specific people over time; data provided as the basis for the attainment example is sampled data from annual Census Bureau surveys.
- **policy action:** an intervention consisting of a desired outcome for a targeted subset of the population; RW represents a policy as a set of actions, each of which has a start and end date and a percentage of the policy's resources. An example policy action is: "Employers provide free online courses to employees for up to 25% of the credits required to obtain an Associate's degree." The amount of resources applied to this action will determine how many employees it is offered to.
- **well-being:** a combination of how well people perceive their lives to be progressing and what they wish their lives to be; for the purpose of modeling behavior in RW, well-being is defined as a combination of basic needs, education, financial, health, freedom of movement, social, security, political, and religious well-being. Consider financial well-being in the context of the attainment example: a female employee's financial well-being could be represented as the gap between her salary and the average salary of her peers.

## REFERENCE WORLD

### Key Features

- **Creates a Continuous Flow from Data to Model.** Semantic relationships are used to determine the data flow from fields of diverse data sources to static and dynamic traits of modeled entities in a virtual representation of society called Reference World (RW).
- **Based on an Underlying Behavior Model.** The behavior of entities in RW is modeled based off of well-being theory by Daniel Kahneman. The behavior of society emerges as its constituents interact with their environment and each other.
- **Incorporates Intervention Actions to Enable Interactive Policy Design**. Arbitrary policy actions are represented as desired outcomes on targeted subsets of the population. Subsets of the population are defined using selections of static traits of RW entities.
- **Enables Quantitative Comparison of Future Policy Outcomes and Trends**. RW applies status quo actions to create a baseline and allows users to create simulations of scenarios with their own selection of actions for comparison.

### Four Phases from Data Ingestion to Forecasting the Future

RW consists of the following four phases which integrate data, models, user-selected actions, and simulation.

1. **Data Onboarding:** simultaneous self-integration of data using semantic relationships;
2. **Constructing Behavior Maps:** self-extraction of behaviors using neural networks;

3. **Computing the Impact of Actions:** self-extraction of policies through a data-driven process;

4. **Simulate Future States:** forecast changes in behavior by applying machine learning techniques.

**Data Onboarding**

Data Onboarding consists of parsing data sets to extract fields and linking the fields to static and dynamic traits of individual agents using semantics, as illustrated in Figure 1. In the attainment example, annual Census Bureau surveys provide samples of a national population with such fields as age, race, sex, citizenship, highest education level, salary, and occupation. Individual agents have static traits, such as age, race, sex, citizenship, and dynamic traits, such as highest education level, salary, and occupation.

Once the data sources have been tied to semantics of the Individual Agent, an Agent Population dataset, illustrated in Figure 2, is created. The Agent Population is distributed into groups defined by unique sets of static traits. One Agent group in the Attainment Example is the African American male between the ages of 25 and 34, a widowed parent of one child, living in a suburban region, is a U.S. citizen who lives in a household that makes $65,000-75,000 annual income.

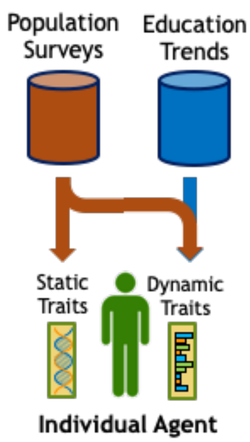


**Figure 1. Semantic Tie between Data and Model.** Data fields are related to traits of the Individual Agent, which represents a person in RW, through semantic definitions.

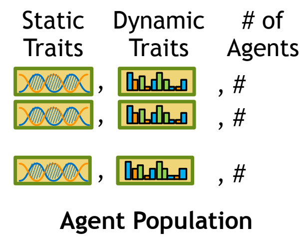


**Figure 2. The Agent Population.** An Agent Population consists of dynamic traits and a number of agents for each unique set of static traits.

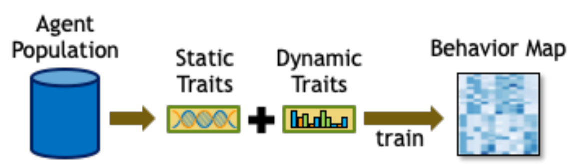


**Figure 3. Constructing a Behavior Map.** Behavior Maps are neural networks trained using the traits of the Agent Population.

**Constructing Behavior Maps**

RW relies on neural network techniques to capture the relationships between the static and dynamic traits and the rate at which these relationships change year by year. The traits of the Agent Population are used to train neural networks, referred to as Behavior Maps, illustrated in Figure 3.

There will be at least one Behavior Map per year. There is a trade off in creating more Behavior Maps per year to capture correlations specific to one trait, such as race: creating a Behavior Map per race causes the resulting behavior of minorities to be unaffected by the data for the majority race but loses the relationship one racial group's behavior has on another. In the Attainment Example, we created one Behavior Map per state per year.

**Computing the Impact of Actions**

Intervention actions are represented in the RW model as desired outcomes on distributions of the population. In the Attainment Example, one intervention action is to provide "reverse transfer credit," which is an action education institutions take to allow students who are considering dropping out of a 4-year Bachelor degree program to use completed credits towards obtaining an Associate's degree, certificate, or license. Data from education institutions reveal the percentage of students in 4-year institutions who complete one or two years of studies and then drop out of college. This data is used as a basis to build an Action Transformation with the desired outcome of making the education attainment levels of Associate's, certificates, and licenses more likely and making college drop outs less likely. The distribution of the population for this action is based on the enrollment distribution in 4-year institutions and the percentage of enrollees who typically drop out after one or two years.

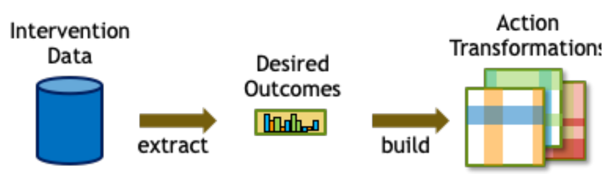


**Figure 4. Computing Transformations for Actions.** Action Transformations are tied to data describing the desired outcomes of the action in terms of increasing or decreasing the probability of various dynamic traits and the distribution of the population the action affects.

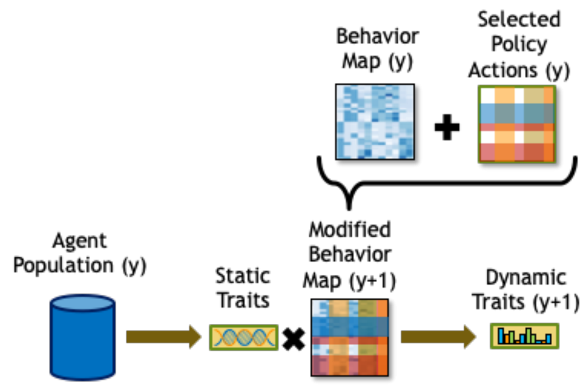


**Figure 5. Projecting Dynamic Traits Forward in Time.**

**Simulate Future States**

RW uses machine learning techniques to combine the Behavior Maps for year y and Action Transformations representing actions applied throughout year y to create the Behavior Map for year y+1. The process of computing the Behavior Maps and the dynamic traits for subsequent years is illustrated in Figure 5.

The static traits of the agent population do not change year by year, meaning the static traits for the Agent Population in year *y+1* will be the same as those in year *y*. Only the dynamic traits of agents change. For instance, in the Attainment Example, the Agent Population has data for 34-year-old Asian females in both year *y* and year *y+1*. The probability that 34-year-old Asian females have obtained Bachelor's degrees may be different in year *y+1* than in year *y* due to changes in behavior among this population segment.

The Behavior Map for year *y+1* is combined with the static traits of agents from the Agent Population to compute the Dynamic Traits for each unique set of static traits.

To capture the change due to population growth, the sizes of the Agent Population segments with corresponding static traits is modified using Population Projections. The population projections by age, race, and sex from the Census Bureau are used in the Attainment Example.

The projected agent population for year *y+1* is gathered using the static traits for a segment of the Agent Population, dynamic traits for year *y+1*, and change in the number of agents with the given set of static traits. This is illustrated in Figure 6.

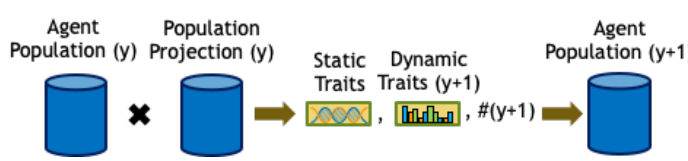


**Figure 6. Project the Agent Population Forward in Time.**

## EXEMPLARY RESULTS FOR INCREASING EDUCATION ATTAINMENT

We use a commercial implementation of RW applied to the real world challenge described as the Attainment Example to illustrate insights that can be gained using the modeling and simulation approach described earlier. For this application, the primary metrics of interest are post-secondary education attainment, overall well-being, and racial inequality—the goal is to identify the strategies that improve overall well-being and increase post-secondary education attainment in work eligible people between the ages of 17 and 64 across the nation. The strategies to be selected as best-in-class must reduce racial inequality and be cost effective.

### Baseline Forecast

The Baseline Forecast scenario (referred to as Baseline) is provided as a reference point. It represents the impact of continuing to apply the policies that were in place in 2017 forward until 2025. Baseline's interventions include the average tuition subsidy for students through financial aid, trends in employer benefits involving subsidizing tuition costs, and the use of traditional and social media to attract new students and inform people of the available education opportunities in 2018. Tuition funding for Baseline is propagated forward in time by increasing the amount by the nation's projected inflation rate.

### Mixed Approach Strategy

In this example, we will target employed adults (25 to 64 years old) who are currently employed in military and information occupations and industries and who likely end up with a maximum education level of a high school diploma (or equivalent), certificate or license, or an Associate's degree. The target population includes people who dropped out of college without any post-secondary credentials.

We will incrementally construct a strategy and analyze the strategy using multi-factor causal inference. A cause-and-effect relation occurs when we can infer that a particular intervention or policy (cause) leads to specific changes in the metrics of interest (effect). Since this scenario involves a mixed approach that includes multiple types of actions, we will create three scenarios to show the impact of (1) tuition funding, (2) actions involving employment activity and social media, and (3) the combination of parts (1) and (2).

#### *Intervention 1: Traditional Tuition Funding*

The scenario named "2-Year Funding for Military and Information Fields" includes actions by institutions to increase funding for adult students enrolling in 2-year institutions from \$8,200/student to \$9,100/student to offset other fees, such as costs for books, room and board, transportation, or other needs, such as child care.

#### *Intervention 2: Integrated Employment Activity and Social Media*

The scenario named "2-Year Credits for Work Experience for Military and Information Fields" includes actions by employers to offer programs where employees can obtain academic credits for specific work experience which can be transferred towards a certificate, license, or an Associate's degree from 2-year institutions. The amount of credits is capped at 30 semester credits, which is approximately half of what is required for an Associate's degree.

In this scenario, we also apply an increased use of social media to inform candidate students of the opportunity to obtain transferable credits for work experience.

#### *Intervention 3: Mixed Tuition and Employment Activity Policy*

The scenario named "2-Year Institution and Employer Funding for Military and Information Fields" includes the increased tuition funding by institutions as well as employer supported programs where transferable credits can be obtained for work experience. This scenario includes increased social media informing candidate students of both the increased tuition funding and employer programs.

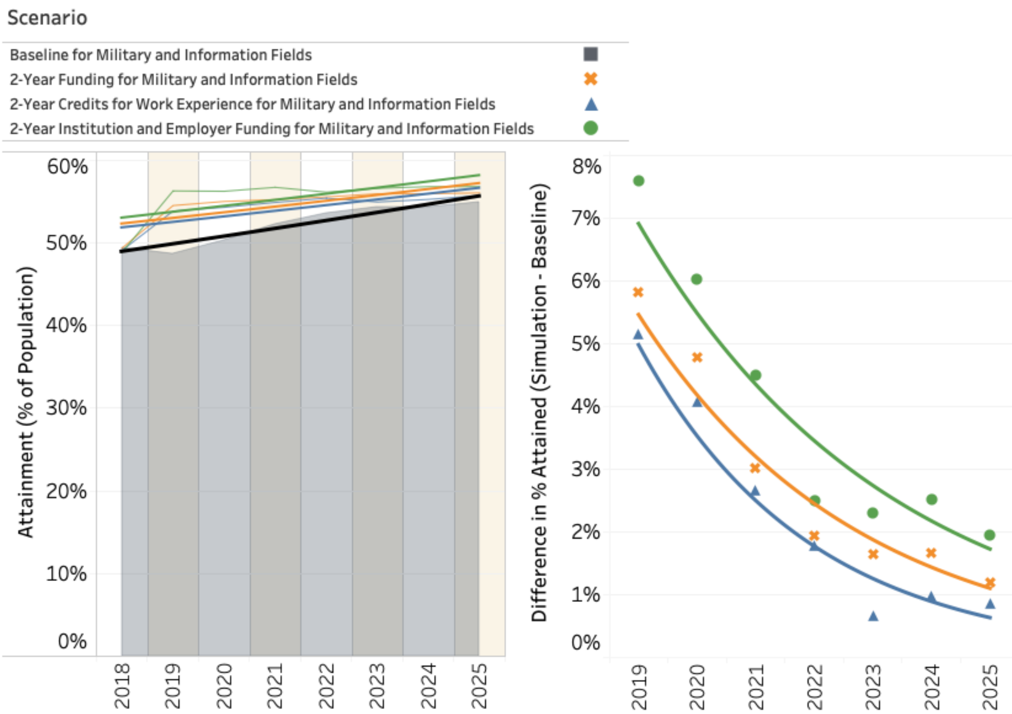


**Figure 7. Absolute Attainment and Attainment over Baseline.**

The left chart shows the attainment of the three intervention scenarios and Baseline. The darker lines show the trends and the lighter lines show the data points. The right chart is each intervention's attainment minus Baseline's attainment, showing the impact of the interventions over the status quo policies projected forward in time.

**Results Analysis**

The charts in Figure 7 shows that all interventions result in an initial jump in attainment as adults take advantage of the opportunities to obtain post-secondary credentials, but the increased attainment converges to Baseline over time. We can also see that Intervention 1 and 2 have very similar impacts. Intervention 3 has an increase over each of its parts, indicating that Interventions 1 and 2 are influencing different individuals.

A closer look at the education levels in Figure 8 shows that the combined intervention causes an increase in employed adults obtaining certificates and a decrease in people who end up with some college but no degree, a high school or equivalent diploma, or no high school diploma. The Credits for Work Experience intervention has a slight decrease in those who obtain an Associate's degree, perhaps because it is more attractive to some to obtain a certificate or license that is practically paid for by their employer.

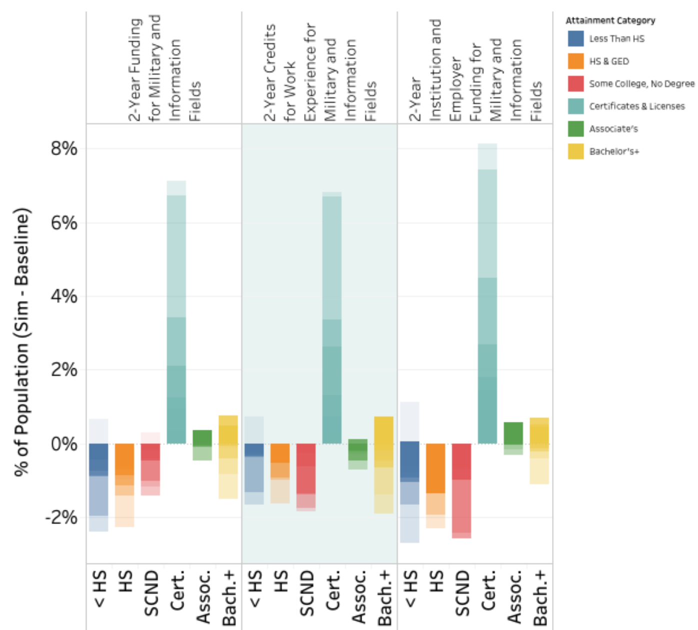


**Figure 8. Difference with Baseline by Education Level.**

This chart shows the difference of each intervention with respect to Baseline. The bars are more opaque as time progresses from 2018 to 2025. A positive value means there are more people in an education level due to an intervention than are in Baseline.

Another metric of interest is the overall differences in attainment by race. Ideally, the gap will be narrowed by the interventions to a greater degree than what is observed in the status quo policies. Figure 9 shows that all three scenarios reduce the inequality among races in terms of post-secondary attainment. The two components of combined interventions scenario work in an additive way, reducing the racial inequality by a greater degree than each of the components alone.

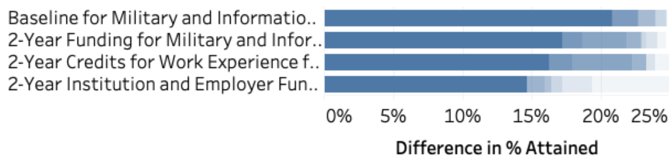


**Figure 9. Racial Inequality by Scenario.**
The maximum – minimum attainment across races is given per scenario. The bars become more opaque as the years progress from 2018 to 2025.

**Table 1. Underlying Data for Observed Well-Being.**

| Component | Data |
|---|---|
| education | education level |
| finance | salary |
| health | who provides health insurance (no health insurance, self, third party) |
| freedom of movement | urban/rural locale & number of vehicles |

Comparing attainment to well-being reveals how the interventions are impacting other aspects of individuals' lives, including finance, health, and freedom of movement. The data used to measure well-being is normalized across all racial groups and years. Table 1 identifies the underlying data used to quantify an individual's well-being. An increase or decrease in the well-being metrics indicates a significant change along one or more of the dimensions in Table 1.

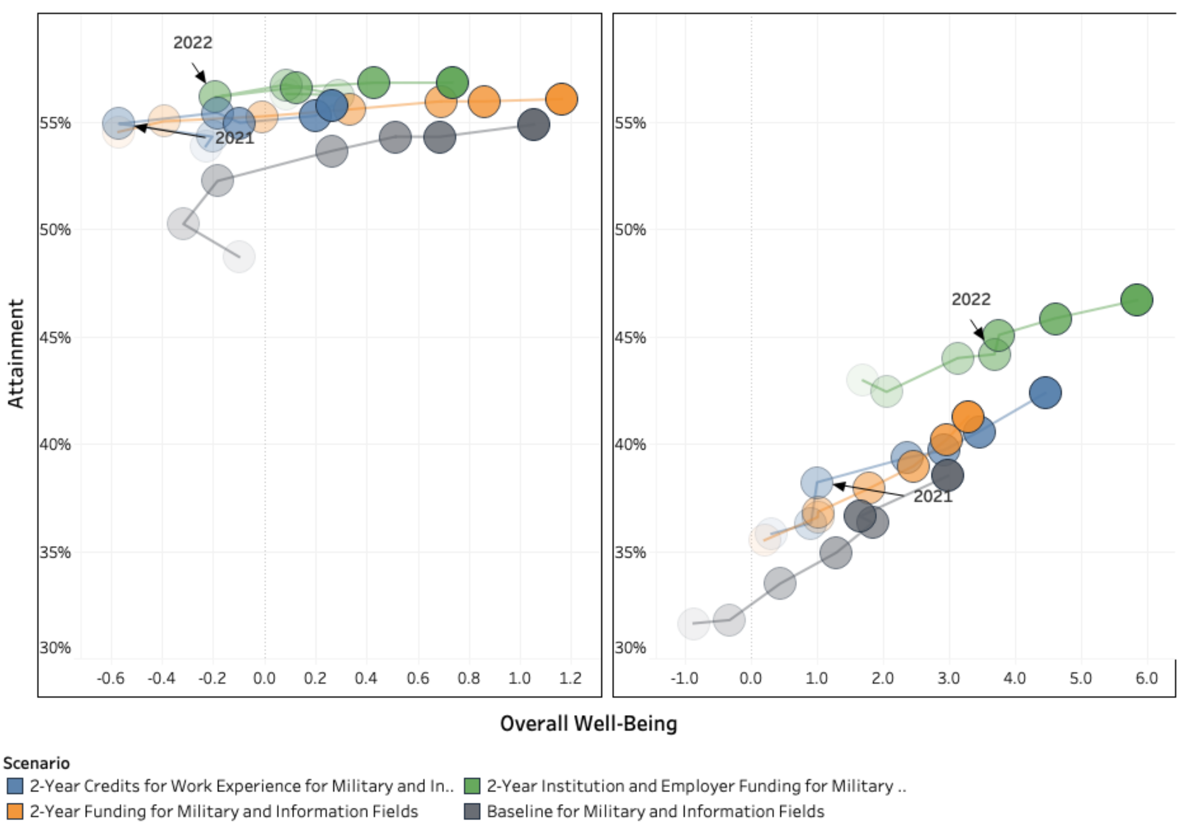


**Figure 10. Attainment vs. Well-Being for the Entire Population (left) and Only Minorities (right).**
The three interventions and Baseline are plotted with attainment as the y-axis and well-being as the x-axis. Their paths over time are shown with the circles becoming more opaque as we go forward in time from 2019 to 2025. The absolute value of well-being is not meaningful, but the relative value from one year to the next or between scenarios is significant. The graph on the right focuses only on minority races.

Figure 10, which compares attainment to overall well-being, shows that there is a general increase in both attainment and well-being over time, but there are different reactions to the various interventions as seen in the well-being. Focusing first on the whole population, we see that the employer action to provide credits for work experience (blue) increases attainment in the overall population, but has an initial decrease in well-being until 2021. This negative trend indicates that though the intervention is providing opportunities for education, it is not effective at addressing other pressing needs in the target individuals. The funding scenario (orange) does not have this decrease. The combined intervention (green) initially follows the trend of the second intervention where well-being decreases as attainment increases, but then increases after 2022.

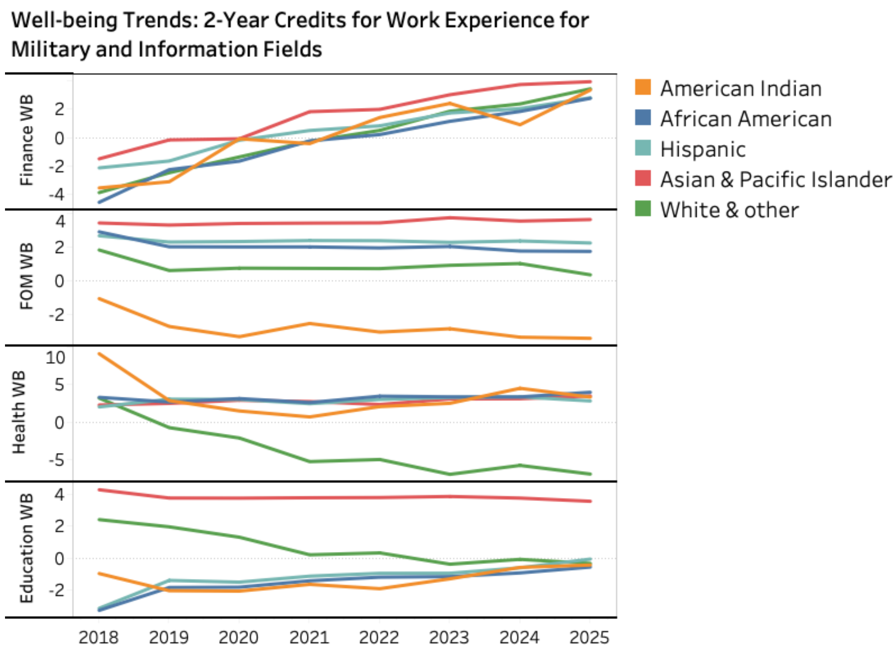


**Figure 11. A Closer Look at the Impact of the Credits for Work Intervention on Well-Being.** This chart shows the trends in well-being over time for each of the well-being components. Each line is one racial group. A drop in well-being indicates a developing need, such as the drop in health well-being for the White and other population indicates less health coverage is being provided. The trends in the education well-being indicate that the observed gap among racial groups is being reduced.

The chart on the right of  focuses on Hispanic, African American, and American Indian racial groups. The impact of the interventions on this population segment is significantly different than the general population. We do not see the negative trend in well-being and the mixed approach involving both employer and institution actions has a greater impact on both well-being and attainment than the combination of the other interventions (Credits for Work and Institution Funding.) The conclusion is that the mixed scenario is addressing roadblocks the minority groups face in education attainment and bridging the racial discrepancy gap in overall well-being.

To gain better insight to why the Credits for Work intervention has a negative impact on well-being, Figure 11 shows the trends in each of the well-being components by population segment. The most significant changes in well-being occur in health well-being among the general population and the impact on education well-being for all but Asians and Pacific Islanders. The negative impact on health indicates that the general population has less health coverage when taking the route of obtaining credentials through a credits for work program than they would otherwise. The negative trend for the general population in education well-being is because more individuals settle with a lower form of attainment when obtaining credits for work instead of going through traditional 2-year or 4-year degree programs.

However, the credits for work provides the ability for minorities to improve their education well-being over what would normally be possible. The different reaction to the mixed intervention among the racial groups is also an indication that the most effective strategies will involve a multi-faceted approach that micro-targets population segments to address the unique needs of each.

## Conclusions

RW has proven to be a flexible platform in which diverse data sources can be integrated to model the dynamic and complex behavior of a population. The generic modeling framework and accepted machine learning and neural network techniques enable analysts to view the virtual population of RW from multiple perspectives, focusing on a variety of traits and metrics. We demonstrated how policies can be built incrementally and analyzed using multi-factor causal inference to identify insights into policy design.

## Direction for the Future

The example application of RW described in this paper has many directions that will be interesting to pursue, including creating models of individual education institutions and analyzing institution behavior to capture their diversity and creativity in attracting and retaining students. We will also expand this scenario to include models of employers with applications directed at developing interventions for job market.

Because RW is based on ABM, we are also investigating building an app for individuals to use to forecast the impact of education and employment choices and teaming with education institutions and job-search related organizations to offer guidance on how individuals can reach their goals.

## REFERENCES

Chaturvedi, A,. Chaturvedi, R., & Armstrong, B. (2014) "Securing the Food Supply Chain: Understanding Complex Interdependence through Agent-Based Simulation." *Health and Technology*: Vol. 4, Issue 2, pp 159-169.

Chaturvedi, A. R., Dolk, D. R, & Drnevich, P. L. (2011) "Design Principles for Virtual Worlds," *MIS Quarterly*, Vol. 35, No. 3, pp 673-684.

Chaturvedi, R., Armstrong, B., Chaturvedi, A., Dolk, D., & Drnevich P. (2013). "Got a Problem? Agent-based modeling becomes mainstream," *Global Economics and Management Review*, Vol 18:2, pp 33-39.

Duan, M., Li, K., Yang, C., & Li, K. (2018). A hybrid deep learning CNN–ELM for age and gender classification. *Neurocomputing*, *275*, 448-461.

Garcia, L. M., & Bayer, A. E. (2005). Variations between Latino groups in US post-secondary educational attainment. *Research in Higher Education*, *46*(5), 511-533.

Gu, X., & Blackmore, K. L. (2015). A systematic review of agent-based modelling and simulation applications in the higher education domain. *Higher Education Research & Development*, *34*(5), 883-898.

Kahneman, D. (1999). Objective Happiness. *Well Being: The Foundations of Hedonic Psychology*, ed. Kahneman, Diener, & Schwarz, New York: Russell Sage Foundation.

McLendon, M. K., & Perna, L. W. (2014). State policies and higher education attainment. *The ANNALS of the American Academy of Political and Social Science*, *655*(1), 6-15.

Sullivan, D. P., Winsnes, C. F., Åkesson, L., Hjelmare, M., Wiking, M., Schutten, R., Campbell, L., Leifsson, H., Rhodes, S., Nordgren, A., Smith, K., Revaz, B., Finnbogason, B., Szantner, A., & Lundberg, E. (2018). Deep learning is combined with massive-scale citizen science to improve large-scale image classification. *Nature biotechnology*, *36*(9), 820-828.

Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., & Maida, A. (2019). Deep learning in spiking neural networks. *Neural Networks*, *111*, 47-63.

Wen, H., Shi, J., Chen, W., & Liu, Z. (2018). Transferring and generalizing deep-learning-based neural encoding models across subjects. *NeuroImage*, *176*, 152-163.